Date: Thu Jul 01, 1993 6:59 am PST
Subject: Re: Generalizing; statistics vs generative models

From Tom Bourbon [930701.0912]

>[From Bill Powers (930629.1845 MDT)]
>
>Gary Cziko (930629.1935 UTC),
>Tom Bourbon (930629.1551)
>
>>IT IS IMPOSSIBLE TO DERIVE GENERAL PRINCIPLES FROM SINGLE
>>EXAMPLES.
>
>What an event! Theoretical biology commits suicide before our
>very eyes! For if it is impossible to derive general principles
>from single examples, then it is altogether impossible to derive
>general principles. There is only one universe.
>...

>Many scientists seem to confuse this search for general rules
>with the search for mechanisms underlying phenomena. Some, of
>course, are perfectly clear about the difference, and come down
>solidly on the side of generative modeling. Tom, can you pick out
>that citation (I think it was in "Worlds") and post it again --
>the one about how the laws of motion would have been treated
>under the strategy of generalization?

I believe the item you want was posted by Gary Cziko, probably over a year ago. It was in a post in which he said he had found something by "the other Taylor," to distinguish the source from Martin. It was a nice selection on how futile it would be to study group statistics on planets as a way to predict the motions pf any particulart planet, or to discover "laws" of motion. I have printouts of the post, but not immediately at hand.

Gary, can you remember the post and the source?

Until later, Tom

Date: Thu Jul 01, 1993 9:27 am PST
Subject: Conrolling Via Language; PCT & Other Theories

[From Hank Folson (930701)]

Has anyone done a diagram showing how a person controls through the medium of language/writing? Aren't the controlling actions parallel, multileveled, and indirect, unlike for controlling direct actions such as controlling the position of the knot in a rubber band by pulling on it, or moving a joy stick to control a dot on a screen?

At one level, the control action is perceived as performed when I hear myself say the words, or see them on the computer monitor. But I am not just controlling to hear myself say something. My primary want is for the other person to respond to my communication by doing what I am really interested in, which is for him to do what I think I communicated to him.

Doesn't a complete diagram have to include my understanding of the aspects of language relating to the communication and my perception of how the recipient perceives language?

If I do not perceive a reduction in the error signal, it is because the other person did not help me to achieve what I want, and I will continue to control. But there are two independent reasons why the other person did not do what I wanted: Either what I wanted was in conflict with the other's wants, OR, because the communication was faulty, the other person perceived something different than I intended. How do we diagram these independent options? How do we select which path to follow?

Controlling variables can become a very slow process for both parties because of communication problems. I wonder if applying the principles of PCT can help avoid communication problems, or if PCT can only make us more aware of the difficulties of controlling through the use of language.

Bill Powers (930623.0700 MDT) on PCT and other theories:

>What is it that's been taking up so much of our time?
>By and large, it's .... arguing with people who have other
>points of view to promote.

Isn't this to be expected? Doesn't PCT say that people will resist change because they are living control systems? Doesn't PCT also say that if they perceive that they are successfully controlling with and for their present theories, they will not only have no interest in a different point of view, but they will control to prevent this disturbance from affecting the variable (their favorite theory) that they are controlling?

>I am coming close to concluding that this is a waste of time.
>
>Does it really matter what (other theories) have to say about control
>theory? .....is that important to the development of PCT?

When I first joined the CSGnet, I thought it was. I liked the openness and willingness to compare and contrast PCT to other theories. It indicated a level of confidence, and a desire to really be sure PCT is correct. But I soon found that other theories were accepted at face value as having been proved to a greater extent than PCT, when in fact they appeared to be on much shakier theoretical ground and did not deserve the amount of time spent on them (and, as you note, on futile attempts to convert the believers). As an example, I recently saw an old interview with B.F. Skinner, whose thoughts have appeared here many times, on "educational" TV, showing pigeons 'reinforced' so they turned around and around to get food. Skinner's statements included the verbal equivalent of a lawyer's fine print: "We can make the pigeon turn around AS LONG AS IT IS HUNGRY." "The pigeon will stop turning when we do not give it food BECAUSE WE HAVE STOPPED PROVIDING THE REINFORCEMENT." For those new to PCT, our view is that the pigeon, trapped in Skinner's unnatural cage environment, has found that it must turn around to get food when it is hungry and wants to eat. So it is the pigeon that decides when it wants to turn around (to get food), and not B.F. Skinner. Skinner only created an environment in which turning around is the only controlling action that leads to food; He couldn't make a satiated pigeon do anything, which might explain why no applications of this research were discussed.

>The development of PCT rests on putting it to experimental test,
>discovering its flaws, and working out ways to correct the flaws by
>improving the theory.
>
>Trying to work out its relationship to all the other theories of
>behavioral organization that exist is a waste of time, because by
>the time PCT has developed enough to deal with the kinds of
>problems implied by these other theories, assuming it ever does,
>NONE OF THESE OTHER THEORIES WILL EXIST ANY LONGER.

I think that, while research must continue, we can justify putting a lot of effort putting into practice what we already know about PCT, as Ed Ford and others are already doing. In 350 years, no one has yet proved that Descartes was right. Communism was applied for 70 years before they decided it didn't work. PCT can do no worse than these two examples!
What do you think?

If PCT is basically correct, proper pragmatic applications should be much more successful than competing approaches, proving PCT's superiority at a practical level even though a lot of theoretical work remains to be done. This will attract more interest, which in turn should attract more researchers and more money for research.

>Why, I thought, are PCTers asked so often what PCT has to say about this
>fact or that theory or the other explanation of behavior?
>
>Obviously, because they don't understand PCT ...
>
>they deal with orthogonal subject matters...

> the next thing to do isn't to try to convince
>the listener of the truth of the answer. It's to ask "Do you want
>to know how I got to that answer?" If the reply is "no" the
>conversation is finished: all the person wants is for his
>conclusion to win.

>we shouldn't keep on pitting the conclusions of PCT against
>other people's conclusions.

A person can believe in the Theory of Relativity and in PCT. But, as you are clearly pointing out in your post, Bill, people can not understand and believe in PCT AND in another theory about living organisms, too. If PCT is right, the competition must be wrong. In the end, your theory is based on the physical structure of organisms, while all the other theories are based on somebody's interpretations of observations of behavior. There is no common ground. [It just occurs to me that the reason the other theories can grudgingly co-exist, and can not completely disprove one another is because they are all based on the same faulty premises, and so have too much in common.] All of this helps explain why it is so difficult to win over people who hold other beliefs.

>All of our problems come from trying to talk with people who
>don't understand PCT. The solution is to get them to understand
>PCT. The rest will take care of itself.

Sort of. People will come to PCT only when they are ready for change. For example, all of Ed Ford's excellent skills and knowledge and empathy are not enough for him or anyone else to succeed in reaching school teachers. Educators pick up on his PCT methods because, if they are caring teachers, they are in a state of reorganization from seeing all the problems that children and education (and in turn, teachers) have these days. Anyone who thinks (rightly or wrongly) that his controlling actions are helping him to reach his goals will not be open to change. Teachers are frustrated, and Ed Ford's ability to present PCT combines with their wants to produce change.

Summary: My feeling is that people will resist even a perfectly proven PCT for the reasons covered in our posts, so perfecting PCT is only a part of the answer. I think that the other part of the answer is to begin applying PCT where appropriate. Positive results of PCT approaches will then create more general interest in learning about PCT.

Hank Folson
Henry James Bicycles, Inc. 704 Elvira Avenue, Redondo Beach, CA 90277
310-540-1552 (Day & Eve.) MCI MAIL: 509-6370 Internet: 5096370@MCIMAIL.COM

Date: Thu Jul 01, 1993 12:53 pm PST
Subject: J. G. Taylor on Newton

[from Gary Cziko 930701.1940 UTC]

As requested by Bill Powers and Tom Bourbon. The following is taken from my _Educational Researher_ paper published last year.
==

It must be noted, however, that PCT models are designed to account for the behavior of individuals, and it is only by obtaining a considerable amount of data from individuals that these models can be tested and refined. Working within the framework of PCT makes it obvious that collecting and statistically processing a small amount of data from each individual of a large group can provide no understanding of individuals as purposeful systems (or any other type of system for that matter). J. G. Taylor (1958) made this point well:

"If Newton had had at his disposal not a vast amount of detailed information about a single solar system but a much smaller number of facts about each of a thousand solar systems, collected by a thousand observatories, he might conceivably have developed statistical methods for organizing this material. He might have found correlations between such variables as the number of planets in the system, the average number of satellites per planet, the average distance of the planets from the sun, and the like. He would, by this means, have learned a good deal about solar systems in general, but he could not have calculated the time and place of the next eclipse of the sun, and he could not have arrived at an understanding of the laws of planetary motion. He would have learned a lot about the ways in which solar systems differ from one another, but nothing about the ways in which any one of them works. For this it was necessary to know as much as possible about one system. Fortunately Newton had no alternative, and the result of his labours was the construction of a theory that survived until the advent of Einstein's theory of relativity." (p. 109)

Taylor provides another example to emphasize this point, an example which is more explicit and which may come closer to the type of research engaged in by quantitative educational researchers:

"Suppose that an investigator, knowing nothing about the construction of a motor car, decided to choose as his area of research the behaviour of the speedometer needle, and to this end took a series of readings in each of a hundred different models. Just to make the problem more like a real one we shall suppose that the speedometer dials are not provided with scales, but that the investigator can measure the angular deviation of the needle. Among the variables he might be expected to record are the distances of the accelerator and brake pedals from the floor, the position of the gear lever, the gradient of the road, the direction and velocity of the wind, and, of course, the speedometer reading. He takes a succession of simultaneous readings of all those variable in each car, and then proceeds to examine his data in the hope of solving the riddle of the speedometer needle. At first the material looks completely chaotic. There is no single independent variable that is functionally related to the dependent variable, and he decides to have recourse to statistical analysis. He finds negative correlations between the speedometer reading and (a) the distance of the accelerator pedal from the floor, and (b) the gradient of the road; and positive correlations with (c) the position of the gear lever, and (d) the distance of the brake pedal from the floor. He finds significant differences between the speedometer readings when the gear lever is in first, second, third and fourth positions, but the distributions overlap extensively. He now decides to record additional data, such as the weight of the car and its consumption of petrol, but the riddle remains unsolved. Of course we know the answer. If our investigator will only take independent measurements of the speed of the car he will find that in each system (car) the speedometer reading is a function of speed, but not necessarily the same function in all systems. He will find, moreover, that he can now dispense with statistical methods and can examine each system, considered as a matrix of pointer readings representing the several recorded variables, to determine how it hangs together. He will discover that what he at first took to be evidence of arbitrariness or caprice in his data was actually an artifact arising from the simultaneous examination of pointer readings taken from a hundred different systems. He will find that the same general principles apply to all the systems, but each of them has its own specific set of parameters." (pp. 110-111)

If for the speedometer reading we substitute some quantitative measure of an educational "outcome" (e.g., achievement test score, level of motivation, absenteeism) and for the various physical variables substitute educational "predictors" such as measures of social class, personality, and cognitive abilities, Taylor's comment can easily be seen to apply to quantitative, group-statistics-based educational research. But while this example points out an important difficulty with such an approach to educational research, it fails to capture the important characteristics of purposeful behavior since in the systems described by Taylor (automobiles) the action of the speedometer indicator can be considered as a computed output variable which depends on the values of the independent input variables. That is, unlike a human being (and unlike a car with a functioning cruise control system), it is a nonpurposeful system characterized by one-way, external causation.

Reference

Taylor, J. G. (1958). Experimental design: A cloak for intellectual sterility. British Journal of Psychology, 49, 106-116.

Date: Thu Jul 01, 1993 3:15 pm PST
Subject: Re: Conrolling Via Language; PCT & Other Theories

From Tom Bourbon [930701.1746]

>[From Hank Folson (930701)]
>
>Has anyone done a diagram showing how a person controls through the medium
>of language/writing? Aren't the controlling actions parallel,
>multileveled, and indirect, unlike for controlling direct actions such as
>controlling the position of the knot in a rubber band by pulling on it, or
>moving a joy stick to control a dot on a screen?

Hank, early in June I posted something that included a few diagrams of social interactions. One diagram, in particular, showed the place language might play in an interaction. I am late for dinner, but I will locate the file and repost that diagram tomorrow.

Until later, Tom Bourbon

Department of Neurosurgry
University of Texas Houston Medical School Phone: 713-792-5760
6431 Fannin, Suite 7.138 Fax: 713-794-5084
Houston, TX 77030 USA tbourbon@heart.med.uth.tmc.edu

Date: Thu Jul 01, 1993 3:36 pm PST
Subject: Applied PCT

[From Rick Marken (930701.1400)] Hank Folson (930701)

>Has anyone done a diagram showing how a person controls through the
>medium of language/writing?

I think I posted a diagram of one side of a conversation; it was back when I was still in my old office so it must have been around the end of last year.

Maybe Greg W. can find it. I recall that Bill P. liked it but I don't think it made much of an impression on anyone else.

>Summary: My feeling is that people will resist even a perfectly proven
>PCT for the reasons covered in our posts, so perfecting PCT is only a
>part of the answer. I think that the other part of the answer is to
>begin applying PCT where appropriate. Positive results of PCT approaches
>will then create more general interest in learning about PCT.

I really enjoyed your post, Hank. I find myself agreeing strongly with the first part of your summary; I think many people will resist PCT even when it is "proven" as conclusively as Newton's laws or relativity. And I think this is expected (as you say) if people are control systems; they will act to keep their perception of what constitutes an explanation of behavior where they want it. This will be especially (but, of course, not exclusively) true of people who don't undertand the scientific approach to understanding (observe, model, test). That is why I try to bring PCT to the attention of what is purportedly a scientific audience; at least there is the hope that I can point to an observation that they cannot model (at least, with a model that will pass tests that compare its behavior to the behavior of the system being modelled). I kind of agree with Dag that people who are not already convinced that PCT is right (for whatever reason) and who don't have at least a rudimentary understanding of science and mathematics simply cannot be convinced of the value of PCT.

I find myself wondering about your second point, though; that a way to interest people in PCT is to start showing what can be done when it is applied. I don't understand what you mean by "applied". How do you "apply" PCT? What is the evidence that the "application" of PCT produces demonstrably better results than not applying it? I guess I don't like the idea of evaluting PCT by its applied success because I don't think there is any evidence of its applied success; at least, no evidence that it is any better than that for any other behavioral theory. Is there QUANTITATIVE evidence for the success of clinical applications of PCT that is any better than that for the success of any other clinical approach? Is there QUANTITATIVE evidence for the success of management applications of PCT that is any better than that for the success of any other management approach? I expect these questions to be a disurbance to Ed Ford and Dag Forssell. When I say evidence, by the way, I mean EVIDENCE -- like the evidence that people control perceptual input variables (.99+ correlations between human and model data).

PCT shows the futility of trying to control other people; it explains (precisely, but in general terms) how people come into conflict with others and with themselves. In short, PCT shows the problems that can be expected when we try to get people to behave the way they "should" So I suppose one is applying PCT when one knows what it means to be controlling another person and when one voluntarily doesn't do such controlling. But it seems like it would be hard to show the benefits of this application of PCT, which accrue mainly to the person who manages to avoid conflict for him or herself.

I think the main application of PCT is to oneself; PCT can make one's own life better in ways that are measurable in terms of one's own perceptual experience -- you just feel better. There is less stress; there is the joy of understanding others and knowing how do deal with them in an almost conflict-free manner. PCT has techniques that can help one regain control (by solving internal conflicts), accept change (reorganization -- it happens) and deal with emotional perceptions. But these are personal benefits of the application of PCT. If people are happy without PCT (and most people apparently are) then talking about these benefits is hardly a selling point.

So I think I am against the idea of selling PCT in terms of its applications, unless its in terms of its applications to' ...the motorcycle that we work on all the time - ourselves" (or whatever it was that Pirsig said so well in "Zen and the Art of Motorcycle Maintenance"). But I can be convinced. Where is the evidence that the results of applying PCT are any better than the results of applying any other theory of behavior?

Best Rick

Date: Thu Jul 01, 1993 3:43 pm PST
TO: * Dag Forssell / MCI ID: 474-2580
Subject: info

Jag vet att du sa att alla haller pa med TQM, men jag tankte enda ge dig namn pa chefen for tqm har pa nasa, eftersom jag horde tal av henne haromdagen och det verkade som om de precis satter igang har, dvs att det inte har ett stort program ennu, men forsoker fa stort program inom kort.

Hon heter Jana Coleman och har t. nummer 604 5085.

OKej, hejda! Lisa

Date: Thu Jul 01, 1993 4:49 pm PST
Subject: Controlling by language

[From Bill Powers (930701.1815 MDT)] Hank Folson (930701)

How does a person control through the medium of language or writing? Good question. This is an excellent opportunity to see if you understand PCT. If I tell you, you will say "Oh," and think up another question. If you try to work it out for yourself, with any reminders you may need, you will be able to answer the next question yourself.

It's best to start with a specific example. So make up one, and try to answer the following questions:

1. Control what perception?

2. With respect to what reference condition?

3. Through what actions?

4. With results observed how?

Best, Bill P.

Date: Thu Jul 01, 1993 10:58 pm PST
Subject: Applications of PCT

[From Dag Forssell (930701 2340)] Rick Marken (930701.1400)

>Is there QUANTITATIVE evidence for the success of clinical
>applications of PCT that is any better than that for the success
>of any other clinical approach? Is there QUANTITATIVE evidence for
>the success of management applications of PCT that is any better
>than that for the success of any other management approach? I
>expect these questions to be a disturbance to Ed Ford and Dag
>Forssell. When I say evidence, by the way, I mean EVIDENCE --
>like the evidence that people control perceptual input variables
>(.99+ correlations between human and model data).

Not much disturbance. What other clinical approach? What reigns in the real world is mass confusion, as I see it.
....
>I think the main application of PCT is to oneself; PCT can make
>one's own life better in ways that are measurable in terms of
>one's own perceptual experience -- you just feel better. There is
>less stress; there is the joy of understanding others and knowing
>how do deal with them in an almost conflict-free manner. PCT has
>techniques that can help one regain control (by solving internal
>conflicts), accept change (reorganization -- it happens) and deal
>with emotional perceptions.

Rick, you are answering your own question. Now quantify your own measurable perceptual experiences.

>But these are personal benefits of the application of PCT. If
>people are happy without PCT (and most people apparently are) then
^^^^^^^^^^^^^^^^^^^^^^^^^^^^^^
>talking about these benefits is hardly a selling point.

How do you know that people are happy? Have you quantified it?

>So I think I am against the idea of selling PCT in terms of its
>applications,..

Any idea totally without applications is of little interest in my book. You have lived too long among "scientists" whose work is useless and has no application because it is wrong.

>.. unless its in terms of its applications to' ...the motorcycle
>that we work on all the time - ourselves" (or whatever it was that
>Pirsig said so well in "Zen and the Art of Motorcycle Maintenance").

You are speaking out of both sides of your mouth, just to get an argument going on this net. (I agree that it is quiet. Perhaps Gary has pulled the plug on us to test our reference for net activity). Your own interest and your own statement above provides evidence. What is wrong with working on your own motorcycle? What is wrong with teaching others how to work on their own motorcycle? This is the application of PCT that every personal and interpersonal problem reduces to.

>...But I can be convinced....

No you can't. Only you can convince yourself. Are you really interested in applications of PCT? It seems to me that all you have to do is to observe yourself when you apply it.

Best, Dag

Date: Fri Jul 02, 1993 10:10 am PST
Subject: Hancock experiment, PCT applications

[From Rick Marken (930702.1030)] Tom Hancock --

>My questions about understanding are rooted in this paper
>but are also an attempt to go beyond this present work.

OK. Here are some quick comments on the paper from a PCT perspective.

>Thus when a student in a class or a subject in an experiment
> perceives a feedback message

>2. continues to spend time with the message (probably executing
>"processing" programs) as long as the message continues to appear
>helpful in reducing the magnitude of discrepancy.

The perception of something about the message IS the controlled perception, according to your hypothesis. So the message can't be "helpful" in reducing the magnitude of the discrepency -- it is a perception and it is controlled. Your statement above is (unfortunately) completely consistent with the idea that there is information in perception about what to do to reduce discrepencies between perceptual and reference signals. This is NOT how control works -- so this research is not based on the perceptual control model. If you really imagine that subjects are controlling something about the message, then why not test to see what it is?

>We propose that a means of testing a subjectUs higher level control
> is by observing his/her use of the feedback message.

There's that "use of feedback" again. Control systems don't use feedback; they control it. Higher level systems "use" lower level systems by varying references for their perceptions (feedback). You can test for higher level controlled variables by observing how lower level references are varied -- but you must also monitor the distubances that are affecting the higher order variable and/or the state of the higher order variable itself.

>By looking at

>the variance of feedback frame latencies across levels of
> metacognitive judgments

>we hope to determine whether each subjectUs
> higher level control systems are efficiently related to learning
>to respond correctly.

I don't know what a "frame latency" is -- but, more important, I don't see why the variance of these frame latencies is relevant to determining the subject's higher level goals. Why just "hope" to determine what the higher level systems are controlling? Why not approach the problem systematically. I think a diagram of your model would be a good start.

>We assume that a subject who rates his/her certainty of
> responding correctly is thereby providing evidence of the
> amount of discrepancy that persists in control systems
> related to responding correctly.

> Based on these assumptions, we would hypothesize the following:

These are pretty specific predictions, though I'm not exactly sure how they were derived from the model. I hope they all came out exactly as expected for every subject. [Now that I have read ahead I see that none were confirmed so I guess they can be dropped now, right?]

>Each feedback screen displayed response sensitive feedback information.

Isn't feedback, by definition, response sensitive? It is the effect of one's own actions (responses) on one's own perceptual input.

>With the same analyses conducted separately on each subject there
> were the following trends:

Did the independent variable account for 99% of the variance in the dependent variable for ANY subject? If not, I'd forget about publishing this until I got some real results.

You seem to be very attached to this kind of experiment -- where a person rates his or her certainty that their response was or will be correct. It strikes me as a very peculiar control task -- what can the subject control? Whether or not s/he participates at all could be one controlled variable. Another might be whether s/he hears "right" or "wrong". What else? Think of yourself doing the experiment; what might you be trying to control? What are potentially variable aspects of the situation that you might want to keep in a particular state? When I think about this task it just seems like there is not that much to control. I can't really control the relationship between ratings and correctness of answers because I don't know how correct my answers REALLY are. I might be controlling the relationship between my ratings and my imagined confidence in the relationship between what I say and what I see; but it's hard to test for control when one aspect of the controlled variable is an imagination. But it seems like that's what you are interested in -- what subject's do about perceptions that have a large imagination component. "Confidence" about the correctness of a response is a perception, but it must be largely self generated -- an imagination. But, maybe it is based on perception of error signals, as you suggest . In that case, one can disturb this perception by giving false reports about the status of the subject's answers. But it's still tough to know whether the subject will buy them -- since you don't know how well the subject actually knows each answer.

Conclusion: Based on the data you report there seems to be no phenomenon here to explain. The group subject data is random noise and I see no evidence that the individual subject data is any better. I would suggest that you work with ONE person, keep trying variants of your study until you can identify a variable that is unquestionably controlled by this subject, and then develop a working model of this control process. New questions will probably suggest themselves once you get to that point.

Dag Forssell (930701 2340) --

>You are speaking out of both sides of your mouth, just to get an
>argument going on this net.

Caught me! Actually, I was just not being clear. There are two senses of the phrase "application of PCT " for me. One suggests the use of PCT as a tool to get others to behave properly; the other suggests the use of PCT as a scientific model (for understanding) and for self-improvement. Obviously, I am a big fan of the latter; I think PCT itself shows that the former is illusory and, ultimately, destructive.

>What is wrong with working on your own motorcycle? What is wrong
>with teaching others how to work on their own motorcycle? This is
>the application of PCT that every personal and interpersonal
>problem reduces to.

I agree completely, of course.

Let me try to be explicit about the point I was trying to make. I was responding to Hank's suggestion that positive results of applications of PCT will be the best way to get people's attention. This suggests that you can point to what people typically regard as "positive results" (like charts showing productivity going up exponentially or recidivism rates after treatment going down or whatever) and say "look, PCT works". This leads to the idea that the success of PCT can be presented in something like a corporate report; "Profits up, costs down after application of PCT!" It is this kind report of the results of the "application of PCT" that I object to. This is not what PCT is about. The concensual, public "success" of the application of PCT can only be judged in scientific terms -- does the model explain the phenomenon of purposeful behavior as it is manifested in living and artificial systems? If this understanding gives one personal satisfaction or if it helps one deal with life better than that is a successful personal "application" of PCT. But I don't think it will satisfy Hank's goal, since it is not a public display of the success of PCT (in fact, the person who has "successfully" applied PCT to their own life might not be judged a success -- especially by those who see no value in PCT anyway).

It's great to teach people PCT if they want to learn it; and, of course, I think PCT is an extraordinary concept and a powerful scientific model. I also think people can benefit personally from their knowledge of PCT. But I think it's a mistake to try to "advertise" PCT in terms of its benefits to one's "bottom line". Some people are happily controlling for other things besides bottom lines; many people who are controlling for bottom lines are as likely to get through the gates of PCT as a camel is to get through the eye of a needle.

PCT is not about helping companies move into the fortune 500 or helping people live like Ozzie and Harriet. It's about how people actually work. Based on this knowledge, people will have to figure out for themselves whether their own interest in making megabucks or living in imagination mode is the best way to get their perceptions of their society and themselves to be as they want them.

Best Rick

Date: Fri Jul 02, 1993 10:12 am PST
Subject: Conference registration

[from Mary Powers 930702]

I'm happy to say that CSG conference registration got quite lively over the last two weeks - we will have about 28 participants and 5 guests. It is not too late to register - up to July 10 (since I sign the final version of the contract with FLC on the 13th). Full payment for the conference is due BEFORE the meeting, mainly because I don't want to have to mess around with bookkeeping and banking while the meeting is going on. So get those checks in the mail, please.

Mary P.

Date: Fri Jul 02, 1993 10:31 am PST
Subject: A modeling question

From Kent McClelland (930630) Reply to Tom Bourbon (930623.1306)

Thanks, Tom, for your thought-provoking answer to my earlier post (Kent McClelland [930622]) on the possibility of "Collective Controlled Variables." I've taken a while to get back to you on this, but I've needed some time to consider the points you made.

In response to my argument that some "alignment" of reference levels between independent control systems is necessary for cooperative actions to occur, you first point out that cooperation often requires individuals to adopt different reference levels, not the same ones:

>What matters is that the participants adopt reference
>signals (that they adopt goals which we model as reference signals -- Hans?)
>that result in each acting in a way that produces a match between personal

>goals and present perceptions. The same constraints apply to my two
>hemispheres when "they cooperate" to produce one-person performance of what
>can also be a two-person cooperative task: they need not adopt similar
>reference signals; all that is necessary is that each adopt reference
>perceptions that result in the intended perceptions. . . .
>Cooperation can, probably always does, entail some necessarily different
>reference perceptions in the participants.

In sociological terms, we're talking about "division of labor," and your point is well taken. I'm a little unclear, though, about what you mean by "adopting goals" and how that's different from my "practical alignment" of reference levels, if the goals are necessarily shared in order for cooperation to take place. You go on to suggest that even when people perceive themselves to be sharing the same goals, they probably aren't, at least not in any detail:

>. . . that kind of agreement says nothing about the
>specific contents of the participants' heads. The socially affected and
>approved options all reside in individual heads, as products of individual
>interpretations or understandings. Necessary and important, to be sure; the
>same in each person, probably not.

A good many sociologists would heartily agree with you on this point, especially those who call themselves "ethnomethodologists." This band of maverick sociologists have made their careers by studying the methods people use to convince themselves that their perceptions of reality match everybody else's perceptions, in other words, that "social reality" is really real. The "ethnomethods" described by ethnomethodologists are techniques by which people succeed in glossing over ambiguities and putting off any discussion of apparent conflicts. In various ways, such folk techniques make it possible for "normal" social interaction to take place on the mistaken assumption that everybody understands each other and sees the world in the same way. The ethnomethodologists themselves apparently believe that the semblance of social reality achieved by these techniques is merely a shared illusion and that, were it not for the use of ethnomethods, social life would be even more chaotic than it already is.

As I've listened at various times to discussions of epistemology on the net, I've thought that the most ardent PCTers might enjoy reading a little ethnomethodology. And I recall now that Rick Marken described to me how he once studied under one of the founding fathers of ethnomethodology (Schegloff, was it, or Sacks?), which might explain a lot about Rick! These founding fathers (especially Harold Garfinkel) are mostly unreadable, but a recent book on ethnomethodology I'd recommend is MUNDANE REASON by Melvin Pollner (Cambridge University Press, 1987).

Anyhow, my sociological conclusion about what you're saying is that you are undoubtedly right in emphasizing that each individual's perceptions are different, but I still think that at some abstract level people must have their control systems at least crudely aligned in order for cooperation to take place and that any such alignment has got to be an important sociological issue.

In the last part of your post, you dismissed my suggestion that independent control systems when aligned can become "super powerful", if there are enough of them. You described in tantalizingly brief terms some cooperation experiments that you and Michelle Duggins will be presenting at the CSG conference. In your words. . .

>In none of the cases I have modeled did a quartet take on super powers. . .

I'm curious about what exactly this means, because of some spreadsheet simulations I've been playing with. I find that when I couple two 50-gain control systems to the same "external variable" and I give them identical reference levels, their joint performance (whether measured as their perceptual variables or their summed output) is indistinguishable from the performance of a single 100-gain system given the same reference level and the same pattern of disturbances. The correlation between either the input or the summed output of the two coupled systems and the input or the output of the single "double-power" system is 1.000. Actually, any combination of two or three systems seems to work the same way, as long as the gains add up to 100 (or whatever the gain of the single comparison system is set to be), and as long as all the slowing factors are the same.

I take from these simulations (which would no doubt be easier to do in Simcon if I weren't working on a Mac), that "perfect" alignment of reference levels has the practical effect of adding up the gains of the participating systems. I find it easy (maybe too easy) to make the jump from this highly abstract simulation to a human social world where alignment is always imperfect but nevertheless there seems to be strength in numbers. Does this make any sense to you?

Till later, Kent

PS Tom, I hope you'll soon be inspired to share your thoughts on "imitation."

PPS It's raining again in Iowa tonight. Pretty soon the whole state will have washed down the Mississippi, and all of us Iowans will be moving to the Louisiana delta country whether we like it or not!

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: Fri Jul 02, 1993 2:22 pm PST
Subject: Speech

From Tom Bourbon [930702.1655]

SAYING /di/ and /du/

Andy Papanicolaou and Tom Bourbon

Many months ago, we floated (1) some thoughts about what kinds of facts a PCT model of speech perception-production must account for and (2) some questions about how best to proceed in designing such a functioning PCT model. Circumstances on our end prevented us from joining in the discussion that followed, and in others that continued on-and-off regarding special issues in modeling speech. Recently, our circumstances and our access to the net have become more favorable and we would like to try again.

The ideas expressed in our post have emerged during many hours of animated conversation and they reflect some of our ideas about how and why many people reject the idea that PCT can contribute to a better explanation or understanding of speech. We have tried to insert disclaimers and reminders in parts of the post where we present our interpretations of "those people." Hint: There is no need to remind us that "their" ideas are different from "ours" in PCT. (A note about the production of this post: We do a lot of collective talking and handwaving, then Andy summarizes it on paper and Tom does the computer stuff, with both of us editing on the fly.)

Going over what has been said on the net so far, we realized that too many related issues had surfaced -- too many, that is, to be addressed at the same time. So we thought it might help to back up a little and focus on just one topic that many people describe as a special issue in speech -- one special (or apparently special) fact that has been used by people who argue in favor of motor plans as the most reasonable or intuitively appealing explanation for speech perception-production. The apparently special fact, or feature of speech, goes by the name of "context- conditioned variation."

First, let us go over what many people accept as evidence that this is a fact, so we (any of us on csg-l who are interested) can try to understand what it is about this "fact" that suggests to people the necessity of invoking "plans," then we (all of us) can try to demonstrate that the same "special" fact can be explained much more efficiently by PCT. Please remember that, in the beginning, we will be trying to summarize the ideas of people who accept plan-driven models as explanations of speech; we will state our own positions later in the post.

"Context conditioned variation" (ccv) is the name Liberman (1970, *Cognitive Psychology*, 1:301-323) gave to the variation in a part of the acoustic signal (as displayed in speech spectrograms) depending on segments of the signal that either precede or follow the part in question. Consider this example after Liberman; the production of /di/ and /du/. (Idealized acoustic signals for both are show in Figure 1.) The acoustic signal that corresponds to the perceived /di/ is said to have two parts: formant transitions (corresponding to /d/) during the first 50 msec or so, and steady-state formants (corresponding to /i/). The acoustic signal that corresponds to the perceived /du/ is also said to be made up of two parts. As practically everyone expects, the steady-state parts of /di/ and /du/ are obviously different sounds. However, contrary to the expectations of many people, the formant transitions for the perceived /d/ preceding those steady states also differ. In the case of the /d/ in /di/, we have two rising formants; in the case of /du/, one rising and one falling. If we change any of the transitions, for instance by changing the falling transition in /du/ to a rising one like we find in /di/, we no linger hear /d/ at all, when we play the spectrogram.

```
=====================================================

Figure 1.  ASCII drawings of idealized spectrograms for /di/ and /du/ follow:

                 /di/                /du/

          |  **************
     high | *
          |*
          |                      *
          |                       *
frequency |                        **************
          |
          |  **************        **************
          | *                     *
     low  |*                     *
          |__________________    __________________
           0             300     0             300
                        time in msec

==================================================
```

Now, there are all kinds of implications of this phenomenon. The one we think is most relevant here is the following: It appears to many people as if, depending on what steady-state formants the articulators *are going* to produce, they modify the transitions that they *do* produce, and the acoustic signal that results in either case is perceived as /d/.

(Note: This apparent phenomenon, context conditioned variation, used to be thought of as unique to speech, but we understand that many people say it is evident in various places, for example, in "anticipatory" postural adjustments that "precede" gross movements such as taking a step or standing still and stretching an arm out in front of you. In fact, the phrase, context conditioned variation or variability, is now used to label any variability, in otherwise "automatic" or "stereotypic" actions, due to changes in the environmental setting in which the actions occur. But let's not stray. The spectrographic evidence about /di/s and /du/s is more than enough for today.)

The question we want to address is this: Why is it that, once they see data like those on /di/ and /du/, so many people immediately think (positively) of plans and commands as explanations, and why are they so reluctant to think about PCT loops as explanations? The motivation seems pretty obvious: If you have a prefabricated plan for the whole /di/ and another for the whole /du/, once you want to say one of the sounds, specific commands go out in parallel and in the proper sequence and you produce, every time, exactly what you want to produce. That's one scenario at any rate. (Of course, there are others.) One reason the scenario seems plausible to so many people is that you don't need a billion such plans. After all, there are only so many consonant-vowel combinations in English (or in Albanian), and no more. (We recognize that one reason the idea seems plausible is that advocates of plans often ignore, or deny the importance of, disturbances that occur while the commands are being released or after they have gone out to the articulators.)

Next, why might people think PCT cannot explain the spectrographic evidence, besides the fact that PCT loops are incompatible with the intuitively appealing plans and commands? For starters, because once you intend /di/, your articulators apparently produce it right away. For many people, it seems counterintuitive that such rapid and accurate production could occur without specific instructions from above; how could the simple error signals in PCT produce the right results every time? And about those error signals: how could they *initiate* the production of the desired sound? After all, they occur only *after* the articulators have produced the *wrong* sound. But the articulators seem to do their job unerringly, without hesitations, as if they *knew*, as if they were *told* the right thing to do, from the very beginning. (Remember, first we are trying to summarize the views of people who advocate plans. The ideas are not our own.)

To make matters worse, a plan advocate might say, imagine that you intend to say /di du/. The articulators produce the first /d/. Doesn't it seem reasonable, since the reference signal in PCT would call for another /d/ right after the first one, that the articulators would do the same thing they just did 100 msec ago, producing the *same* perceptual signal to match the *same* reference signal, /d/? But that is not what happens. If the system tried to do that, the result would be /di -- something else -- u/, rather than /di du/. There are, again, other reasons why this evidence leads people toward plans and away from PCT, but before we address those, it seems absolutely necessary to address the timing issue: I intend /di/ and right away /di/ happens; I intend /du/ and right away and without errors /du/ happens. The phoneme production rate is phenomenal -- in fact it was that rate that led Al Liberman's group at Haskins Laboratories, and Lashley before that, to reject all S-R accounts of speech production and to embrace the notion of plans.)

If PCT is correct, then how is it, the linguist asks, that without up-front guidance, and only on the basis of "after-the- fact" guidance provided by error signals, the muscles so quickly do the right thing? Notice that what for most linguists seems to constitute a paradox (apparently dispelled by the invocation of plans), in the form of the phenomenally rapid -- seemingly instantaneous -- rate with which the muscles assume complex configurations simultaneously and successively and "do the right thing, also provides the most natural opening for us to introduce the concept of the PCT loop.

So then, how is it that several articulatory muscles contract in the appropriate manner, degree, and sequence as soon as the reference signal is instituted? Or do they? What better way could than there be to proceed in demonstrating that PCT loops can *explain* the phenomenon, than to actually construct a PCT model that *produces* the phenomenon?

THE ISSUE OF REFERENCE SIGNALS

So, how do we build a PCT model that perceives-produces speech? First, we believe that it can be convincingly demonstrated that the timing problem is not insurmountable and that PCT loops *can do* what "plans" *are supposed to do* and that the loops can do it better. For example, PCT loops can overcome disturbances while plans can not. Also, error signals in PCT loops arise immediately when a new reference signal is instituted and affect behavior at that moment; they are not signals about after-the- fact knowledge of results. There are many other ways in which we believe PCT loops will outperform plan-driven models when it comes to the issue of the "timing problem."

However, we did find a problem that, though not insurmountable *in principle*, is currently insurmountable *in practice*, or so it seems to us. We are talking about the question of choosing the appropriate reference signal in a PCT model of speech. Hold on, here it goes.

For a model to produce /du/s and /di/s, it must be able to recognize /du/s and /di/s. That's obvious. And the model must control its perceptions of /du/s and /di/s. In order to build or program the PCT model, let us see how *we* control our perceptions of /du/s and /di/s.

Say that I am listening to a radio program and, control theorist that I am, I adopt for my self an unusual reference signal: Every time I hear the /d/ sound, in any context whatever that it may occur during the program, I will say /di/. For example:

```
said:        /di/       /di/       /di/
heard:   the doors were drawn closed...)
```

In other words I don't set myself the task to control for the meaning of the program or the quality of the announcer's voice or anything else but for perceiving /d/ and saying /di/. And, not surprisingly, I can do the task very well.

Now we want to build a CT model that will do the same thing. Obviously one of the first things we need is to put into it some reference signals. Certainly one for /d/. After all we want the model to control for /d/. And here is where the problem is: we know perfectly well, subjectively, what it means to control for /d/. We are controlling for a particular experience, in this case auditory, which corresponds reliably to something outside our system, in this case to what comes out of the radio, to some aspects of the stream of sound that the announcer makes. Beyond the knowledge that we are controlling for that particular experience there is not much more that we know. But certainly, we can not put "an experience" into our computer model. We can only put into it a signal (a number, a voltage), which in some ways will have to correspond to the signal (number, voltage) coming out of our model's input function, the p signal that is present when our model achieves its task of controlling for /d/. Equally certainly, the reference signal must correspond to some aspect of the stream of sound (that is to those portions of the acoustic signal that we hear as /d/). We have to construct our model in such a fashion that when its sensors encounter the same "something" in the stream of sound that we hear as /d/, the model should, like us, say "di".

So, to what feature of the sound stream should our reference signal be made to correspond? The answer is, but of course, to that portion of the sound stream that, when it hits our ears, we experience /d/. So we record the sound stream and we have a time series (amplitude over time). Then we go back and identify those spots where when played we hear /d/. When we do that though, we discover, as Liberman and associates at Haskins discovered years ago, that no two portions of the time series that correspond to our experience /d/ are alike. Therefore we can not tell our model, "Please control for this amplitude or amplitude variation," because there is no such invariant amplitude variation that repeats itself whenever we experience /d/.

The next thing we could do (again following the Haskins route), is to transform the recorded time series into a long speech spectrogram (frequency over time). We play that back and lo and behold, here and there we see particular configurations of formants and formant transitions that correspond to our experience of /d/. But here and there is not always. In fact here and there is, at best, so seldom that if we build our model to control for those particular configurations, then unlike us, the model will miss most of the /d/s, and it will produce a lot of false positive responses.

We, therefore, like the Haskins crowd before us, may conclude, after decades of trying, that there are no invariant features in the acoustic signal that our model can control for. Or we may be more reserved and instead of making an existential or ontological assertion, "there are no invariant features in the acoustic signal that correspond to the experience of /d/," we may opt for an epistemological one like, "we have not found yet any such invariant features." Which means that we may in the future, if, for instance, we transform the raw acoustic signal in another way (how does phase space sound?). But the reality is that, right now, we don't know what our model should control for.

Some of the Haskins folks would be ready, at this point, to make a suggestion. They will say: Aren't you guys looking for something in the world to which the experience /d/ corresponds? If so we think we have something for you to use in your model. Namely we have discovered that the only thing, thus far, that corresponds to the experience /d/ is a pattern of articulatory movement. True these movements are not metrically the same every time one produces an acoustic signal that is heard to contain /d/, but there are some aspects of these movements that are always the same (we can show movies of them, if you don't believe us) every time /d/ is heard. For example, every time /d/ is heard, the speaker's vocal cords vibrate (in the case of /d/, neither the frequency or amplitude matter) and at the same time the tip of the tongue is pressed either against the front of the palate or the back of the upper teeth (where it obstructs the air flow), then the tongue moves, releasing the flow of air. (Notice that there can be variability in the details, but relative timings of vibrations and release of air must be invariant.)

So, they would continue, we propose that what corresponds to the perceptual experience of /d/ is an articulatory plan that the hearer activates in his head - against which he compares the stream of sound from the radio. (He controls for sounds that would be sufficiently similar to the sounds that would be produced if the particular articulatory plan were to be executed).

We hear that suggestion and we recoil at the thought that we should endow our model with "nascent plans" as reference signals. It is perceptual signals that our model should control for, we insist. And that's as it should be. But, being pragmatists rather than religious fanatics or day-dreamers, we ought to assess what we are faced with in our attempt to construct a model that will be like us in one respect: that it will be able to respond like us with /di/ every time it "hears" a /d/ in the stream of speech.

What we are faced with is this:

(1) There are no known acoustic invariances that our model can be made to control for.

(2) There are articulatory invariances - i.e. particular sequences of movements of articulators. These are unacceptable, or appear to be, for reference signals. One reason they are unacceptable is because we may refuse to believe that movement invariances exist at all in spite of the evidence. In the long run, that refusal may prove to be a good thing because the data may be faulty (it has happened before); or it may prove to be detrimental to our efforts if the data are solid. Maybe the linguists can tell us what is most likely the case here.

A second reason we find the suggestion of articulatory invariances unacceptable may be because we perceive the danger of beginning to think that, since they exist, we might as well use them to drive the articulators. But that outcome need not happen. Ever.

A third reason is that we see no use for them. After all, the model should control for perception not for movement. But the fact that we see no use for the reported articulatory invariances does not preclude our finding one for them in the future. In fact we have at least two options at this point in time, and we could pursue them both: (a) to insist on looking for invariances on the acoustic signal side (the business of the phase-space was not a joke); and (b) to explore the possibility that a variant of "motor theory", where articulatory invariances are used to explain perception, may help us identify the reference signals we should program into our PCT model.

There may be other options than the ones we have identified, but we have to do something. We can handwave perception only so long. And we can't handwave it at all, if we want to create a working model that says /di/ every time it "hears" /d/, which must correspond to every time *we* experience /d/ and say /di/.

To our colleagues in the States, have a nice Fourth of July.

Andy and Tom

Date: Fri Jul 02, 1993 2:54 pm PST
Subject: Aligning goals

[From Bill Powers (930702.1530 MDT)] Kent McClelland (930630) --

Some thoughts about "aligning reference levels" that might prove useful:

Suppose that four people undertake a task together. They are to lift a folded card table by its corners and hold it level. Someone else holds a hand out and says "hold it about here" and then removes the hand.

Person A thinks of the goal as holding the table 3 feet above the floor.

Person B thinks of the task as holding the table 5 feet below the ceiling.

Person C thinks of the task as holding the table about level with his belt-buckle.

Person D thinks of the task as holding the table about twice as far above the floor as the seat of a chair.

Now, are these reference conditions "aligned?" Each person is controlling a different perception. Each reference condition is, of course, conceived of in the way relevant to that perception. Yet the result is for each person to want the table to be in about the same position. They will each agree that the table is in the right place in space when they finish.

Now consider what can happen when conditions change. The four people are moved into another room and told to "do it again," but the ceiling in the new room is two feet higher and the only "chair" in the room is a footstool. The "aligned" goals are suddenly no longer aligned, and conflict will result. "You're holding your corner too low!" "No, you're holding yours too high!"

This bears some thinking about, doesn't it?

Best, Bill P.

Date: Fri Jul 02, 1993 3:35 pm PST
Subject: The Practical Value of PCT

from Ed Ford (930702:1350) Rick Marken (930701.1400)

>Is there QUANTITATIVE evidence for the success of clinical applications
>of PCT that is any better than that for the success of any other
>clinical approach?

You might ask those to whom I have taught PCT and have shown its practical applications in their areas of interest and who pay me large sums of money (more than I'm worth, but I don't argue) and who keep asking me back for more instruction and help. Everything they've been doing has been based on S-R psychology. I would say that QUANTITATIVE evidence comes in various sizes and shapes.

>I expect these questions to be a disturbance to Ed Ford and Dag Forssell.

Not at all. I think I experience the same thing you do when you have someone try to tell you the limits or drawbacks to PCT and they've never really learned what it's all about. But unlike you, I consider the source and then say to myself, "Do I have control over how this person thinks?" Obviously, I don't. I only deal with those things over which I have at least some control.

Rick, get out of your ivory tower office and come on into the trenches and find out for yourself. Try working on a factory floor, in a business office, or in a school setting. Work with real people having real problems and try to apply PCT. Get out where the rubber hits the road. If not, try reading Freedom From Stress and provide the net a detailed report of why the pracitcal applications found in the book don't apply. And what's the purpose of all this talk of PCT if we can't apply it. The modeling and theorizing of PCT are extremely valuable to me and others trying to apply these ideas, but everything we say about PCT is really worthless unless these ideas can guide us toward a better way of helping others live their lives more effectively, in a more satisfying way.

Finally, my thanks to Dag for his remarks. Well done, good friend!

Best, Ed

Date: Fri Jul 02, 1993 4:57 pm PST
Subject: Application

[From Dag Forssell (930702 1610)] Rick Marken (930702.1030)

>Let me try to be explicit about the point I was trying to make.
>I was responding to Hank's suggestion that positive results of
>applications of PCT will be the best way to get people's
>attention. This suggests that you can point to what people
>typically regard as "positive results" (like charts showing
>productivity going up exponentially or recidivism rates after
>treatment going down or whatever) and say "look, PCT works". This
>leads to the idea that the success of PCT can be presented in
>something like a corporate report; "Profits up, costs down after
>application of PCT!" It is this kind report of the results of the
>"application of PCT" that I object to. This is not what PCT is about.

In your subjective opinion.

>It's great to teach people PCT if they want to learn it...

>It's about how people actually work. Based on this knowledge,
>people will have to figure out for themselves ...

Agreed, but again, you speak out of both corners of your mouth. Where does your "..if they want to learn it.." come from. Precisely from the things Hank mentioned and you (mistakenly) argue against.

Hank has learned from PCT that people respond (take action) when they have error signals. When he writes advertising copy, he recognizes that it is wise to guess very carefully about what the consumers of your product really control for; what they want and how they perceive, and describe your offering in those terms. Here, Hank is applying his insight, not teaching it. PCT has value for Hank, and he continues to study it, to gain even more advantage.

Thus, PCT can be "sold" to advertisers based on its effectiveness on their bottom line. Certainly, I am selling PCT to companies based on the favorable impact on their bottom line from increased productivity, cooperation and satisfaction among all employees. Ed is selling PCT to parents and teachers based on the favorable impact the application of PCT has on their work environment.

As long as people are satisfied with what they know, they have no reason whatsoever to look at PCT. Applications are the only reason for people to come to: "..if they want to learn it.."

Best, Dag

Date: Fri Jul 02, 1993 10:39 pm PST
Subject: PC-Eudora and NUPop

[from Gary Cziko 930703.0600 UTC]

The following is an article from a local campus newsletter which describes two e-mail programs for IBM PCs and compatibles which CSGnetters may find of interest. You will see below that the price is certainly reasonable (i.e., free).

I have been using Eudora for the Mac since 1990 and can't imagine managing CSGnet without it. I can provide info on this Mac program to anyone who asks. But remember, your mainframe link to the Internet has to be running POP to use these programs.

Also note that (a) PC Eudora currently requires a direct network connection, but NUPop can be used over a modem, as can Eudora for Mac (as I'm doing now), and (b) both programs can send and receive binary files as attachments to e-mail messages (Bill P. and Greg W., are you listening?).--Gary

P.S. I will be very impressed if anyone can tell me why one of these programs is called "Eudora" (people down South USA like Greg Williams and Chuck Tucker should have an advantage for this one). "NUPop" is too easy to figure out to impress me.

==

PC's are POPping with PC Eudora and NUPop

When Eudora, the well-known e-mail client for the Macintosh, was released at UIUC a little over two years ago, it literally revolutionized the way Mac users on the campus network (and eventually all over the world) processed their electronic mail. Eudora afforded people the luxury of preparing outgoing messages and reading and organizing incoming messages within the familiar confines of the graphical, menu-driven Macintosh desktop. For most Mac users--even those well acquainted with UNIX e-mail software and text editors--it was time to say, "Good-bye Elm and vi, hello Eudora."

Shortly after Eudora was introduced, PC users began clamoring for a similar e-mail package for DOS machines. Several DOS-based e-mail clients entered the public domain, but until recently, most were either unstable, unsupported, or lacking the extraordinary functionality and ease-of-use that made Eudora an overnight sensation. However, two packages--NUPop and PC Eudora--previously available in beta (test) versions, are now production software and hold much promise for the DOS user base. NUPop is a text-based application developed at Northwestern University that can run on both old and new PCs. For high-end PCs and compatibles, QUALCOMM, Inc. (current developer of Mac and PC Eudora) offers a Microsoft Windows version of Eudora that looks and feels very much like its Macintosh cousin.

Like Eudora for the Mac, both NUPop and PC Eudora are POP (Post Office Protocol) clients. In order to use a POP client, you must have an e-mail account on a multi-user computer that is running POP server software (all CCSO-administered mainframes including uxa, uxh, ux1, ux4 and VMD support the POP protocol). Most of the actions you will perform with the client (NUPop or PC Eudora) do not involve network communications. For example, when you prepare a new message, you use a text-processor built in to the POP client. Reading, organizing, and replying to incoming mail are also desktop operations. The only time the network comes into play is when you send or retrieve mail (or use special network utilities such as ph, finger, etc.).

PC Eudora and NUPop each possess a rich set of features and configuration options comparable to those found in Eudora for the Mac. Some of the more outstanding of these are:

- Integrated Text Processor. Creating new messages or replies is a breeze with the intuitive word processing capabilities built in to NUPop and PC Eudora. Word wrap and the ability to select, cut, copy, and paste blocks of text are standard features. Both packages also allow the user to open existing text files in order to paste all or part(s) of the file into a new message.

- Nicknames/Groups. You only need to type an Internet or BITNET address once to realize that such addresses are often long and difficult to remember. With PC Eudora and NUPop you can create aliases (called Nicknames in Eudora and Groups in NUPop) for the people or groups with whom you correspond regularly. Creating an alias involves entering a user's full e-mail address and a short, easy-to-remember nickname that corresponds to the full address. For example, an alias for the address of a colleague with the e-mail address mgeg8538@uxa.cso.uiuc.edu might be "Mel." Once an alias is created, you can use it in the To: or cc: field of the message instead of the full e-mail address. The client software takes care of expanding the alias to the full address when it sends the message to the recipient.

  Nicknames or groups can also consist of multiple e-mail addresses or even other nicknames. For example, the nickname "happy hour" might include the full e-mail address (or alias) for each member of the gang you meet after work on Friday afternoons. Helpful Hint: You can use the copy and paste feature to copy an e-mail address from the header of a message or the ph query window into your list of nicknames or groups.

- Binary and ASCII File Attachments. One of the shortcomings of electronic mail is that many systems can only process plain ASCII text, and thus it is often not possible to send binary data such as formatted word-processing files, spreadsheets, executable programs, etc., by e-mail. NUPop and PC Eudora get around this problem by encoding binary files in the format known as BinHex (for more information about BinHex format, see the November 1992 issue of UIUCnet, vol. 5 no. 7). The binary data is converted to a format that uses only ASCII characters and is attached to the body of the e-mail message. When a BinHexed attachment is received by PC Eudora or NUPop, the attachment is automatically restored to its original state and placed in a directory designated by the recipient. If the recipient uses an e-mail package that cannot decode BinHexed files, a special utility can be used to convert the file back to its native format.

ASCII text files can be attached to any NUPop or PC Eudora message and do not require special handling. They simply appear as part of the body of the message.

- Mailboxes and Folders. If you receive lots of e-mail, it's very helpful to be able to store it according to your own organizational scheme. This need is addressed in PC Eudora and NUPop with the Mailbox menu. Both packages come preconfigured with a few essential mailboxes such as IN, OUT, and Trash. New mailboxes can be created by selecting the appropriate command from the Mailbox menu. Once a new mailbox has been created, messages can be moved from one mailbox to another. When creating a mailbox in PC Eudora, you have the option of designating it as a folder, which is a mailbox that contains other mailboxes rather than messages. Mailboxes can thus be neatly organized into categories that form a hierarchical filing system.

- Automatic Message Retrieval. NUPop and PC Eudora can be configured to check for and download new mail from a POP account at regular time intervals. A pop-up window and/or alarm signals the arrival of new incoming messages. (Note: because NUPop is a DOS application, it is only convenient to take advantage of timed message acquisition when running NUPop as a background process under multitasking software such as MS Windows.)

- Built-in Ph Client. Standard queries can be sent to the CCSO Nameserver database with the built-in ph client found under the Special menu in PC Eudora or the Utilities menu in NUPop. Ph makes it easy to look up e-mail addresses, phone numbers, and many other types of other information about people and units on campus. You can also use it to get weather, area code, campus timetable, and local restaurant information. Also included in NUPop are a finger client (for looking up information about a user account on a specific host), the IP Finder (an nslookup client, which will return the numeric IP address associated with any fully-qualified domain name), a simple telnet client for logging in to remote hosts, and a webster client for looking up information on Webster Dictionary servers.

System Requirements

If you decide to run a POP client on your PC, which application should you use, PC Eudora or NUPop? Your choice will be partially influenced (if not dictated) by your PC hardware and software. The system requirements for PC Eudora are not met by many PCs on the campus network. For optimal performance, PC Eudora requires a 386 processor or better, at least 2 MB of RAM, a sizable hard disk (PC Eudora only requires 750 KB, but MS Windows uses roughly 10 MB without any other applications installed), mouse, and a color VGA or better resolution video subsystem (PC Eudora can run on a 286 machine with a Hercules monochrome graphics or EGA video card, but performance will suffer). On the software side, two commercial products-Microsoft Windows version 3.1 and FTP Software's PC/TCP release 2.1 or higher-must be installed on your machine in order to run PC Eudora (for more information on PC/TCP see the article entitled PC/TCP 2.11 Available-2.2 around the Corner on page ? of this issue).

NUPop, on the other hand, can run on both low- and high-end machines. The pseudo-graphical user interface (consisting of windows, buttons with drop-shadows, and pull-down menus) uses only characters from the ASCII and IBM extended ASCII character sets. Thus, even a plain monochrome video card and monitor will suffice. The NUPop executable and associated files take up roughly 700 KB of hard disk space and require 470 KB of free RAM after DOS and memory resident programs are loaded. The only special software required by NUPop is DOS version 3.0 or higher and a packet driver for your network interface card. (Packet drivers for common brands and models of network cards are free and ftpable from many locations on the Internet. In any case, most PCs on UIUCnet already have a packet driver installed. For more information on packet drivers, see the NetWord and Administrivia columns in this issue.) Power users can run NUPop in a DOS window as a background operation under MS Windows. In short, just about any PC with a hard disk and at least 640 KB of RAM can run NUPop.

Apart from system requirements, there are a few other items to consider before committing to NUPop or PC Eudora. Although the two programs are roughly equal in terms of overall functionality, each possesses its own unique set of features and flaws.

More NUPop Facts and Features

NUPop supports both plain serial and SLIP (Serial Line IP) connections by way of the CCSO terminal servers, a feature not yet available in PC Eudora. This means you can run NUPop from home or office with just a PC and modem (no direct network connection is required). If configured for serial communication, NUPop automatically dials the terminal server and connects to your mail server when you issue the command to send or retrieve mail. Once the transfer is complete, the connection is closed. Special NUPop script files customized for the CCSO terminal servers are available with the NUPop distribution files at the CCSO Resource Center.

Although both NUPop and PC Eudora are much easier to use with a mouse, most NUPop commands can be easily managed from the keyboard (running PC Eudora without a mouse, on the other hand, is extremely cumbersome). Menu and button commands are executed by pressing the <Alt> key in conjunction with the highlighted letter in the command name itself. For example, to open the Options menu, use the key combination <Alt - O> (see Figure 1 on page ?). The space bar acts as a toggle for selecting and unselecting messages within a mailbox. Additionally, a set of "hot-keys" have been defined for frequently used commands.

NUPop Flaws

Few public domain software applications are perfect and NUPop is no exception. Until you get used to it, mastering NUPop's multiple windows is quite confusing. At start-up, the Index (IN mailbox), Composer (editor for creating new messages or replies), and Viewer (used for viewing received messages) windows appear successively on top of one another, and it is not quite clear what is happening until the action stops. Invoking the ph and/or finger utilities causes additional full-screen windows to open up. Some NUPop windows have scroll bars and others have a close box. There is actually some method behind this window madness, but, at first, it is not obvious to the user.

In general, NUPop 1.0.3 is a fairly stable piece of software, but there are still a few bugs to be ironed out. For example, occasionally NUPop will behave in an unpredictable manner or lock-up completely. Such problems do not typically result in the loss of any data but are certainly alarming when they do occur.

How to Get NUPop

There are several ways to retrieve the NUPop software. CCSO staff have prepared a special distribution of NUPop that includes custom scripts for use with a modem and the CCSO Terminal Servers. This distribution is available on the Novell file server volume mounted on the PCs at the CCSO Resource Center, 1420 DCL (Resource Center staff can help you copy the appropriate files onto your own disks). The following files are available:

nupopx.zip - ZIP archive of NUPop (where the variable x stands for the current version number [nupop103.zip as of this writing]) nupopps.zip - ZIP archive of the PostScript documentation (only useful if you have a PostScript printer) connect.scr - Dialing script to connect to the UIUC terminal server with standard serial access slipdial.scr - Dialing script to connect to the UIUC terminal server with SLIP access (only necessary if NUPop is the only SLIP application you use)

Both the application and documentation are in ZIP archived format. If you do not already have a utility to unzip the files, a free file-extraction utility called UNZIP is available on the anonymous ftp host ftp.acns.nwu.edu in the directory pub/nupop (retrieve the file named unzip.exe). When extracting the ZIP archives, you may be asked whether you want to overwrite the files named connect.scr and slipdial.scr. Enter "n" for no so that the custom UIUC script files will not be overwritten with files for modem users at Northwestern University.

The authoritative anonymous ftp source for NUPop is ftp.acns.nwu.edu. Northwestern University maintains the most up-to-date versions of the software and documentation, including test versions of new releases. The NUPop application and related files are located in the /pub/nupop directory and subdirectories. You may download and use these files, but CCSO staff will not offer assistance on test versions of the software.

Installation, Documentation, and Support

The NUPop installation procedure may be difficult for network novices. However, an excellent installation manual and draft user's guide are available in PostScript format from the CCSO Resource Center and come highly recommended. If you do not have access to a PostScript printer, you can buy a spiral-bound, printed copy of the documentation for $5.00 from the CCSO Accounting and Distribution Desk, 1420 DCL. This manual should answer most of your questions about installing and configuring NUPop. If you need further assistance, try contacting your building or departmental network administrator. CCSO can also offer assistance with the proper configuration and operation of NUPop. Send requests for NUPop support by e-mail to nupop@uiuc.edu. Bug reports and suggestions should be sent by e-mail directly to NUPop developer Philip R. Burns at pib@nwu.edu.

More PC Eudora Facts and Features

Although the system requirements of PC Eudora are steep, the learning curve is not. Many of the commands in the File and Edit menus (as well as the mouse movements and keystrokes for managing individual windows) are identical to those in other MS Windows applications. The menus and commands unique to the Eudora application are largely self-explanatory and more or less identical to the Mac version.

Although PC and Mac Eudora are very similar, the Windows version offers a convenient feature not yet available on the Mac. An icon bar appears across the top of every open mailbox window, providing quick access to message management commands such as Reply, Reply All, Forward, Redirect, Print, and Trash. Once one or more messages are selected from the mailbox, clicking on the appropriate icon will produce the desired result (see Figure 2 on page ?).

Finally, PC Eudora takes full advantage of the Windows multitasking environment. As a background application, PC Eudora can check for mail regularly without any direct user intervention. For users who are heavily dependent on e-mail and want a package that will notify them as soon as new mail has arrived, PC Eudora is a good pick.

PC Eudora Flaws

PC Eudora is also not without shortcomings. Several options that appear in various menus have not yet been implemented, notably the Undo/Cut/Copy/Paste commands under the Edit menu (the standard Windows shortcut keys for these functions do work) and some of the switches under the Special menu. Also, at present, PC Eudora offers no support for modem users.

One of the most disconcerting aspects of PC Eudora is that it does not conform to the Windows Multiple Document Interface (MDI) standard. Every Eudora entity (the menu bar, each open message and mailbox, the configuration menu, etc.) exists in an independent window and appears in the Windows Task List as a separate task. If you open a message, for example, the opened message window may hide Eudora's main menu or the icon bar on the Mailbox window. Thus, it's often necessary to bring the application or a Mailbox window into the foreground in order to reply to or print an open message. In MDI-compliant applications, documents and other items generated by an application are all contained within a single window and constitute a single task. It is impossible to cover up the menu bar with a document created by the application. According to PC Eudora developers, this problem, which was a limitation of the software tool kit used to build PC Eudora, will be addressed in a future release.

How to Get PC Eudora

PC Eudora can be downloaded from the anonymous ftp server ftp.cso.uiuc.edu. The application and related files are found in the pc/pc-eudora/windows directory. The PC Eudora distribution includes the following files:

pcex.exe - Self-extracting archive of the PC Eudora application (where the variable x stands for the current version number [pce10.exe as of this writing])

README.TXT - A file containing basic information about the PC Eudora distribution files (if you download this file , it's a good idea to rename it because the archived application file [pcex.exe] also includes a file with the name readme.txt)

ftpvar.lst - A list of FTP Software Value Added Retailers located outside the US

wsocket.dll - An MS Windows dynamic link library required to run PC Eudora (this file is included with PC/TCP 2.11)

If you have PC/TCP version 2.11 or later, the only file you should need is pcex.exe. Place this file in its own directory (e.g. , C:\EUDORA) and decompress it by typing its full name. You should see an executable file named pcapp.exe and a text file called readme.txt. The readme.txt file gives further information for installation.

The authoritative ftp source for PC Eudora is ftp.qualcomm.com. The most recent production and test versions of PC Eudora are found in the pceudora directory on this host. CCSO will not provide help on alpha or beta versions of the software.

Installation, Documentation, and Support

Instructions for installing PC Eudora are given in the readme.txt file included in the self-extracting archive pcex.exe. As with NUPop, installing and configuring PC Eudora may be too difficult for beginners, especially if PC/TCP has not yet been installed. If you do not have experience modifying your autoexec.bat file, creating directories and plain text files, and installing new Windows applications manually, seek out the help of an experienced user or your network administrator. The readme.txt file describes the information that should be entered into the configuration window under PC Eudora's Special menu once the software has been installed.

Aside from the readme.txt file, there is, as of yet, no documentation written specifically for PC Eudora. However, PC Eudora and Eudora for the Mac are so similar that the Mac manual should answer most PC Eudora questions (except those regarding installation). The manual is available in PostScript format from the ftp host ftp.qualcomm.com (change to the pceudora/windows directory and download the self-extracting archive called 1_3EUMAN.EXE). Spiral bound copies of the Mac Eudora user's guide are sold at the CCSO Accounting and Distribution Desk, 1420 DCL, at $6.00 per copy.

Although CCSO is not yet officially supporting PC Eudora, staff members will do their best to help users with the installation and operation of the software. Send e-mail requests for help to eudora@uiuc.edu, and someone will get back to you as quickly as possible. Bug reports and comments or suggestions about the software should be sent directly to the PC Eudora developers at the e-mail address pc-eudora-bugs@qualcomm.com.

Date: Sat Jul 03, 1993 3:00 am PST
Subject: PC-Eudora and NUPop

[from Bill Silvert, 930703]

>[from Gary Cziko 930703.0600 UTC]
>
>The following is an article from a local campus newsletter which describes
>two e-mail programs for IBM PCs and compatibles which CSGnetters may find
>of interest. You will see below that the price is certainly reasonable
>(i.e., free).

Keep in mind that many of the features of these programs are already available in standard mailers. For example, I am using the free ELM software on my workstation which supports virtually every feature of Eudora and NUPop. From a PC or Mac one simply telnets to the work- station to send mail. I'm not saying that this is better, but it does save me the problem of installing additional software.
--
Bill Silvert

Date: Sat Jul 03, 1993 4:57 am PST
Subject: Misaligned herrings

From Greg Williams (930703) Bill Powers (930702.1530 MDT)

>Some thoughts about "aligning reference levels" that might prove
>useful:

>Suppose that four people undertake a task together. They are to
>lift a folded card table by its corners and hold it level.
>Someone else holds a hand out and says "hold it about here" and
>then removes the hand.

>Person A thinks of the goal as holding the table 3 feet above the floor.

>Person B thinks of the task as holding the table 5 feet below the ceiling.

>Person C thinks of the task as holding the table about level with
>his belt-buckle.

>Person D thinks of the task as holding the table about twice as
>far above the floor as the seat of a chair.

>Now, are these reference conditions "aligned?" Each person is
>controlling a different perception. Each reference condition is,
>of course, conceived of in the way relevant to that perception.
>Yet the result is for each person to want the table to be in
>about the same position. They will each agree that the table is
>in the right place in space when they finish.

>Now consider what can happen when conditions change. The four
>people are moved into another room and told to "do it again," but
>the ceiling in the new room is two feet higher and the only
>"chair" in the room is a footstool. The "aligned" goals are
>suddenly no longer aligned, and conflict will result. "You're
>holding your corner too low!" "No, you're holding yours too
>high!"

Apparently, A-D were initially told to adopt the (higher-level) goal of seeing the card table being held level, otherwise they wouldn't be upset about the tilt in room 2. I would call this alignment of a goal. In room 1, each individual sets a different sort of lower-level goal in support of the higher-order goal; these are obviously not in close alignment, but the higher-order goal is achieved. In room 2, each individual retains the same lower-level goal and finds that the higher-level goal (still in alignment with the higher-level goals of the others) cannot be achieved. If A-D are to achieve their higher-level goals (in alignment) in room 2, then they must alter their lower-level goals. Given altered environmental disturbances, lower-level reference signals of course must change to enable higher-level reference signals to be met. Pointing out non-alignment of low-level reference signals is a red herring which doesn't argue against alignment of goals at higher levels being necessary for consistent achievement of joint tasks.

As ever, Greg

Date: Sat Jul 03, 1993 7:55 am PST
Subject: Re: misaligned herrings

[From Bill Powers (930703.0830 MDT)] Greg Williams (930703) --

RE: Misaligned herrings

>Apparently, A-D were initially told to adopt the (higher-level)
>goal of seeing the card table being held level, otherwise they
>wouldn't be upset about the tilt in room 2. I would call this
>alignment of a goal.

Yes, I would too. There are two goals: hold it level, and hold it "about here." The "level" goal isn't the problem.

>In room 1, each individual sets a different sort of lower-level

>goal in support of the higher-order goal; these are obviously
>not in close alignment, but the higher-order goal is achieved.

The specified vertical position of the level table isn't a lower-order goal. The levelness and the vertical position are independent goals at the same hierarchical level. It isn't necessary to hold the table at a certain height in order to hold it level, or vice versa. The table actually has six static degrees of freedom: the x,y, and z coordinates of its center, and three rotational angles. Each of these can be controlled with respect to an arbitrary reference level independently of the others. None is subordinate to the others.

>Pointing out non-alignment of low-level reference signals is a
>red herring which doesn't argue against alignment of goals at
>higher levels being necessary for consistent achievement of
>joint tasks.

My point is that in any specific situation, the only way to find out whether goals are aligned is to see if all cooperating people can agree that the "common" goal is being achieved. If each one's perception matches that person's reference perception for the common task, there is no basis for suspecting that other people are actually controlling different perceptions. In the first phase of my example, such agreement would be reached. In the second phase, however, the difference in the four controlled perceptions (relating to height) becomes evident because of a change in background conditions that two of the people were incorporating into their definitions of the "common" perception. On moving to a new room, two of the people experience a change in the controlled perception, and two do not. This is what leads to conflict among people who had thought they were cooperating successfully.

Consider the "shared" goal of furthering the development of PCT. For some people on the net, the controlled perception is what other people in other disciplines think of PCT. Some would see the potential of PCT improved if PCT could be absorbed into the discipline in which they are already expert. Others would see furtherance of PCT if it became commercially successful, showing that people would pay money to use it. Still others would see furtherance in learning how to use PCT to get others to do what they want: learn faster, respond more favorably to advertisements, abide by social rules, become less violent or more caring, and so forth. Some see it in testing PCT against experiment to find out where it is wrong and how it can be improved. Some see it in connecting PCT to a set of very broad general principles of mathematics.

So we all agree that we are interested in the furtherance of PCT -- yet there are as many different perceptions and goals involved as there are individuals. Under some conditions, all of the people involved can experience their own perceptions as matching what they conceive of as the common goal. But that perception of happy agreement can quickly be disturbed when conditions change and reveal how different the individual perceptions and goals really are. Conversations on this network are concerned with little else.

Simply stating verbal agreement on a goal is far from guaranteeing that the parties to the agreement are controlling the same variables relative to the same reference levels, even when all the parties are satisfied that the group efforts are nicely aligned. Each person's perceptions are based on private experiences of different kinds, and must fit into a different hierarchy of goals and ways of perceiving. Each person assumes different facts, different meanings of terms. For every public description of a perception or a goal, there are myriads of unspoken assumptions and imagined conditions that were never explicitly mentioned by anyone yet which are an essential part of each perception. These are brought out only when conditions change, disturbing some people but not others, violating some conceptions of the goal but not others.

To reach effective alignment of goals, it's necessary for everyone to realize that this is not something that can be achieved through simple verbal agreement. It's necessary to keep varying the conditions, to look for changes that are seen as errors by some but not by others. The conflict that results is mute evidence that the goals are NOT aligned. Only by exploring so-called alignments with great skepticism can these differences be brought into the open, so that the range of conditions over which alignment is experienced can be increased.

People agree with each other much too casually. They assume that everyone else experiences the world as they do and makes the same assumptions they do. They find apparent agreement in one narrow circumstance, and then are shocked, hurt, and angry when the other person seems to violate the agreement. They seldom spend much time finding out what the other thinks is being agreed to, or bringing to consciousness their own background assumptions and desires to see whether they, too, are shared. Reaching agreement is actually very difficult and takes a long time, and much self-understanding. It requires testing the agreement under as many different conditions as possible. It requires repeatedly challenging the agreement to bring out unanticipated conflicts.

On an historical time-scale, human social institutions have the life of a May-fly. A few hundred years, and poof! Everyone realizes, or learns, that cooperation and alignment are better than competition and conflict. People keep trying to organize themselves into effective groups. But because people assume that everyone understands words in the same way, perceives the same world, and imagines the same unspoken assumptions, all groups eventually dissolve into rivalries and conflicts, and break up or change beyond recognition. It will be interesting to see whether PCT has in it any ideas that will keep its proponents together any longer than any other group has survived. Maybe all groups are temporary expedients.

Best, Bill P.

Date: Sat Jul 03, 1993 10:14 am PST
Subject: Practical PCT, /di/ /du/, Harvey Sacks

[From Rick Marken (930703.1100)]

I asked:

>Is there QUANTITATIVE evidence for the success of clinical applications
>of PCT that is any better than that for the success of any other

>clinical approach?

Ed Ford (930702: 1350) replies:

>You might ask those to whom I have taught PCT and have shown its
>practical applications in their areas of interest and who pay me large
>sums of money

I bet my old neighbor, Art Janov, is hauling down a bundle, too. Does that put "Primal Scream Therapy" (PST) on a par with PCT? And how about Dyer, Bradshaw, et al? Glasser's not doing too bad, either. I submit that money paid for services is not a good measure of the success of PCT or PCT loses hands down. And satisfaction with services is not a great measure either. David Koresh's followers (who survived) say they are still quite satisfied with the services he provided. Same is true for followers of Janov, the Maharishi, etc. People who pay a high price (in whatever form) for services very often are (or say they are) quite happy with those services.

I said:

>I expect these questions to be a disturbance to Ed Ford and Dag Forssell.

Ed (and Dag) said:

>Not at all.

Then why was there an action (you both replied to the post)? If there were no disturbance, there would have been no change in output (your output changed from not posting replies to my posts to posting replies)? A disturbance is not necessarily what an observer might judge to be something that has a negative effect on a controlled variable. I, for example, responded to the disturbance created by Dag's nice post about the biologist's book -- a "good" disturbance.

>Rick, get out of your ivory tower office and come on into the trenches
>and find out for yourself.

Unfortunately, I am no longer working in the ivory tower. I have spent enough time in the "trenches" to know that I don't want to be in them. In fact, I'm controlling for NOT being in the trenches. I don't want to be where the "rubber meets the road" either -- it stinks down there and its messy; I prefer the driver's seat (with a nice Mozart piano concerto on the CD) where it smells nice, it's clean and I'm in control.

You seem to have the idea that there are two worlds out there; a pleasant, serene rational one and a grubby, ugly, violent one where "the rubber meets the road". I think there is just one world out there -- represented as perceptual variables that can be in various states. I prefer my perceptual variables to be in the states I like them to be in. The perceptual states that you call "the trenches" are possible states of my perceptual variables; I have experienced those levels of my perceptions and I even spent a period during my youth when I was controlling for having some of those variables in "trench"-like states (hey, it was the 60s). I am not impressed by the lessons of controlling perceptions at "trench-level" reference states. I am also not terribly impressed by those who spend their time in the "trenches" (am I the only person in the world who finds Mother Teresa insufferably annoying?) helping the "downtrodden" instead of working to promote understanding (like the understanding of exponential population growth) of the reasons for the existence of those trenches in the first place.

>The modeling and theorizing of PCT are
>extremely valuable to me and others trying to apply these ideas, but
>everything we say about PCT is really worthless unless these ideas can
>guide us toward a better way of helping others live their lives more
>effectively, in a more satisfying way.

Well, here's where we completely disagree. I think the only thing of value is the modeling and testing -- the science is the only approoapriate measure of the worth of PCT. It's nice if you can help yourself and help others with PCT ideas. I think it's wonderful that you are using PCT in this way. What I disagree with is the idea that "practical success" is the real measure of the worth of PCT. If you measure success in terms of money, testamonials, number of adherents, etc. then PCT, even if it "catches on big" will eventually fade away like every other fad therapy -- there will always be a new therapy in a better wrapper that will work just as well (or better) in terms of the measures you propose.

The ONLY way to measure the worth of PCT is in terms of the science; the model either captures the nature of purposeful behavior, or it doesn't. When reality (rather than the marketplace) is the arbiter of the value of a theory, then the judgement is permanent. Newton's little laws (within the constraints in which they apply) will ALWAYS be right -- not because Lockheed Missiles and Space Systems "bought them" but because nature says they are right.

Dag Forssell (930702 1610)--

>Thus, PCT can be "sold" to advertisers based on its effectiveness
>on their bottom line.

Again, I am not against attempts to apply PCT; I am just against using the purported "success" of these applications as evidence of the worth of PCT. The only evidence of the worth of PCT is scientific evidence.

You may not know this but John B. Watson, the founder of Behaviorism, went into advertising after being kicked out of Johns Hopkins for having an affair with a student. Watson married the student and made a FORTUNE from his incredibly successful advertising campaigns. Watson was born into poverty and retired with more wealth than you or I together will probably ever see. Is that proof of the worth of Behaviorism?
B. F. Skinner raised two wonderful daughters, one of whom spent the first year or so of her life in what was basically an operant conditioning chamber. Both girls are quite successful and they loved and admired their father. Is this another proof of the worth of Behaviorism?

>Applications are the only reason
>for people to come to: "..if they want to learn it.."

If that were true, then why the hell did I want to learn it?

Again, its great if people want to apply PCT. I think that there may be personal benefits from learning PCT. But I think that using practical success as a measure of the worth of PCT will lead to great disappointment -- especially given the current level of development of PCT. Eventually, PCT may be able to be of clear, practical benefit -- demonstrably and reliably superior to other approaches. But that's not the case yet -- not even close.

Remember, Galileo's equations of motion, though correct, were of no use to the military people who had to figure out how to point their canons to get the canon balls to fall in the right place. Canon balls don't fall in the vacuum assumed by Galileo. Attempts to apply Galileo's laws would have produced a huge practical loss; but Galileo's laws were right; the military "laws of motion" were wrong.

Tom Bourbon (930702.1655)--

Very interesting post on /di/ and /du/. I spent most of yesterday afternoon saying these two syllables with all kinds of different mouth configurations -- the Gary Cziko test. I was able to recognize the /d/ quite well, even when it was articulated very differently in both cases. I think the Cziko test rules out analysis by synthesis (A-S) as even a plausible consideration. But I think A-S is also ruled out in terms of modelling too; how would you build an A-S model that recognizes /d/ when it is said by someone else, for example?. The input to the model is x. This x would presumably set off the "synthesis" routines, only one of which produces x as the result. But this is just the same as analysis by transduction. Look:

```
Synthesis model of recognition      Analysis model of recognition

        S1---> x                            A1 --> 20
x --->  S2 --> y                      x --->A2 --> 1
        Sn---> z                            An --> 2
```

In the synthesis model, the input sets off all the pheneme articulation synthesizers. The S1 synthesizer is the articulation pattern for /d/, S2 and S3 and the articulation patterns for other phomemes. Now, how does the system know which articulation pattern is the one what would have produced x? Presumably because it produces (in imagination) a perception that corresponds to the actual input. So the output of S1 is "imagined" x and it matches the input so the input is recognized. But suppose the input is y and the phoneme is still /d/ (as in di, du). Then S2 must also be a /d/ recognizer -- it is the same articulation pattern as S1 but it must result in y so y will be recognized as /d/ also. So analysis by synthesis is really just proposing a separate articulatory recognizer for each acoustic input that represents the same phoneme. In fact, S1 and S2 don't need to be there; all we need are the results -- x and y -- which are compared to the input. So A-S boils down to a template matching system; the synthesis part, when you actually start trying to build the model, is irrelevant.

The analysis model is a hell of a lot more elegant; the problem is designing the A1, A2...An analysis functions. I believe that these funtions are at the event level -- we don't necessarily have /d/ detectors. Anyway, the A1 system would put out a bigger signal than the others when an input came in that was /d/ like -- either x or y. A1 could just have S1 and S2 templates in it and put out a signal that corresponds to the degree of match to either template.

I think this is worth a little discussion at the meeting; I bet we get a lot of people walking around saying /di/ and /du/ with tongue in cheek. Hmmm. "Tongue in cheek" -- a good name for the "Cziko test".

Kent McClelland --

It was Harvey Sacks -- the Bob Dylan of sociology. He had a great act and I really enjoyed it. But I'm afraid I was the way I am even before meeting him.

Best Rick

Date: Sat Jul 03, 1993 1:40 pm PST
Subject: CELLS - SYSTEMS - RKC

[From Bob Clark (930703.0530 EDT)] Bill Powers (930621.0840 MDT)

You remark:

>...I don't think any cells in the body are merely passive participants.
>Kidney cells control independently for concentrations of many substances
>in the blood by varying the rate of their removal; one kidney cell
>contains many control systems.

I think your generalization is too broad. Certainly there are many physiological control systems, some, at least, seem to involve more than one organ (identifiable group of cells). Several cells within the organ may be involved, possibly as individual cells, possibly as groups of cells. Some cells within the organ may "merely" have structural functions. These are distinguished by Oncologists (secretory vs connective tissue) both regarding diagnosis and treatment. In addition, an organ may operate as one component of a control system, perhaps an Output Function. Even groups of cells, lacking identifiable means for transmission of signals, are not necessarily control systems. The operation of the Venus Fly Trap can be described without need for control systems. Coelenterates (from the Greek, "coel," cavity) are described as having some sort of "pouch" into which samples of the surroundings are drawn and later expelled after absorbing nutrients. This could be a sort of relaxation oscillator, with no need for a control system.

Signals can be transmitted by surrounding fluids, circulatory systems, and nerves, acting as specialized transmitting systems. An organism without some way to transmit signals from one part to another surely cannot act as a negative feedback control system.

In any case, the body as composed of cells is not a good example of a "social control system." The individual cells are constrained by their internal structure and/or their surroundings. They cannot "learn," they cannot reorganize, they have no mechanism for memory! To the extent that they may operate individually as control systems, they appear to be systems of a single level.

To consider a single cell as an independent set of control systems calls for careful application of The Test, particularly with regard to possible alternative explanations in terms of physiological and biochemical actions, such as osmotic pressure, semi-permeable membranes, chemical reaction equilibria, etc. As you point out, in other posts, a ball "seeking" the bottom of a bowl is not an example of a control system.

Speaking of The Test -- not long ago, perhaps a few weeks, you summarized The Test in the form of _four questions_. I thought them highly pertinent and useful. Otherwise the only definite statements of The Test I have found are in BCP, pp 47, 53f, and 232f. These are directed exclusively to the experimental application of a Test Disturbance. But I think the additional questions are essential.

I missed downloading them at the time, and I think they are well worth repeating, perhaps even as a separate Post.

I will post other topics separately, for convenience of referral.

Regards, Bob Clark

Date: Sat Jul 03, 1993 1:44 pm PST
Subject: Re: Aligning Goals

From Kent McClelland (930703)

Bill Powers (930702.1530 MDT) (930703.0830 MDT) Greg Williams (930703)

Your table example, Bill, is an interesting one, and I'm heartily in agreement (as I perceive it!) with your description of the difficulties inherent in maintaining cooperative groups. I'm pleased that both of you see some merit in talking in terms of alignment of perceptions or goals, which seems to me a crucial concept as we try to develop a PCT sociology.

I think that Greg's right in saying some higher level goals must be involved in the table example. The higher-level goal that seems to me most relevant is a self-conscious commitment to cooperation in the project of lifting the table. The point of the example, as I read it, is that a perception of cooperation with others cannot be maintained simply on the basis of plans or some announced intention to cooperate. Like other perceptions it must be controlled by resisting disturbances, in this case the disturbance of moving to a room where the "background conditions" have changed. The conflict Bill describes ("You're holding your corner too low!") is, I would say, an attempt (maybe a crude one) to reestablish a perception of cooperation by verbally checking the reference levels involved. Bill, of course, makes this and similar points more fully in his second post (on the difficulties of cooperation). It strikes me again that many of the things Bill says are similar to the objections that ethnomethodologists have raised against conventional sociology.

This exchange has helped me to see that I've been using the term 'alignment' a little too glibly. It doesn't help much to say that several people's reference levels are approximately aligned without saying what levels of perception we're talking about and what are the higher-level intentions of the people involved. There's a difference between "cooperative alignment," where the participants consciously coordinate their goals, and "competitive alignment," where everybody has the "same" goal but that goal involves exclusive and individual use of some piece of the environment, e.g., the freeway at rush hour. Both kinds of alignment may have important social consequences, but the consequences of course are different. Yet another kind of alignment I see is "cultural alignment," where people share the same goals or reference levels or perceptual organization simply by virtue of belonging to the same cultural group. Language would be one example of cultural alignment, but it also might include things like "institutionalized racism." I guess the anthropological concept of 'culture' implies alignment. Things cultural are by definition shared.

Enough musing for the weekend!

Kent

Date: Sat Jul 03, 1993 2:05 pm PST
Subject: Toward better alignment?

From Greg Williams (930703 - 2) Bill Powers (930703.0830 MDT)

>My point is that in any specific situation, the only way to find
>out whether goals are aligned is to see if all cooperating people
>can agree that the "common" goal is being achieved.

Goals also can be aligned sometimes if all cooperating people agree that the "common" goal is NOT being achieved. Like failing to move a big rock to a particular place (agreed upon by all) because the rock was too heavy for the folks to lift in unison. But your basic point is well-taken.

>To reach effective alignment of goals, it's necessary for
>everyone to realize that this is not something that can be
>achieved through simple verbal agreement. It's necessary to keep
>varying the conditions, to look for changes that are seen as
>errors by some but not by others. The conflict that results is
>mute evidence that the goals are NOT aligned. Only by exploring
>so-called alignments with great skepticism can these differences
>be brought into the open, so that the range of conditions over
>which alignment is experienced can be increased.

Yes. This is what business managers and families contend with every day. Tests for goal-alignments can be based on either really or hypothetically altered conditions: "I've decided that from now on, we'll do it like ___" vs. "Would you have a problem with ___?" It is no great surprise that much of people's conversations are filled with informal attempts to apply The Test to others -- "What do you think about ___?" -- to help avoid surprises about what their reference signals are.

>People agree with each other much too casually. They assume that
>everyone else experiences the world as they do and makes the same
>assumptions they do. They find apparent agreement in one narrow
>circumstance, and then are shocked, hurt, and angry when the
>other person seems to violate the agreement. They seldom spend
>much time finding out what the other thinks is being agreed to,
>or bringing to consciousness their own background assumptions and
>desires to see whether they, too, are shared.

In my experience, this characterizes some people I've met, but by no means all, or even an overwhelming percentage. But your sample must be different. That's certainly a problem with statistics, isn't it? (ONE problem with statistics!)

>Reaching agreement
>is actually very difficult and takes a long time, and much self-
>understanding. It requires testing the agreement under as many
>different conditions as possible. It requires repeatedly
>challenging the agreement to bring out unanticipated conflicts.

Again, I think many, many individuals have recognized this for a very long time. How long have professional attorneys been around?

>People keep trying to organize
>themselves into effective groups. But because people assume that
>everyone understands words in the same way, perceives the same
>world, and imagines the same unspoken assumptions, all groups
>eventually dissolve into rivalries and conflicts, and break up or
>change beyond recognition.

More descriptive statistics. My own sampling suggests that there are many other reasons besides naive assumptions about others' perceptions leading to dissolution of groups. Of course, "eventually" is a long time, so if a group hasn't had problems of other sorts, you could argue that it will sometime have the kinds of problems you describe. But perhaps some (even many) individuals are less naive than you allow. Time to collect more data, I suppose.

As ever, Greg

Date: Sat Jul 03, 1993 7:55 pm PST
Subject: Equivalence instead of alignment

[From Bill Powers (930703.2100 mdt)] Kent McClelland (930703)

>I think that Greg's right in saying some higher level goals
>must be involved in the table example.

I wasn't arguing with him about that; in the background of my example there must be, as you say, a higher-level intention to cooperate. I was pointing out that the spatial position of something isn't a lower-level aspect of its orientation: those are variables at the same level, since neither depends on the other. Setting any of these reference levels does, of course, depend on the goal of cooperating, and perception of cooperation depends on perceiving lower-level control processes that are working harmoniously together. So that is a hierarchical relationship.

>It doesn't help much to say that several people's reference
>levels are approximately aligned without saying what levels of
>perception we're talking about and what are the higher-level
>intentions of the people involved.

I think I'd rather do without this concept of "alignment" altogether. Instead, I propose that we talk about _equivalent_ perceptions and _equivalent_ reference levels. Perceptions and reference levels are equivalent if one person brings the environment to a state perceived by that person as satisfying that person's goal for it, and another person experiences a perception at its reference level as a result. The equivalence is strengthened if one person acts to resist a disturbance of a perception, and the other person agrees that there was a disturbance and that it was successfully counteracted.

In my example, a disturbance applied directly to the table would be resisted by all four of the people even though they are controlling different perceptions. The equivalance of their perceptions and reference levels would extend to many kinds of disturbances. However, the equivalence is limited because of the unshared aspects of the perceptions. If a footstool is substituted for the chair being used by one person as part of the definition of the table's vertical position, that will disturb that persons's perception, calling for a lowering of the table, but not the perceptions of any of the others, who will resist the lowering. So the equivalence does not include that kind of disturbance.

Both you and Greg pointed out that when what I now call a "failure of equivalence" is discovered, through the unexpected appearance of a conflict, this may well lead to communication and an attempt to discover why the conflict developed. The outcome of the communication, if successful, will be that one or more persons will reorganize a perception or reset a reference level, so that the perceptions become equivalent again under the new conditions as well as the old. For example, the person using the chair as the definition of height might be able to think of some compelling reasons for defining the task in terms of that perception, and the others might agree to adopt it. This changes the definition of one of the tasks to "Keep the table at twice the height of the chair seat above the floor." If all the participants are now controlling this perception, the previous conflicts will disappear, indicating that equivalence has been achieved. The equivalence now extends to rooms with different ceiling heights, belts that have a tendency to sag toward the floor, and even the capricious substitution of footstools for regular chairs.

Communication in words isn't essential for establishing equivalent perceptions and reference levels, but interaction is necessary. The signal of nonequivalence is the appearance of conflict when no conflict is intended. Two people stacking lumber according to grade can discover a nonequivalence in the perception of grade when they find themselves pulling in opposite directions on one piece of lumber. If one person assumes that the other knows more about lumber grades, the first little tug may be enough to signal the need to revise the perception of grade, so that person immediately defers to the other, and that disagreement doesn't arise again -- all without a word being spoken. The perceptions still may not be the same, but with respect to which pile in which to place a piece of lumber they are equivalent. Apprentices probably learn most of what they know in this way from their working supervisors.

The reason I want to say "equivalent" instead of "aligned" is that the concept of alignment begs a question that nobody is in a position to answer: are my perceptions like yours even when we agree we are perceiving the same thing? Equivalence is violated often enough that the most probable answer is "no." And as a practical matter, actual alignment of perceptions and reference levels isn't required; all that is required for cooperative action and effective communication -- within a limited domain -- is equivalence. Equivalence can be achieved even when perceptions are radically different.

Best, Bill P.

Date: Sun Jul 04, 1993 12:39 pm PST
Subject: Batting practice

[From Bill Powers (930704.1015 MDT)] Greg Williams (930703 - 2) --

>Goals also can be aligned sometimes if all cooperating people
>agree that the "common" goal is NOT being achieved. Like failing
>to move a big rock to a particular place (agreed upon by all)
>because the rock was too heavy for the folks to lift in unison.

If the the people never even get started toward achieving the goal, how will they ever discover whether they had different goals? To say that the goal is "agreed by all" doesn't mean that they all understand it the same way. They agree on the verbal description of the goal, but the real test comes in carrying "it" out. That's when "it" can easily become "them." Of course if everybody TRIES to move the rock, at least you can claim that their individual goals of TRYING to lift the rock were equivalent, provided that all of them were actually lifting. Maybe some of them just wanted to see what was under it, but that situation never arose since they couldn't move the rock at all.

If they actually lifted it and moved it, the effort could easily break down over the destination, as one "co-worker" after another said "Not in MY back yard."

>>They seldom spend much time finding out what the other thinks
>>is being agreed to, or bringing to consciousness their own
>>background assumptions and desires to see whether they, too,
>>are shared.

>In my experience, this characterizes some people I've met, but
>by no means all, or even an overwhelming percentage. But your
>sample must be different.

I guess I disagree. I see people reaching agreements all the time, being satisfied when their words match but not pursuing the matter of what each person meant by the words. Even intellectually sophisticated people are often surprised, after reaching a seeming agreement, to find that the other person immediately does something in violation of it. Sometimes this leads to a resolution, but often it simply devolves to an argument over who was right. I don't find that intellectuals are any less prone to that kind of argument than other people.

Oh, well. I'm just trying to pitch some ideas about group efforts. If you want to treat this as batting practice, that's up to you.

Best, Bill P.

Date: Sun Jul 04, 1993 12:46 pm PST
Subject: Flytraps, kidneys, etc.

[Mary Powers 930704] Bob Clark (930703) says

>I think your (Bill's) generalization (about PCT) is too broad.

But this is the point: the broad assertion that every living thing is a control system. It is an audacious challenge, and the point of it isn't to say "what about coelenterates, or Venus Fly Traps", but to look at these creatures, and at the cells that compose them, and kidneys, and whatever, from this new perspective. A kidney cell may look like it simply has an output function, in the context of the organ of which it is a part. But the cell itself lives an active metabolic life, maintaining itself as a kidney cell. Its output, along with all the other kidney cells, rids the organism of substances that would be poisonous, but that overall effect is a side effect of what the kidney cell itself is controlling for.

>The operation of the Venus Fly Trap can be described without
>need for control systems.

So can the operation of human beings. But is it a good description? PCT says no, look again. How does apparent S-R behavior really work?

Ditto relaxation oscillators in coelenterates. I don't know what a relaxation oscillator is, but to say cells operate as though that is what they are begs the question of how they go about acting that way.

Since nobody on the net is an invertebrate physiologist, we'll be wondering about that for a while. But I think the question "suppose it's control systems all the way down?" is more interesting than saying "it's only a .... ".

Mary Powers

Date: Sun Jul 04, 1993 1:45 pm PST
Subject: Kudos for Practical PCT, Fly Traps

[From Rick Marken (930704.1400)]

I called the Powers' last night to talk about travel to the conference and, as an unpleasant side-effect, learned from Bill that my posts on "Practical PCT" sounded (to him, so probably to others too) like I was belittling the work of Ed Ford, Dag Forssell and others who are trying to apply PCT. So let me try it again, from the top:

The MOST IMPORTANT, THE MOST VALUABLE, THE MOST WORTHWHILE WORK on PCT is being done by those who are trying to apply it (honestly and acurately) in practical situations. This is the work of Ed Ford, Dag Forssell, Hank Folson, etc -- the PRACTICAL PCTers. The work of the scientific types is valuable in a different way.

For reasons that have to do with problems I have encountered in my very limited dealings with nonPCTers interested in the practical value of PCT, I have been arguing against the use of "standard" measures of "success" as a basis for evaluating the practical value of PCT -- that's all. There is no question in my mind that clinical practice, management practice, business practice, etc can be improved enormously by the application of PCT. I just think that typical measures of the success of these practices will NOT produce convincing, reliable evidence of the superiority of PCT. I would prefer that the practical value of the application of PCT be "backed up" by the scientific evidence. And I think that is what is typically done. Ed, Dag, Hank and other practical PCTers participate in the science (for example, by attending CSG meetings) so that they can present the science as evidence of the value of their approach. So they are already doing the "right" thing (from my perspective, of course). Perhaps I reacted strongly to Hank's post that set this thing off because I didn't want the teensey little contributions of the scientists to be ignored as part of the efforts to produce a convincing piture of the practical value of PCT.

I know that Practical PCT works -- from a personal perspective and from successful applications that I know of. So please don't take my earlier comments as anything other than a suggestion (probabaly not worth the bits it is typed with) regarding strategies for advertising PCT applications.

Bob Clark said:

>The operation of the Venus Fly Trap can be described without
>need for control systems.

Mary Powers (930704) replies:

>So can the operation of human beings. But is it a good description?
>PCT says no, look again. How does apparent S-R behavior really work?

Yes, indeed. This harks back to the earlier discussion of the apparently "open loop" moth drop. I don't think that one has to be an entymologist to see that the behavior of the flytrap involves control. All you have to see is that the apparent stimulus for flytrap behavior is affected by that very behavior -- there is a closed loop. The fly affects some sensor on the flytrap (this is where the entymologist comes in; s/he could tell us what sensor). The result is a response -- closing the trap as well as (yuck) the secreting stuff that dissolves the fly (I suppose). So the response affects the ultimate cause of the response -- the sensory effect of the fly. This is a slow loop, but it is a loop nonetheless, and the sign of the feedback is almost certainly negative (the response reduces the effect of the cause of the response). So the flytrap is controlling the sensory effects of the fly (the disturbance). Flytrap behavior is the control of perception; now the entymologists can tell us which perceptions are likely to be controlled -- and how.

Best Rick

Date: Sun Jul 04, 1993 4:06 pm PST
Subject: Teambuilding & goal alignment

Bill Cunningham (930704.1950)

There is more to collective goal achievement than aligning goals, and it falls under the general heading of "teambuilding". One of the criteria that distinguishes a team from a bunch of folks with a common objective is that the team is collectively responsible, _to an external standard_ for the result. Mother Nature can provide the standard, as for a mountain climbing team.

The real point is that the external standard and collective responsibility force subordination of all the individual goals. Those who have played or worked on teams will understand this ethos exactly. If you are looking for literature on the subject, try "leadership". You won't find PCT mentioned, but you will find the discussion of "team chemistry" to be entirely consistent with a PCT description of successful collective behavior. It wouldn't take much effort to call teambuilding (an essential leadership skill) a practical application of PCT.

As a thought for Independence Day, teambuilding involves giving up a degree of indepedence and generating interdependence.

Bill C.

Date: Sun Jul 04, 1993 5:44 pm PST
Subject: Fielding practice

From Greg Williams (930704) Bill Powers (930704.1015 MDT)]

GW>>Goals also can be aligned sometimes if all cooperating people
GW>>agree that the "common" goal is NOT being achieved. Like failing
GW>>to move a big rock to a particular place (agreed upon by all)
GW>>because the rock was too heavy for the folks to lift in unison.

>If the the people never even get started toward achieving the
>goal, how will they ever discover whether they had different goals?

They won't. In my example, I said that the folks "failed." They DID try, and they found that the rock was too heavy to lift in unison.

>To say that the goal is "agreed by all" doesn't mean that
>they all understand it the same way. They agree on the verbal
>description of the goal, but the real test comes in carrying "it"
>out. That's when "it" can easily become "them."

Correct. I didn't say: "Goals MUST be aligned..."; I said "Goals... can be aligned sometimes..." You sound like you're trying to pick a fight with YOUR misreading of my words. Nothing better to do on the 4th?

BP>>They seldom spend much time finding out what the other thinks
BP>>is being agreed to, or bringing to consciousness their own
BP>>background assumptions and desires to see whether they, too,
BP>>are shared.

>>In my experience, this characterizes some people I've met, but
>>by no means all, or even an overwhelming percentage. But your
>>sample must be different.

>I guess I disagree.

That our samples are different? I don't see how they could be the same.

>I see people reaching agreements all the time, being satisfied when
>their words match but not pursuing the matter of what each person
>meant by the words.

"All the time," eh? When do you sleep?

>Even intellectually sophisticated people are often surprised, after
>reaching a seeming agreement, to find that the other person
>immediately does something in violation of it.

Ah, statistics. How about quantifying your "often" a little? Are your data on-disk? If so, I could run them through a neat program I have to determine their significance.

>Sometimes this leads to a resolution, but often it simply devolves to
>an argument over who was right.

"Sometimes" and another "often." I could run those data, too.

>I don't find that intellectuals are any less prone to that kind of

>argument than other people.

Hey, I agree. I don't find that, either, because I didn't gather the necessary data and analyze it.

>Oh, well. I'm just trying to pitch some ideas about group
>efforts. If you want to treat this as batting practice, that's up to you.

I thought it was fielding practice for the pitcher. NOW you tell me!

As ever, Greg

Date: Sun Jul 04, 1993 6:02 pm PST
Subject: BOOT CAMP SYSTS - RKC

[From Bob Clark (930704.9:45 pm EDT)] Bill Powers (930621.0840 MDT)

You remark:

>"Tasks not well done" does not mean the same as "Tasks doomed to
>failure," which is what I was talking about. I don't mind joining
>in tasks that are being done suboptimally. It's the other kind I
>have a hard time getting enthusiastic about.

I see what you mean. It's something like the way I feel when someone is trying to interest me in some form of perpetual motion machine.

>We're close indeed about social control systems.

Yes, good, I think so, too. In fact, I think we are very close in most areas, but sometimes we come from different viewpoints.

>... I've been
>thinking lately about how social cooperation toward controlling
>variables of mutual interest (how's that?) is actually carried
>out. I was musing about a few halcyon weeks in Boot Camp (now you
>have me doing it: I mean boot camp) in the Navy when I suddenly
>became a 17-year-old platoon leader with about 40 people to march
>around in the compound, practicing close-order drills.

{My use of capitals is intended to improve clarity and/or emphasis. Does it work?}

Your amusing report from bOOT cAMP (HA!) is a good illustration of the way I look at these interactions. Each participant has his own view of the situation, but they each include enough over-lapping similarity to result in the observed interactions.

>All this does give some notion of how a social control project
>could be set up to work. The control isn't arbitrarily imposed
>the way it is inside a person. It requires all the people
>involved to understand most of what is going on, at least to the

>extent that it affects something they can perceive and control by
>themselves. It was up to me to think up the pattern and the
>strategy in the marching drill, but it was up to the individuals
>to make it work.

Of course it was all within the larger framework of the government and military, as represented, perhaps differently, in each individual's memory.

How to "set up" a "social control project?" But we already have very large numbers of such projects in operation right now. These consist of families, businesses, social clubs, governments, etc. I am now observing in detail a multi-level "project" in the form of the Government of the city of Forest Park. In addition to the one or two level social system of the close-order drill, at least five or six levels are in operation in the City organization. The various individuals and their own goals, structures, purposes, etc may be considered separately. It is fascinating to see the various orders of control in operation as they are outlined in our analysis of the organization of individuals.

>Also, the speed of control is very slow in such a group effort.

Notice the time scales and control of temporal variables required in the following descriptions.

>The marching cadence and the fixed meaning of the commands,
>thoroughly understood by each individual and translated reliably
>into shift of directional reference levels, left a minimum time
>resolution of four steps, about two to three seconds depending on
>the pace (one thing that the platoon and I often disagreed about;
>I ended up following them on that point). The commands had to be
>issued over a period of at least four counts in advance of the
>actual execution, even so, to allow for processing time. "First
>column tatharear HAR!" takes four steps, and is executed on the
>fifth step (assuming I started on the correct foot for the
>direction of pivot). Some commands like "First column column-
>half-left MARCH" were too long to do in four steps and had to be
>started another step in advance to be understood.

In the following, memory is continually and repeatedly involved, to recognize the command, to select and activate the reference levels required to produce the desired results. Selections among interpretations of the audio signals and selections among projected effects of possible sets of reference levels are equally involved.

>.... At each word,
>all the people had to (a) recognize that they were in the group
>to which the command applied, (b) interpret exact technical
>meaning of the term and turn it into a reference level for an
>actual physical act, and finally (c) wait for the MARCH command
>and execute the procedure already set up with a delay of one step
>(if correctly given) or two steps (if issued on the wrong foot,
>blush). At the same time, each individual was maintaining a
>cadence in synchronism with the nearest visible people,

>maintaining spacing and line straightness, and imagining the move
>that would result in the correct perception when the MARCH
>occurred. Furious mental activity in each person.

As I conceive these events, the DME is involved in all these selection processes. That is, a set of current perceptions is compared to a set of similar memories. The memory with the closest match is selected (by the DME) for comparison with memories of activity. Some combination of these memories is selected (by the DME) to use for sets of reference signals. This combination is then "activated" (by the DME) to produce the selected activity.

This seems to be a rather lengthy description of events requiring, at most, only a few seconds. However, in larger scale, more complex social situations, days, months, etc may be needed.

When social systems are expressed in terms of the participating individuals, a very large number of variables is involved. These are likely to include several Orders of Control. This is one situation where I think that my suggested definitions for some of the Orders may be helpful.

Thus, to the commander, when selecting his commands, the soldiers were essentially "objects" to be moved around. In addition, the commands had to have a specific sequence, and be issued with a suitable timing.

Thus we have Control of Sequence Variables and Control of Temporal Variables. These can be combined to form Control of Skills (here to be considered "Muscle Skills), that is, those skills needed to produce Close-Order Drills. Up to this point, the men are essentially objects to be controlled.

In your report of the pacing of the commands (above), you note that you and the platoon disagreed -- here you illustrate control of "Interpersonal Relationships" both for yourself and each of your men. You treated the men as independent (more or less) entities, and worked out a compromise. You worked with them as a "Cooperating Group." An alternative "Interpersonal Relationship" might have been a "Leader-Follower Group:" that of a commander giving "orders" to his subordinates.

Yes, Bill, of course there are many forms of Relationship -- but I think that category is much too broad to use here. And it seems to me that "Interpersonal" alone, without modifier, is also too broad -- but perhaps it would be a better name for the category. Even if restricted to those variables requiring more than one individual, "Interpersonal Relationships" is very broad. I find I tend to think of sub-categories like, "Games," "Strategies," "Leader," "Follower," etc etc. Yet I think these are very meaningful and useful concepts. How about it?

>I'm impressed with how much practice is needed to create a
>precise group control process, and how simple it has to be to work at all.

It is interesting to observe the relation between time and precision -- this seems to be a very general trade-off: the MORE ACCURACY the SLOWER THE RESPONSE, and vice versa.

Others on the Net also seem to see the situation similarly. For example:

Kent McClelland (930622. 8:29 pm EST)

>I'm willing to concede Tom and Rick's point that each human control system
>has its own reference levels, and that no super-ordinate social-control
>system exists to impose reference signals on independent individuals. Two
>individuals cannot literally share the same reference level.
>.....
>For "practical purposes" two people's reference levels can be the
>same--good enough for government work, in the case of Bill's Naval
>comrades.
>.....
>The term I like to use to designate this practical similarity of reference
>levels is "alignment"--a concept that I think is key for doing any kind of
>social analysis with PCT.

>Several interesting questions arise when we think about social interaction
>in terms of alignment:

>1. How do people succeed in getting their reference levels into alignment?

>2. How good is good enough when it comes to alignment?

>3. What are the consequences when lots of control systems, not just two,
>are effectively aligned in some coordinated effort?

"ALIGNMENT" -- an interesting concept. Perhaps you have others to offer.

And these are very interesting questions, Kent -- and there are many more, especially when "lots of control systems" are "coordinated" into becoming an organized group. In this case, it seems to me that such a group tends to establish some form of hierarchy of control, much like the government of the City of Forest Park that I'm currently observing.

Regards, Bob Clark

Date: Sun Jul 04, 1993 7:06 pm PST
Subject: Cooperative Aligning

from Ed Ford (930704:2000) To Greg, Kent, and Bill ....

Speaking of coopertive alignment, several years ago I was consulting in several businesses where I was able to put this into practice. I consulted in a carpet mill in Canada and the shipping department, consisting of a foreman and six or more tow motor drivers, was making two mistakes per day shipping carpets by truck to customers. The drivers were very independent, very territorial, and very uncooperative. I met with the foreman and suggested a plan. He would meet with the drivers every morning for about 15 minutes. He created a chart comparing the number of mistakes per week. He asked the drivers if they felt this was satisfactory and they all agreed it should be better. He asked them if they would meet with him every morning for the first 15 minutes of the shift, review the prior day's performance with him, and offer any suggestions for improvement they might have and any improvements or successes they noticed in the operations. He had coffee and donuts available to them during the meeting.

After two months of daily meetings, two important changes occurred. First, the number of mistakes went from two a day to two per month. The second was that the drivers began to cooperate more with each other, lending a hand when needed (two serviced the manufacturing floor, two storage, and two shipping). It might be noted that the chart was updated on a daily basis and was referred to at each meeting.

The second example had to do with housekeeping department in a hotel. The eight to ten women who cleaned the rooms were averaging 42 minutes per room, and the acceptable standard was 30 minutes per room. The head housekeeper and I worked out a plan where she would meet with her staff at the beginning of the shift (about 8 a.m.) and go over a chart which compared the average amount of time per room for each day. They had all agreed to work toward the hotel management's goal and thought it was reasonable. The ladies would come into the little office of the head housekeeper, look at the chart on the wall to see if they had "gone over the hill" which meant were they over 30 minute per room goal. Within two weeks, they were at or below 30 minutes and stayed consistently at this over the years. We also devised a chart for maintaining the quality of the cleaning of the room and found success in this area as well. Again, the ladies became more cooperative with each other, reporting off of work when sick (which had been a problem).

In both the above cases (and there are others), the perception of the group both of each other and their attitude toward their supervisor seemed to improve. They all seemed happier and more cooperative. The hotel manager was so happy he threw a pizza party every Friday after work for housekeeping. When I found out that the front desk was having problems, we created the same kind of feedback chart with daily meetings for them.

I think that the daily meetings allow for continual updating of goals and the controlled variable or feedback. In both the above cases, they were all dealing with "how they were doing." In the first case, although they serviced different areas of the plant, they became much more cooperative. In the second, if one of the housekeepers got finished early, she'd help someone else.

Best, Ed

Date: Mon Jul 05, 1993 6:47 am PST
Subject: Replying to CSGnet

[from Gary Cziko 930704.1822]

When I reply to a message that has come to me via CSGnet, I expect that my reply will also go to CSGnet, and not just directly back to the person to whom I am replying. But I have learned that this is not always the case.

In replying recently to messages sent by Tom Bourbon and Bill Silvert via CSGnet, my replies have gone back directly to them because their own personal e-mail address is included in their "Reply To" header and not CSGnet's address, as is more normally the case.

I do not know why Tom's and Bill S's messages are different in this respect and will try to find out. In the meantime, it is always good practice to double-check the "To:" field of your message before sending it, especially if you are using an automatic reply function. While having a message intended for CSGnet go to just one person is no disaster and can be sent again to the net, having a private message show up on a network can be embarrassing, to say the least. Checking your "To:" field before sending is the only way to make sure that your message will get to where you intended.--Gary

Date: Mon Jul 05, 1993 6:47 am PST
Subject: Re: Equivalence instead of alignment

[from Gary Cziko 930705.1430 UTC] Bill Powers (930703.2100 mdt) said:

>I think I'd rather do without this concept of "alignment"
>altogether. Instead, I propose that we talk about _equivalent_
>perceptions and _equivalent_ reference levels. Perceptions and
>reference levels are equivalent if one person brings the
>environment to a state perceived by that person as satisfying
>that person's goal for it, and another person experiences a
>perception at its reference level as a result. The equivalence is
>strengthened if one person acts to resist a disturbance of a
>perception, and the other person agrees that there was a
>disturbance and that it was successfully counteracted.

This analysis appeals to me, but wouldn't it be better to use the term "functional equivalence" or use some other qualifier? "Equivalence" all by itself implies to mean that the reference levels really are in some "objective" sense the same (as two thermostats of the same manufacturer and model set to the same temperature). "Functional equivalence" implies instead that the two reference levels may not be the same but that in a given environment they function equivalently.--Gary

P.S. Notice how two different thermostats may also be functionally equivalent--they may be set to the same temperature, but one may use a bimetal strip and mercury switch and the other a thermistor and so they are really controlling quite different sensory readings. But they can work together anyway if set to the same temperature.

Gary Cziko

Date: Mon Jul 05, 1993 7:51 am PST
Subject: goal alignment, teambuilding & harmony

Bill Cunningham (930705.1115)

One of the features of teams that I failed to mention was that team members have distinct roles with respect to team function. This implies individual goal sets that contribute to the collective goal.

Ed Ford (930704.2000) provides two excellent examples of teambuilding, both subtle as viewed by the team members. The case of the housekeepers involves a number of team members with essentially the same role (except for different room assignments). The drivers had three sets of roles, and therefore quite different job related goals, in addition to their individual differences. One of the functions of teambuilding is to foster an appreciation for the contribution of the other team members, and their individual goals.

The term "harmonization" is used among international workgroups. This term recognizes there will be very different national goals which may be reflected in the instructions to workgroup members. The leadership task is to find subsets of goals that are "harmonious", in turn identifying areas where real progress is likely.

I've not seen this term used in the literature of teambuilding, but think it is a good one. It solves Gary Cziko's terminology problem over thermostats with functional equivalence, and it admits to constructive functional nonequivalence.

I also failed the degree of mutual trust that goes into team "chemistry." The requirement to development this trust leads to activity that appears to be unrelated to the team goal. Ed Ford's examples include some of this. From a PCT perspective, I think Martin Taylor's layered protocols apply, especially the bottom line that communication is the control of the sender's beliefs about the recipient.

We can also point to the situation where the warm fuzzy of team harmony replaces the original team goal. That's when everybody sits around agreeing with each other, or addressing only topics for which there is no obvious need to resolve conflict. This beautifully (and hilariously) described in videotape put out by a U Texas sociologist. Title is "The Abilene Experience", and it is the origin of the expression "We're on a trip to Abilene." Can't remember the author.

Bill C.

Date: Mon Jul 05, 1993 1:14 pm PST
Subject: DETECTORS - RKC

[From Bob Clark (930705. 4:55 PM EDT)] Bill Powers (930621.0840 MDT)

Your discussions of "contrast" (per Bruce Nevin (Mon 930621 09:44:18 EDT)), "phonemes" and "differences" call for corresponding kinds of "detectors."

I agree, but the question remains, how does such a detector work? As I see it, there must be some way for the incoming set of perceptions to be compared with some kind of "reference set" of perceptions. That is, the incoming set is compared with remembered perceptions.

This can be demonstrated with existing hardware used for "pattern recognition." There are optical, and magnetic, character readers that work very well. These devices are passive and do not use feedback systems. They also do not need any form of spectroscopic analysis or other analytic equipment. A related concept is sometimes used, I think, for analysis of star fields (correct me, Bill, if I am wrong here). This uses a reference photograph and a "new" photograph with the viewing device alternating between the two. This is a remarkably sensitive detector. This method is also used in a variety of situations for the detection of small differences in contrast, brightness and other variables.

Given the concept of memory discussed in BCP, an equivalent is readily conceived as operating within the brain. Indeed, alternative memories can also be compared in the same manner. However, for this to work, there has to be some Entity to perceive, and report, the presence, or absence, of any differences.

From time to time I am surprised by the omission of memory from the on-going discussions on the Net. Surely this is a very important aspect of the Hierarchy?

Regards, Bob Clark

Date: Mon Jul 05, 1993 1:20 pm PST
Subject: GENERALIZING - RKC

[From Bob Clark (930705. 5:05 PM EDT)] Bill Powers (930629.1845 MDT)

From time to time you have remarked on this, and the related topics, included in the cited post. Generally (if I may "generalize from a few samples") I have agreed with your position.

>...The essence of probability is to observe what has happened and to
>compute the chances of its happening again, with no basis for predictions
>except experience at the level of phenomena. All kinds of generalizations
>can be made about probabilities, single, jointly, and in bunches.
>Mathematical theorems can be developed to bring out interesting properties
>of probabilistic calculations. But behind all these calculations, there is
>the simple fact that we must use probabilities mainly when we have no model
>of the underlying order behind phenomena.

>To characterize anything in nature as being intrinsically probabilistic is
>simply to give up the search for a generative model.

How does one "search for a generative model?" And exactly what is a "generative model?"

By this term, you seem to mean some representation of the data that can be used to derive related, but different observations. If the only data available consist of statistical samples, these data can be represented in various ways. But these data, alone, are not sufficient for logical/mathematical derivation of other relationships.

Thus, as has been pointed out by others, Galileo's stone-dropping from the Pisa Tower provides no method for deriving additional relationships. Indeed, without additional data about air resistance, dropping feathers would really louse up his data!

Without stopping to review Copernicus, Ptolemy, etc, consider Newton. Too often it is suggested that Newton "derived" his "laws" by "generalization from observation." As far as I'm concerned, he arbitrarily proposed his "laws," resulting in the _definitions_ of force (that which changes the state of motion of an object) and mass (the ratio of force to acceleration). Acceleration needed no further definition. These definitions were of such a nature that pertinent measurements could be made and other relationships (such as kinetic energy) could be calculated. Among the most important aspects of these definitions is that time was included, at least implicitly, through the definition of acceleration.

This combination of definitions with time included as an intrinsic independent variable, created the type of "generative theory" you seem to be seeking.

That theory, as is well known, works very well when applied to the physical world as perceived by most of those who know about those laws. If it had not, it would have been replaced long ago.

It has even been extended, with the addition of only rather few additional observations, through relativistic quantum electrodynamics into current high energy physics. Likewise, it has been applied with great success in such areas as solid state, chemical physics, statistical mechanics and others.

In my view, HPCT, when time is included as an independent variable, becomes such a "generative theory." The elements of the negative feedback loop are the same type as used in Newton's Laws. That is, each is defined in terms of the relation between its input(s) and its output(s) with time as an independent variable.

What is needed, I think, is some revisions and extensions to areas where it can clarify and contribute to many existing "people problems."

Enough for now. Bob Clark
Date: Mon Jul 05, 1993 1:25 pm PST
Subject: UNDERSTANDING - RKC

[From Bob Clark (930705. 5:15 PM EDT)]

Bill Powers (930623.0700 MDT) Ed Ford (930623:1100)

I've tried to write a general discussion of "Understanding" several times. It is harder than it looks. Let me try this:

When I say I "understand" something, I mean that I am able to fit it into my general working framework, and that I find no internal inconsistencies and/or contradictions neither within the "something" nor in regard to my own framework. Commonly, I am concerned primarily with my verbal structures, theories, etc, but I can also "understand" something that I do not accept within my own structures when I perceive it as something expressed by another person in his own terms.

When I ask whether you "understand" something, I expect that you will check the something against your corresponding framework and verbal structures and will report the existence (seeming or actual) of inconsistencies or contradictions.

In addition to receiving your "report of understanding," it is often a good idea to ask some key questions to verify that your listener has not over-looked something, or, in fact, inadequately considered the subject versus his previous concepts. Despite a favorable report, "misunderstanding" is quite possible.

This view points out the importance of using the listener's vocabulary (including body language, etc as well as verbal expressions). If there is too large a fraction of new words/concepts, he may easily think he "understands" when his "understanding" differs significantly from that of the speaker/writer.

In presenting HPCT to those lacking elements of understanding the physical and mathematical worlds, it must be presented in more ordinary, every-day language. This is possible, I've been experimenting a bit along such lines.

And so have you, Ed Ford, from your occasional posts. I'd like to know more of your materials.

Regards, Bob Clark

Date: Mon Jul 05, 1993 5:18 pm PST
Subject: ALGNMNT/EQUVLNT - RKC

[From Bob Clark (930705. 9:00 PM EDT)]

10 POSTS RELATING TO ALIGNMENT AND EQUIVALENCE

These posts demonstrate how easy it is to misunderstand even the most simple and direct objectives when trying to establish common goals.

Four people at the corners of a table seems a bit unusual -- how about a couple deciding which movie to watch? Or which restaurant for dinner? Or where to go for vacation?

Decision-making all over the place! And they can't all be automatic by previous agreement -- nor require the Reorganizing Function!

The DME can help here.

I've previously reported my experience with the City Government developing a major conflict with an unusually large number of citizens over selecting one garbage hauler for a franchise. Several meetings, tempers beginning to rise. The Council had become committed to the franchise concept, and the citizens felt they "were being pushed around." Like all good HPCT systems, they resisted strongly to "the push." A good place to demonstrate conflict resolution and decision making. I found I was in a place where I could play both parts: in a ten minute talk I summarized the situation and pointed out an alternative -- "Sponsorship with Permits" for any qualified hauler -- noted the disadvantages of the franchise and the advantages of the "Sponsorship-Permit" system. The vote went from 4 to 3 for franchising to 6 to 1 for Permits.

This is treating the government as a HPCT system, with the addition of a Decision Making Entity -- myself -- resolving the conflict when the Council was locked into its pre-established position.

ALIGNMENT VS EQUIVALENT
It's interesting to find that Sociologists have already been considering such interactions -- Interpersonal Relationships. "Harmonizing" may also be a useful term.

I, personally, prefer Equivalent as simpler, more direct and more readily understood by other people.

It looks like we're getting into the analysis of the behavior of Social Systems, whether or not they are -- or become -- Social Control Systems.

Regards, Bob Clark

Date: Mon Jul 05, 1993 5:54 pm PST
Subject: Direct mail for Bob Clark

For Bob Clark. Sorry to clutter net with direct mail.

Bob, please send me your phone number, direct, and any good files on your DME construct. I was logged off when you posted. Something has come up that may be of interest to you. It takes several days to get the archaic table server updated here, so I will be unable to send direct to you until midweek.

regards, Bill C.

Date: Tue Jul 06, 1993 7:36 am PST
Subject: Practical PCT; Shared Goals

[From Hank Folson (930706)] Rick Marken (930703.1100):

>Ed Ford (930702: 1350) replies:

>>You might ask those to whom I have taught PCT and have shown its

>>practical applications in their areas of interest and who pay me large
>>sums of money

>I bet my old neighbor, Art Janov, is hauling down a bundle, too. Does
>that put "Primal Scream Therapy" (PST) on a par with PCT? And how about
>Dyer, Bradshaw, et al? Glasser's not doing too bad, either. I submit that
>money paid for services is not a good measure of the success of PCT or
PCT loses hands down.

and also:
>It's nice if you can help yourself and help others
>with PCT ideas. I think it's wonderful that you are using PCT in this way
>What I disagree with is the idea that "practical success" is the real
>measure of the worth of PCT. If you measure success in terms of money,
>testimonials, number of adherents, etc. then PCT, even if it "catches on
>big" will eventually fade away like every other fad therapy -- there will
>always be new therapy in a better wrapper that will work just as well (or
>better) in terms of the measures you propose.

As to the money, you're looking at who gets it. Instead, think about who's spending it, for this is where the controlling is really happening. I don't know about others, but I have a very high reference level for keeping what money I have. It takes a very big error signal in some other area of my life to override it!

Why do people pick up on these fads? Because they are not satisfied with the success of their controlling efforts. Why do fads die? Because people soon find that they still have big error signals, and they move on. Would these same people pick up on PCT if it were generally available in a useful form? Yes, but pretty much at random. Would PCT be a short lived fad, too? If it worked for them, no. They would stick with it, because their error signals would be decreasing. People would then be less likely to pick up on these fads that you and I do not like. They will also be less likely to pick up on a new and better truth than PCT, were one to come along.

>And satisfaction with services is not a great measure either.
>David Koresh's followers (who survived) say they are still quite
>satisfied with the services he provided. Same is true for followers of
>Janov, the Maharishi, etc. People who pay a high price (in whatever form)
>for services very often are (or say they are) quite happy with those services.

If what a Koresh or a Janov is selling is what they are controlling for, they will be happy. That is PCT 101. It doesn't matter to them if you or I can scientifically prove that they are wrong. Their reality is what they perceive; this is a limitation of living control systems. How close people can get to Boss Reality (absolute reality) is determined by the quality of our physical sensors and, in my book, by the quality of our mental tools. PCT is a tool that can help us get closer to Boss Reality.

>I said:

>>I expect these questions to be a disturbance to Ed Ford and Dag Forssell.

>Ed (and Dag) said:

>>Not at all.

>Then why was there an action (you both replied to the post)? If there wer
>no disturbance, there would have been no change in output (your output
>changed from not posting replies to my posts to posting replies)? A
>disturbance is not necessarily what an observer might judge to be >something th
at has a negative effect on a controlled variable.

Dag and Ed saved me a lot of thought & time by answering your post [Thanks Dag & Ed], so I will return the favor: They probably perceived your "I expect.." statement as though you intended to say, "These questions will be seen by Dag and Ed as a personal attack on their views." And they replied, "Not at all.", because they did not perceive it as such, and then they rambled on about what they are happily controlling for. My post (Hank Folson (930701)) about the difficulties of controlling via communication was about this sort of thing.

>>The modeling and theorizing of PCT are
>>extremely valuable to me and others trying to apply these ideas, but
>>everything we say about PCT is really worthless unless these ideas can
>>guide us toward a better way of helping others live their lives more
>>effectively, in a more satisfying way.

>Well, here's where we completely disagree. I think the only thing of
>value is the modeling and testing -- the science is the only appropriate
>measure of the worth of PCT.

Well, here's where we completely disagree. I like modeling, so this is a question, not an attack: Does modeling absolutely prevent us from perceiving truth that isn't there? Is the modeling process so rigorous that in his enthusiasm a researcher can not choose to perceive what isn't really there, particularly in the testing phase? Modeling is based on our perceptions of reality which affect how we set up the model. Testing in laboratory conditions is a good start, but it is isolated from Boss Reality. We can never directly perceive Boss Reality, so our best test of our ideas is to expose them to Boss Reality, and see whether our ideas can withstand the unpredictable disturbances Boss Reality throws at them. "the science" isn't Boss Reality, it's only what we, with our limited sensors, perceive as Boss Reality. Haven't we seen people who are very smart believe in, and scientifically prove (by their definition of scientifically), things that turn out to be dead wrong? The history of living organisms is littered with examples.

[From Bill Powers (930703.0830 MDT)]

Yours is a very important post, Bill. My comments here are only nitpicks about matters peripheral to your points about shared goals:

>Consider the "shared" goal of furthering the development of PCT.
>

>Others would see furtherance of PCT if it became commercially successful,
>showing that people would pay money to use it.

My paragraphs above cover my thoughts about what money really means. Controlling for money does not mean you will get it from others. Those with the money must want something bad enough to spend their hard earned dollars. The meaningful and most important control lies with the spender. If people want to spend to reduce their error signals, the recipient does not need to offer much more than hope.

>Still others would see furtherance in learning how to use PCT to get
>others to do what they want:

I think this just needs to be clarified. To "get others to do what they want" is difficult to impossible. What I think PCT practitioners who really understand your theory are trying to do is target groups that are not controlling efficiently because of ignorance of how things really work i.e. through perceptual control. The targeting can be thought of as selfish, and meeting the practitioners wants. Seeing the target group mess up (in the practitioner's eyes) is perceived as creating real or imagined problems (error signals) for the practitioner. The controlling action being tried is to educate these people in the hope that they will begin to control in a way that satisfies the practitioner's goals, while also satisfying the target group's goals (otherwise, they won't do it). Isn't what is happening here a subtle variation of the shared goals scenario?

>(make them) respond more favorably to advertisements, and so forth.

This is what advertisers are traditionally taught to do. What I try to do as a devoted PCT practitioner fits this "shared goals" topic exactly: I try to find ways that what I want to make/sell matches what my selfishly targeted group, serious cyclists, want to buy because of where their reference levels are set and what error signals they have (or that I can create). Their reducing their error signals is to my benefit, even though they are not controlling to reduce Hank Folson's error signals.

>On an historical time-scale, human social institutions have the
>life of a May-fly. A few hundred years, and poof! Everyone
>realizes, or learns, that cooperation and alignment are better
>than competition and conflict. People keep trying to organize
>themselves into effective groups. But because people assume that
>everyone understands words in the same way, perceives the same
>world, and imagines the same unspoken assumptions, all groups
>eventually dissolve into rivalries and conflicts, and break up or
>change beyond recognition. It will be interesting to see whether>PCT has
>in it any ideas that will keep its proponents together
>any longer than any other group has survived. Maybe all groups
>are temporary expedients.

The May-flies are the other theories about living organisms, Bill!

Bill Powers (930703.2100 mdt): on Equivalence instead of alignment

>I think I'd rather do without this concept of "alignment"
>altogether. Instead, I propose that we talk about _equivalent_
>perceptions and _equivalent_ reference levels.

Might "compatible" be a good word to use here?

Keep disturbing, Hank Folson

Date: Tue Jul 06, 1993 8:03 am PST
Subject: Re: PC-Eudora and NUPop

From Tom Bourbon [930706.1037] Gary Cziko 930703.0600 UTC
>
>The following is an article from a local campus newsletter which describes
>two e-mail programs for IBM PCs and compatibles which CSGnetters may find
>of interest. You will see below that the price is certainly reasonable
>(i.e., free).
>
>I have been using Eudora for the Mac since 1990 and can't imagine managing
>CSGnet without it. I can provide info on this Mac program to anyone who
>asks. But remember, your mainframe link to the Internet has to be running
>POP to use these programs.
>
>Also note that (a) PC Eudora currently requires a direct network
>connection, but NUPop can be used over a modem, as can Eudora for Mac (as
>I'm doing now), and (b) both programs can send and receive binary files as
>attachments to e-mail messages (Bill P. and Greg W., are you
>listening?).--Gary

Gary,

In some of my posts during the past couple of months, I mentioned the fact that I finally have convenient and reliable access to the net. One reason for the change is NUPop. It was installed for me by a resident guru who does not use it and who tells me I am the local expert. So much for expertise.

I have found NUPop to do all of the things described in your post -- the good and the bad -- but life was not the same without it. My experience with NUPop is apparently like yours with Eudora: I *do* know I did not get along without it. My endorsement would be complete, if I could only get it to run by modem from home. Maybe I should buy the manual.

Until later, Tom Bourbon

Date: Tue Jul 06, 1993 9:53 am PST
Subject: Re: A modeling question

From Tom Bourbon [930706.1211] Kent McClelland (930630)

Some of you had access to the net during the holiday and the discussion about Kent's "modeling question" concerning alignment of reference signals has gone well beyond his initial remarks. I have just finished reading the accumulated posts will not repeat all of them. Instead, I will return to a few of Kent's comments that were directed to an earlier reply from me.

>Reply to Tom Bourbon (930623.1306)
>
>Thanks, Tom, for your thought-provoking answer to my earlier post (Kent
>McClelland [930622]) on the possibility of "Collective Controlled Variables."

(I include this line only to maintain continuity in the thread of citations.)

>In response to my argument that some "alignment" of reference levels between
>independent control systems is necessary for cooperative actions to occur,
>you first point out that cooperation often requires individuals to adopt
>different reference levels, not the same ones:
>
>>What matters is that the participants adopt reference
>>signals (that they adopt goals which we model as reference signals -- Hans?)
>>that result in each acting in a way that produces a match between personal
>>goals and present perceptions. The same constraints apply to my two
>>hemispheres when "they cooperate" to produce one-person performance of what
>>can also be a two-person cooperative task: they need not adopt similar
>>reference signals; all that is necessary is that each adopt reference
>>perceptions that result in the intended perceptions. . . .
>>Cooperation can, probably always
>>does, entail some necessarily different reference perceptions in the
>>participants.
>
>In sociological terms, we're talking about "division of labor," and your
>point is well taken. I'm a little unclear, though, about what you mean by
>"adopting goals" and how that's different from my "practical alignment" of
>reference levels, if the goals are necessarily shared in order for
>cooperation to take place. You go on to suggest that even when people
>perceive themselves to be sharing the same goals, they probably aren't, at
>least not in any detail:
>
>>. . . that kind of agreement says nothing about the
>>specific contents of the participants' heads. The socially affected and
>>approved options all reside in individual heads, as products of individual
>>interpretations or understandings. Necessary and important, to be sure; the
>>same in each person, probably not.
>
Here, you presented some observations about the similarity of my ideas and those of the ethnomethodologists, a group you also mentioned in a reply to Bill Powers during the holiday. It seems that, all along, PCTers were like ethnomethodologists and symbolic interactionists and never even knew it!

>Anyhow, my sociological conclusion about what you're saying is that you are
>undoubtedly right in emphasizing that each individual's perceptions are
>different, but I still think that at some abstract level people must have
>their control systems at least crudely aligned in order for cooperation to
>take place and that any such alignment has got to be an important
>sociological issue.

What went without saying, but probably should not have, in my earlier post was the necessity of each person having a reference perception for cooperating. Incidental interference and disturbance can occur even when people are ignorant of the existence of those with whom they interfere, and one person can contrive to control the actions of another who is ignorant of the presence and the intentions of the would-be controller, nonetheless the actions of the various individuals can become intricately coordinated. Deliberate cooperation is a different matter and it certainly should be an important issue in sociology and social psychology.

Cooperation requires that each participant have a reference perception to cooperate; beyond that, I stand by my earlier post and add the following: even the reference perceptions to cooperate are individual matters. But that topic was discussed at length during the three-day holiday. I think Bill Powers' comments on the topic were closest to my own thoughts. And I like Gary Cziko's suggestion that, instead of saying the reference signals in cooperating individuals are "aligned," we might say they are "functionally equivalent." That leaves a lot of room for the necessarily independent things people often do while cooperating.

>In the last part of your post, you dismissed my suggestion that independent
>control systems when aligned can become "super powerful", if there are enough
>of them.

I didn't mean to dismiss the idea, merely to raise some questions. I will return to that topic in a later post.

Until later, Tom Bourbon

Date: Tue Jul 06, 1993 2:17 pm PST
Subject: Functional Equivalence, Applied PCT

[From Rick Marken (930706.1330)] Gary Cziko (930705.1430 UTC) --

>This analysis appeals to me, but wouldn't it be better to use the term
>"functional equivalence" or use some other qualifier?

>"Functional equivalence" implies
>instead that the two reference levels may not be the same but that in a
>given environment they function equivalently.

Good point.

Hank Folson (930706) --

>Does modeling absolutely prevent us from perceiving truth that isn't there?

I don't understand the question. But my inclination is to say "no" just based on the "absolutely" part.

>We can never directly perceive Boss Reality, so our best test of

>our ideas is to expose them to Boss Reality, and see whether our ideas can
>withstand the unpredictable disturbances Boss Reality throws at them.

Are you saying that "real world" applications -- clinical practice, business, management, etc -- somehow expose PCT more directly to Boss Reality than do laboratory tests? I would hope that the exposure to Boss Reality is just as direct and challenging in the "real world" as it is in the lab (another part of the real world). Is the study of the control of a line on a computer screen somehow less of a test of Boss Reality then the study of the control of the location of a drug balloon that is being smuggled into a prison?

I think that the lab provides the opportunity to do more structured and poised tests of the nature of Boss Reality -- the same Boss Reality that exists outside the lab. If this were not true, then all science would be truly worthless. I agree that it is not always easy to go from the structured world of the lab to the more chaotic (appearing) world of application -- that is why the application of PCT is so important; the applied people can, hopefully, show how the models tested in the lab relate to the messier world of application where the variables will not stand still while you try to measure and manipulate them -- and where there are no instant replays, as there can be in the lab.

The real world of application is the ultimate test of PCT -- I agree. But this kind of testing must be based on a VERY CLEAR understanding of the models developed in the lab -- and how these models relate to measurable variables. The applied PCTer must know what the scientific PCTer knows -- and then some. The applied PCTer must know the PCT model in detail, its relationship to observable variables AND how to identify these variables ON THE FLY: variables that are possibly controlled, others that are disturbances and outputs. The applied PCTer must be able to identify charateristics of the environment that might determine the nature of the relationship between output variables and input variables, and he or she must be able to notice behavior that does NOT seem to be consistent with the PCT model. Then he or she should be able to describe what happened so that the scientist can invent tests to see what might be going on.

I see the role of the applied PCTer as being something like that of the field geologist. In order to do his or her job, the field geologist must know the basic models (physics, chemistry, plate tectonics) and he or she must know how to identify the relevent variables, often without the benefit of fancy instruments (checking out "volcanic" variables -- where the lava meets the vent -- is probably a lot like doing PCT in prisons --where the rubber meets the road). I agree that PCT field work (applied PCT) is extremely important. Maybe it is even more important than the science. But ultimately, applied PCT, in order to succeed as a "real test" of PCT, must be based on a deep understanding of the model. Otherwise, applied PCT contributes no more to our understanding of human nature than my amateur attempts at field geology contribute to our understanding of "earth nature".

Best Rick

Date: Tue Jul 06, 1993 2:26 pm PST

Subject: Re: Generalizing; statistics vs generative models

[Martin Taylor 930706 17:50]
(Bill Powers 930629.1845) (I've been away for a few days)

After talking about generative models versus generalization...

>This is what is behind my stubborn rejection of "probability" as
>an operative factor in nature. The essence of probability is to
>observe what has happened and to compute the chances of its
>happening again, with no basis for predictions except experience
>at the level of phenomena. All kinds of generalizations can be
>made about probabilities, single, jointly, and in bunches.
>Mathematical theorems can be developed to bring out interesting
>properties of probabilistic calculations. But behind all these
>calculations, there is the simple fact that we must use
>probabilities mainly when we have no model of the underlying
>order behind phenomena.

I would like anyone who agrees with this statement to read (or, I hope, re-read) my posting "Subjective probability intro." (Martin Taylor 930326 16:20). In that posting, I tried to show from first principles why is is ABSOLUTELY NECESSARY to have a model of the underlying order behind the phenomena before you can talk sensibly about probability. To say that "The essence of probability is to observe what has happened and to compute the chances of its happening again," is to mistake one way of estimating a probability value for probability itself. It is like saying "the essence of length is to take a tape measure and put it against an object," or "the essence of intelligence is to present an IQ test."

Measuring how often something happens when the "same" conditions are repeated is itself subjective. You must have a model, in order to assert that the conditions are "the same." It may not be a generative model, but it is a model of what is relevant to your measurement. Being wrong there is to be wrong in how the thing you are measuring works.

To take "probability" as the relative frequency of some event occuring, over an infinite number of opportunities is common. It is both philosophically wrong and practically misleading. It leads to concepts like "significance tests," which are the absolute bane of science. It is rather like taking the essence of information theory as being a measuring tool for transmission channels.

Martin

Date: Tue Jul 06, 1993 2:27 pm PST
Subject: Re: speech: /di/ /du/

From Tom Bourbon [930706.1624] and Andy Papanicolaou

>Rick Marken (930703.1100)] Tom Bourbon (930702.1655)--
>

>Very interesting post on /di/ and /du/. I spent most of yesterday
>afternoon saying these two syllables with all kinds of different
>mouth configurations -- the Gary Cziko test. I was able to recognize
>the /d/ quite well, even when it was articulated very differently in
>both cases.

We were pleased to learn that someone else spent some time playing /di/s and /du/s along with us. There can be some pretty interesting variations on the theme, can't there?

And we were pleased to see that you had given some serious thought to the questions we posed.

>... I think the Cziko test rules out analysis by synthesis (A-S)
>as even a plausible consideration. But I think A-S is also ruled out
>in terms of modelling too; how would you build an A-S model that
>recognizes /d/ when it is said by someone else, for example?. The input
>to the model is x. This x would presumably set off the "synthesis"
>routines, only one of which produces x as the result. But this is
>just the same as analysis by transduction. Look:

```
>
>Synthesis model of recognition      Analysis model of recognition
>
>        S1---> x                          A1 --> 20
>x --->  S2 --> y                    x --->A2 --> 1
>        Sn---> z                          An --> 2
>
```

>In the synthesis model, the input sets off all the pheneme articulation
>synthesizers. The S1 synthesizer is the articulation pattern for /d/,
>S2 and S3 and the articulation patterns for other phomemes. Now, how
>does the system know which articulation pattern is the one what would
>have produced x? Presumably because it produces (in imagination) a
>perception that corresponds to the actual input. So the output of S1
>is "imagined" x and it matches the input so the input is recognized.
>But suppose the input is y and the phoneme is still /d/ (as in di, du).
>Then S2 must also be a /d/ recognizer -- it is the same articulation pattern
>as S1 but it must result in y so y will be recognized as /d/ also. So analysis
>by synthesis is really just proposing a separate articulatory recognizer
>for each acoustic input that represents the same phoneme. In fact, S1 and S2
>don't need to be there; all we need are the results -- x and y -- which are
>compared to the input. So A-S boils down to a template matching system;
>the synthesis part, when you actually start trying to build the model,
>is irrelevant.

We agree, wholeheartedly, that the A-S model has problems, some of them probably fatal. If it didn't have those problems we probably wouldn't be having this discussion about building a model that recognizes and produces speech; we could all run down to Radio Shack and buy one.

>The analysis model is a hell of a lot more elegant; the problem is
>designing the A1, A2...An analysis functions. I believe that these
>funtions are at the event level -- we don't necessarily have

>/d/ detectors. Anyway, the A1 system would put out a bigger
>signal than the others when an input came in that was /d/ like --
>either x or y. A1 could just have S1 and S2 templates in it
>and put out a signal that corresponds to the degree of match to
>either template.

We also agree that the straight forward Analysis model would be more elegant, but the emphasis is on *would be*. And we agree that there may be no need for /d/ detectors as such and that the Analysis function(s) might have to be placed at the event level. (See how agreeable we can be?)

Our problem arises when we concentrate on building an input-output analysis box, or filter, or Pandemonium demon, or feature analyzer, or whatever, that gives us as its output a perceptual signal, p, that under the right input circumstances matches the current reference signal, p*. To build such an analysis device, we need to know the p*. If we have no idea what p* to put in the model, we can't build an analysis box such that its output signal can match the values specified in p*.

It seems to us that the problem, at this stage, is not what kind of Analysis function to adopt -- this filter, that filter, or no filter or template at all. (Some plausible sounding options have been laid out over the past thirty years or so.) Rather, the problem at this stage seems to be the kind of p* with which we should endow the model. This where we get stuck, every time. As we said in our earlier post (Tom Bourbon [930706.1624]), as perceiving systems, when we listen to a radio program, we can establish listening for the experience /d/ as a reference signal and once we have done that, we can very effectively detect /d/s embeded in the program. How then should we represent that experience as p* in the model? What kind of p, corresponding to what kind of thing in the domain of speech sounds, should the model be made to control for?

>I think this is worth a little discussion at the meeting; I bet we
>get a lot of people walking around saying /di/ and /du/ with tongue in
>cheek. Hmmm. "Tongue in cheek" -- a good name for the "Cziko test".

It could turn out to be a pretty cheeky meeting.

Until later, Andy Papanicolaou and Tom Bourbon

Date: Tue Jul 06, 1993 2:39 pm PST
Subject: Re: Science & PCT

[Martin Taylor 930706 18:00] (Rick Marken 930630.1400)

>It seems to me that many people WANT to believe that human behavior
>is complex. This is another reason PCT has problems; PCT says that,
>underneath it all, human behavior is NOT complex; it is simply the process
>of controlling perceptual input. The basic law of behavior (control of
input)
>is simple -- I think this offends many people. For some reason, people can
>accept the fact that a simple law (inverse square) underlies all the physical

>complexity they see; but the idea that a simple law underlies all the
>behavioral complexity just seems like hubris or something.

Newton proposed more than an inverse square law. Even with the simplicity of his laws, it is impossible to predict the motions of even a set of three bodies into the indefinite future, and when it comes to the motions of the gas molecules in a room, you don't even try.

Are you sure that there is onle ONE simple law underlying the complexity of human behaviour? Do you think that you can use it to predict in detail, when you have to model the interactions of millions of elementary control systems? That's an awfully big projection from using one, or two ECSs to predict tracking data with 99.9% accuracy. It's like saying that we can predict eclipses back a few thousand years, and therefore we can forecast the weather using the same methods.

Martin

Date: Tue Jul 06, 1993 3:20 pm PST
Subject: /di//du/, Subjective Probability

[From Rick Marken (930706.1600)]

Tom Bourbon [930706.1624] and Andy Papanicolaou

>It seems to us that the problem, at this stage, is not what kind of Analysis
>function to adopt -- this filter, that filter, or no filter or template at
>all. (Some plausible sounding options have been laid out over the past
>thirty years or so.) Rather, the problem at this stage seems to be the kind
>of p* with which we should endow the model. This where we get stuck, every
>time.

I need some help on this. I think of p* as just a unidimensional variable. What do you (I presume I'm talking to Andy here) mean by "the problem at this stage seems to be the kind of p* with which we should endow the model"? Do you mean that you have problems determining the type (level) of perceptual variable whose value is specified by p*? Do you mean you have problems determining the higher level outputs that are combined to produce p*?

Also, what is the model going to do? If it's going to say /di/ and /du/ then you do need to specify some p* values since you will probabaly have a hierarchical model. But if your model is just a recognition model -- where the outputs of /di/ and /du/ detectors vary appropriately depending on whether /di/ or /du/ was spoken -- then you do not NECESSARILY need to specify p* values-- unless you feel that imagined lower level perceptions (the result of replaying p* values back into the perceptual channels) can improve recognition performance -- which they might.

Martin Taylor (930706 17:50) --

>I would like anyone who agrees with this statement [of Bill Powers (930629.1845)

I do!

>to read (or, I hope, re-read) my posting "Subjective probability intro." (Martin >Taylor 930326 16:20).

Glad to. Could you re-post it. Thanks.

Best Rick

Date: Tue Jul 06, 1993 7:09 pm PST
Subject: Re: perception,nupop,equivalence, gen models

[From Bill Powers (930607.1800 MDT)] Bob Clark (930705.4:55 PM EDT) --

>Your discussions of "contrast" (per Bruce Nevin (Mon 930621
>09:44:18 EDT)), "phonemes" and "differences" call for
>corresponding kinds of "detectors."

>I agree, but the question remains, how does such a detector
>work? As I see it, there must be some way for the incoming set
>of perceptions to be compared with some kind of "reference set"
>of perceptions. That is, the incoming set is compared with
>remembered perceptions.

The basic PCT model does not use the concept of "templates" (one name for what you are talking about) to explain perception. There are all sorts of practical difficulties with that concept, once you stop thinking of a narrow range of positive examples like recognizing changes in stars fields using a blink comparator.

Consider the problem of perceiving a cube. To use the template idea, you would have to suppose that there are templates for cubes of all sizes, orientations, colors, and lighting conditions, simultaneously applied. Perceiving a spinning cube would be even more daunting by this method, and it's still harder to imagine how the template method could yield a perception of a relationship like "inside."

The model I settled on some time ago, while it has some drawbacks, treats perceptions simply as scalar values of functions of lower-level signals, each of which is also a scalar representation of still lower-level perceptions. By scalar I mean that any perception is indicated strictly in terms of amount or magnitude, by a single one-dimensional neural signal.

Of course what we normally think of as perception is composed of hundreds or thousands of perceptual signals. In the PCT model, each signal represents just one variable attribute of experience which can only be present to a greater or lesser degree. When we abstract patterns from such collections of signals, we do it with a higher-level perceptual function that receives sets of the lower signals, applies a transformation typical of that function, and reports the result, once again, as a single one-dimensional signal. As I wrote last week (or so), any function of multiple input signals creates surfaces of indifference, defining ways in which the set of input signals can vary without altering the higher-level perceptual signal. Those types of variations are the transformations under which the higher signal is invariant. Lines normal to the surfaces of indifference define other ways in which the input set can vary that cause the resulting perceptul signal to become larger or smaller. They define one dimension along which the set can be controlled.

Once this model is adopted, the reference signal becomes trivial: it is just a signal with a certain magnitude. All it has to specify is the magnitude of the perceptual signal that is to be brought about by action through lower-level systems. Memory also becomes trivial, more or less: all that can be recorded are magnitudes that have occurred in the past."A memory," of course, consists of hundreds or thousands of such recordings of magnitudes, perhaps indexed by time or other patterns. To recognize a memory, one must use the same perceptual functions that are used with externally-originated signals.

Modern perceptron research seems to be going down this same path. Of course a single perceptron has multiple outputs, but each output can be seen as the scalar value of a different function of the common set of input signals. It is the fact that one output rather than another is activated that indicates perception of some one input pattern. Any one output can represent only the degree to which one particular pattern is present.

>There are optical, and magnetic, character readers that work
>very well. These devices are passive and do not use feedback
>systems. They also do not need any form of spectroscopic
>analysis or other analytic equipment.

On the contrary, modern optical character readers do a great deal of processing very much like the kind I described. They do not use templates at all. That approach was tried long ago and found to demand extremely rigid and artificial input conditions, which are evidenced in that peculiar set of numbers we used to associated with checkbooks (I guess they're still used -- there was quite an investment in the reading machines). Once that approach was abandoned, the capabilities of character recognizers grew a lot; multiple fonts could be recognized in various sizes with a kind of accuracy that would have been impossible for the old-style recognizers. And new fonts could be learned, if the programs were told the correct identifications of various letters. This method depends mostly on extracting features of the letters, roughly what I mean by attributes, which are simply represented as magnitudes. In multilayer perceptrons, these features are discovered by a random search, and are combined in the top layer to generate a signal at just one output when a particular set of attributes is detected.

>From time to time I am surprised by the omission of memory from
>the on-going discussions on the Net. Surely this is a very
>important aspect of the Hierarchy?

Memory is certainly a phenomenon, but its mechanism may not be at all like what we experience. It doesn't seem a necessary part of the simple working models we've been using. Some day, no doubt, we will find an experimental phenomenon that can't be explained without invoking the concepts of recording and playback. But not yet. I don't want to add it just for the sake of saying it's in the model. I already did that in BCP. You'd really have to do some memory-specific experiments to get into that subject in any meaningful way.

Tom Bourbon, Gary Cziko (various) --

RE: Eudora-PC and NUPop

These programs have to operate with a corresponding program running on the mainframe. In my case, the mainframe doesn't have the corresponding program, so I can't use either one. There doesn't even seem to be enough dial-up activity to warrant getting a faster version of Kermit. I get about one busy signal every couple of months.

The mainframe with which I communicate is a VAX cluster, which automatically shovels mail into a mailbox for me. The mainframe allows for pseudonym files, so that I can store Gary's bitnet address in a file called cziko.dis, and then simply type @cziko for the address. This can also be used to reduce lists of communicants to a single mnemonic name.

My communication program, PC-Plus, can run scripts that send messages, wait for certain messages to be received, and other automatic interactive stuff. When I write something to send, I store a formatted copy of it in a file called "s", exit from the word processor, and type "sendcsg." This starts a DOS batch file which activates PC-Plus with a specific script file. PC-Plus dials up the mainframe and uses Kermit to upload the contents of "s" into a temporary file. It then (controlling the mainframe) exits Kermit, switches to MAIL, does the commands for sending the temporary file to "@csg", which of course contains the actual address, and waits for me to enter a subject. After I enter the subject, the script file sends the message off, deletes the mail, exits, deletes the temporary file, and logs off, returning control to DOS. Under Windows 3.1, a double click of an icon does all this. All I have to do is enter the subject when prompted.

Getting mail is even simpler: I type "mail" or in Windows double click on a mailbox icon. The whole cycle goes automatically, including concatenating the mail into one file and selecting a new DOS filename called "mailyymm.dd[a..z]", derived by looking for existing filenames and getting the system date. After a few minutes the program finishes, and I can call up my word processor and read the mail offline.

I also have setups for direct sends (manually enter @personname and subject), and for sending binary files and coming up logged on ready to send them. There's also a plain "logon" that logs me on.

This doesn't let me do CCs or other fancy stuff, but it's cut down the labor to the point where I don't mind it any more.

Greg Williams (930704) --

>>If the the people never even get started toward achieving the
>>goal, how will they ever discover whether they had different goals?

>They won't. In my example, I said that the folks "failed." They
>DID try, and they found that the rock was too heavy to lift in unison.

I seem to be having great difficulty in expressing myself. I was trying to say that because the people failed, they never found out what goals each person actually had (as opposed to their descriptions of their goals). Only if they get near the goal-state will differences in reference levels, or non-equivalences, show up.

>I didn't say: "Goals MUST be aligned..."; I said
>"Goals... can be aligned sometimes..." You sound like you're
>trying to pick a fight with YOUR misreading of my words.
>Nothing better to do on the 4th?

Well, no. Isn't the 4th of July a day for celebrating warfare and other such bracing activities? Anyway, I was trying to point out that in your example there's no way to tell if the goals are aligned. Failure to get the rock off the ground simply doesn't tell you anything about the intended end-result each person had in mind.

I'm trying to push for an operational definition of alignment, which is why I'm speaking of equivalence, or as Gary suggests, functional equivalence. That can be established only by seeing what happens when people start approaching their goals. If they come into conflict, then maybe failure to reach a goal does give us information about individual disparities (or equivalence) between goals.

>Ah, statistics. How about quantifying your "often" a little?
>Are your data on-disk? If so, I could run them through a neat
>program I have to determine their significance.

Unfortunately, it won't tell either of us their importance.

It would be nice to have some real data on this subject, I agree.

Gary Cziko (9307.1430 UTC) --

>This analysis appeals to me, but wouldn't it be better to use
>the term "functional equivalence" or use some other qualifier?
>"Equivalence" all by itself implies to mean that the reference
>levels really are in some "objective" sense the same ...

"Functionally equivalent" is OK. Also, for variety, how about "operationally equivalent?" Maybe this is just another of those cases where there isn't any one nice word we can assign to have the right meaning.

Any talk of equivalence or alignment seems to imply a third party judging "objectively" whether two other people have the same goals or perceptions. We need some way to say that for ANY observer, goal equivalence is judged, basically, by the Test.

RE: thermostats

Nice illustration. Here's another: one thermostat has its reference dial calibrated in centigrade, the other in Fahrenheit.

Bob Clark (930705.5:05 PM EDT) --

>How does one "search for a generative model?" And exactly what
>is a "generative model?"

A generative model explains observations by appealing to hypothetical processes that are more detailed than the observations. We explain certain observed relationships between environment events and behavioral actions not by classifying them and generalizing from specific instances, but by proposing a mechanism containing a perceptual function, a comparator, and an output function, and providing specific parameters for these components (where possible).

Generative models are inherently quantitative. In order to compare them against a real system, it's necessary to give them quantitative properties that completely determine how the model will operate (through calculation, simulation, or building physical examples). In simulation, we set the parameters, establish initial conditions, and let the model run while simulated environmental effects act on it. The result is a description of the behavior of the model through time. Then we let the real environmental effects act while a real person does the behavior in question. We then compare the model's behavior point-for-point against the real peson's behavior. The differences tell us how to alter the model's properties (and sometimes its organization) to achieve a closer match. The ultimate test of the model is to bring in new conditions, and without changing the model's parameters, show that the model matches the real behavior under the same new conditions.

The control system model is a generative model because out of its own properties it _generates_ behavior that can be compared with real behavior.

>By this term, you seem to mean some representation of the data
>that can be used to derive related, but different observations.
>If the only data available consist of statistical samples,
>these data can be represented in various ways. But these data,
>alone, are not sufficient for logical/mathematical derivation
>of other relationships.

Pretty close. We don't use data that are only statistical, except in the laboratory sense of errors of measurement. The aim of our modeling is to predict data with very high accuracy, so high that "correlation" tends to be a useless measure. Quite often we have been able to get to such high levels of predictivity. Unlike L. Ron Hubbard, we have actually done the experiments; you'll see them at the meeting.

I am convinced that if we simply stick to our standards, we will eventually be able to build up non-statistical understandings of all human behavior, even that which today is thought to be inherently uncertain. Of course in part this will involve changing what it is we think to be predictable. We will never be able to predict individual responses to weather, for example, because it is unlikely that we will ever be able to predict weather with much accuracy. But as we learn to recognize and measure controlled quantities, we should be able to make a great deal of sense of the OUTCOMES that people seek and maintain, no matter what disturbances come along. Even now we can make some good guesses: if it rains, most people will do something that keeps them from getting soaked (sorry, though -- no data on my disc).

>Too often it is suggested that Newton "derived" his "laws" by
>"generalization from observation." As far as I'm concerned, he
>arbitrarily proposed his "laws," resulting in the _definitions_
>of force (that which changes the state of motion of an object)
>and mass (the ratio of force to acceleration).

These definitions were, as you say, abitrary inventions. But Newton's law of gravitation, while also an invention, is a generative model. It proposes that the phenomena of gravity at the observational level are explained by the action of elements in a much more detailed world. The inverse-square law does not in fact hold true for arbitrary physical objects. The underlying model says that every point-mass in the universe attracts every other point-mass with a force calculable from the basic model equation. If these forces are integrated over all points in two masses, the law of attraction at the phenomenal level is predicted -- zonal harmonics for almost-spherical bodies, other laws for disks and rods and cubes, and so forth.

>This combination of definitions with time included as an
>intrinsic independent variable, created the type of "generative
>theory" you seem to be seeking.

That's not quite what I mean by a generative model. A generative model actually proposes things that are not directly observable. It explains observations as outcomes of underlying quantitative processes, which we know ONLY through their outcomes. As technology improves, we often, rather astonishingly, find that the guesses about underlying processes were right: think of the chain of progression that started with the hypothetical abstract gene and ended up with electron-microscope pictures of the DNA helix.

Time isn't what makes a generative model, although it's an essential ingredient. Time can also be an element of descriptive models or generalizations, as in saying that accuracy of recall falls off with time. Descriptive models and generalizations are different from generative models because they go in the opposite direction in searching for explanations. They treat the observations as particular examples of general relationships; they employ classifications and abstractions instead of underlying mechanisms.

Come to think of it, Bob, I believe that all of the modeling efforts we talk about nowadays were developed rather long after our association. You and I did quite a lot of work with control phenomena, but we never got very far with actual working models, even with our analogue computer. I didn't get started until 1973, when I was able to get a home computer, and even then it took several generations of computers for enough computing power to exist to do real-time experiments and create running models. When you come to Durango, I think a lot of things will become clearer to you. There's nothing like seeing an actual working model predicting random-looking behavior. That's what Rick and Tom get so het up about. Once you've seen that sort of model at work, it's pretty hard to take other theories very seriously.

PS: we now have 29 attendees for the meeting, plus half a dozen guests. There will be a MOB at our house Tuesday night.

Best to all, Bill P.

Date: Wed Jul 07, 1993 4:00 am PST
Subject: GA bibliography

Date: Wed Jul 07, 1993 7:22 am PST
Subject: Real life

From Tom Bourbon [930707.0900]

From time to time, discussions on this net turn to drawing distinctions between science, on the one hand, and the real world or real life, on the other. That was the case during the past few days. The next time someone is tempted to draw that distinction, I recommend a wonderful vaccine, prepared by Philip J. Runkel. (How I wish he were on the net!)

"Real Life

There is no such thing as real life. Or unreal life. That is, every condition or setting of human life is as real as every other.

In physical matters, the difference between the laboratory and 'real life' is simply that the workers in the laboratory are usually more skillful (in the necessary ways) than those elsewhere and the tools more precise. Gold behaves the same way outside the laboratory as inside. So do hydrogen, arsenic, and chlorotrifluoromethane. Things balance in the same way, accelerate in the same way, and heat up in the same way in and out of the laboratory.

Similarly, people act to control their perceptions whether they are in the laboratory or out. As a natural kind, the person is as natural in and out of the laboratory as gold or arsenic. ..."

(From page 150 in the section, "Real Life," in "Chapter 12, Social Psychology," in Philip J. Runkel, *Casting Nets and Testing Specimens: Two Grand Methods of Psychology*, New York: Praeger.)

The idea that the laboratory and real life are different and that the two have little or nothing in common sprang from the failures of experimental laboratory psychology: The results of statistical hypothesis testing, as a technique for identifying causal variables in laboratory experiments, have little to do with what happens outside the laboratory. The conclusion reached by psychologists? The laboratory is not the real world and we can't expect our laboratory results to "generalize" to that world. That conclusion rivals, and clashes with, another reached by experimental psychologists: Our experimental and statistical procedures guarantee noise and slop in the data, therefore, behavior is too variable to predict with any degree of precision -- a conclusion that reverses the prior one and asserts that results from the laboratory *do* apply to the real world.

Phil got it right.

On a related topic discussed in recent days, there is a difference between "good science" and "commercial success." This observation is not intended as a rebuke to anyone who applies PCT or markets it as part of a program. Rather, it is a simple observation, offered in a spirit of good will. I think Rick's comments were on the mark when he wrote about how John B. Watson's successes in advertising (and I would add, in publishing immensely popular books and delivering equally popular lectures on how radical behaviorism could improve personal and family life, child rearing, business practices and society at large) provided no evidence either way concerning his theory of behavior. Neither do *my* personal beliefs and experiences, when I think of myself as "a person who knows about PCT and thinks it makes an important contribution to my daily life," provide evidence either way concerning perceptual control theory as a scientific undertaking.

Meanwhile, "There is no such thing as real life. Or unreal life."

Until later, Tom Bourbon

Date: Wed Jul 07, 1993 8:47 am PST
Subject: One law, one special law

[From Rick Marken (930706.0800)]

Martin Taylor (930706 18:00) --

>Are you sure that there is onle ONE simple law underlying the complexity
>of human behaviour?

Well, two:

1) p = r

2) o = -1/g(d)

>Do you think that you can use it to predict in detail, when you have to
>model the interactions of millions of elementary control systems?

We have already used it to predict, in detail, tracking behavior that results from the the operation of thousands (possibly millions) of control systems, all of which probably interact -- at least via environmental effects on each other's perceptions.

>That's an awfully big projection from using one, or two ECSs
>to predict tracking data with 99.9% accuracy.

I think we're talking about different kinds of complexity.

All I meant was that there are people who seem to think that the apparent complexities of human behavior are mirrored by underlying complexities in the organization of brain/nervous system function. Freudian/Jungian type theories, which posit complex interactions -- repression, regression, compensation -- between psychical entities -- id, ego, superego, archetype -- come to mind. I would also put trait theories in this category -- where every noticable aspect of behavior is thought to reflects the existence of a distinct mental trait; so there are traits for being tidy, lazy, smart, glum. There could could even be a "post to the net" trait; I'd score high on that one.

PCT is not unique in assuming that beneath the apparent complexity of behavior is simplicity. I think this assumption is at the heart of the scientific (and even the religious) attitude toward experience. S-R theory says that a simple causal process underlies behavioral complexity (this guess is wrong but, instead of revising it, the S-R types add an extra hypothesis that makes it difficult or impossible to reject their model -- the hypothesis that behavior is, in part, the result of random processes; a neat trick, producing an unfalsifiable model that is still the basic model of the behavioral sciences). Reinforcement theory, a variant of S-R theory, says that a simple selection process underlies behavioral complexity.

Cognitive theory is actually a step backward on the road to simplicity. The basic idea of cognitive theory is simple (output generation based on mental events) but most cognitive theorists make up all kinds of complex mental "mechanisms" to account for every different kind of behavior they study. The "cognitive mind" seems to be like a collection of different Rube Goldberg machines, each designed to produce a different kind of cognitive behavior -- attention, iconic memory, backward masking, short term memory, decision making, problem solving, etc.

I have no doubt that the brain is a very complex device and that the means by which it computes and controls perceptions is very complex. But the basic organizing principle of the brain is quite simple -- and is summarized by laws 1) and 2) above. Like Newton's laws, the laws of PCT are just a start. There will be plenty to do in working out the details of how these laws are implemented in the nervous system and how they apply to all behavior. But the evidence so far is that these are the basic laws of behavior. Of course, precise, negative evidence would be much appreciated (Bourbon and Powers have provided the negative evidence for S-R and cognitive models so when someone does produce the evidence that rejects PCT I hope they have a nice, new alternative ready).

Best Rick

Date: Wed Jul 07, 1993 2:36 pm PST
Subject: Re: Speech

From Tom Bourbon [930707.1653] and Andy Papanicolaou

To follow up on our conversation with Rick Marken, let us try to clarify the problem we see in trying to construct a functioning PCT model for the production of speech. We will do so by directly comparing that kind of model with a common PCT model for tracking.

In a conventional pursuit tracking experiment, say that I intend to keep the cursor "just below the target," then I do that to my satisfaction.

Now I want to construct a PCT model that will duplicate my performance. To duplicate my intention, the model must have a specific value for the reference signal, p*. I decide to define my perceptual experience of "just below the target" in terms of pixels, which are units of vertical resolution on the computer screen. I follow the now-common procedure for estimating p*: I calculate the environmental distance between the cursor and target for every sampling interval during the task (often that is 1800 intervals) and I calculate the average distance. I discover that, during this particular tracking run, the average distance was 5 pixels. I use a scaled analogue of that external physical distance as the reference signal, the p*, in my model. For the sake of simplicity, I assume the physical distance at any moment is converted directly into a perceptual signal, properly scaled. The modeled perceptions, and the reference signal, all are described and measured as direct analogues of some measurable feature of the external world.

I run the model. It works very well; the correlations are greater than +.99, between its simulated positions of the handle and cursor and my positions of handle and cursor.

Now I intend to detect /d/s in a stream of speech and to say /di/ every time I hear a /d/. I do so to my satisfaction.

Now I want to construct a PCT model that will duplicate my performance. To duplicate my intention, the model must have a specific value for the reference signal, p*. I decide to define my perceptual experience of /d/ in terms of ... . To which aspect of the world must our p* conform? That is the question we encounter, and fail to answer, every time we try to reason our way through this problem. The question obviously applies to the perceptual signal, as well as to the reference signal: The perceptual signal for /d/ should be an analogue of ... of what?

Until later, Andy and Tom

Date: Wed Jul 07, 1993 6:26 pm PST
Subject: Speech Model

[From Rick Marken (930707.1900)]

Tom Bourbon [930707.1653] and Andy Papanicolaou --

>To duplicate my intention, the model must have a
>specific value for the reference signal, p*. I decide to define
>my perceptual experience of "just below the target" in terms of
>pixels, which are units of vertical resolution on the computer screen.

I think you are descibing the cart before the horse. The first thing you did was define the perceptual variable, p, by defining the function that computes it:

p = pixels(t-c)

The perceptual variable, p, is the number of pixels between target (t) and cursor (c) -- as vertical distance. p is just a number; your function (which I am calling "pixels"), because of the way it is computed, defines that number as the signed vertical pixel distance between cursor and target. This was the first BIG step in your modelling; you are guessing that the subject is controlling a particular perceptual variable defined by the "pixels" function of the objective situation rather than some other one (for example, you could have defined p as (t+c) or sqrt(t-c), etc.-- pretty silly choices, but possibilities nevertheless). The problem with tracking tasks is that the most profound part of the modelling is the simplest -- defining the perceptual function. If you build models to control other variables -- like size or shape, which are functions of two independent variables, x and y -- then the "definition of p" problem becomes more interesting -- ie. p can plausibly be x*y, y/y or x+y, etc.

Now, to make the model work you have to select a value for p* -- a free parameter. This is just another number -- but it must be selected so that, when p=p*, the corresponding objective distance between t and c is as close as possible to what was observed in the experiment.

>Now I intend to detect /d/s in a stream of speech and to say /di/
>every time I hear a /d/. I do so to my satisfaction.

And, again, the big problem (as I see it) will be to design the perceptual function:

p = /d/(acoustical waveform)

where I have named the perceptual function /d/. Where the "pixels" function transformed t-c into a time varying scaler representing vertical distance in pixels, /d/ transforms the acoustic waveforms into a time varying scaler representing degree of "/d/-ness" in the waveform in unknown units. To pick off the /d/s in continuous speech, I would create a new perceptual variable -- pd -- which is the difference between p and a perception that is a "model" of /d/ -- call it p'. p' is the "imagined" output of the /d/ function -- the result of creating the maximum output that that function can produce. p' is still just a number but it's a number that corresponds to the most /d/ like perceptual signal. I would then make pd the input to another control system -- the "decision" system. The output of this system is the word "/di/" whenever pd is greater than or equal to the reference input to this system. The reference for this system, pd*, is set so that the system responds "/di/" only to those parts of the waveform to which the subject responds "/di/". Of course, this decision making control system (that controls pd) is part of another control system that is trying to match the perceived output ("/di/") to pd; when pd is >= 0, then perceived output should be "/di/", other wise not.

> To which aspect of the world must our p* conform? That is the question
>we encounter, and fail to answer, every time we try to reason our way
>through this problem.

I think you are really asking "to which aspect of the world does p conform?"; p* implicitly corresponds to the aspect of the world represented by p. This may be a fairly academic distinction but it might help focus your efforts -- the problem of modelling speech (in this case) is really a problem of determining the perceptual function, /d/, which may be a function of the outputs of several lower level perceptual funcitons -- equally difficult to model.

> The perceptual signal for /d/ should be an analogue of ... of what?

YES. That's the question.

I have a suspicion that the /d/ function is based on more than just the frequency transition part of the syllabic waveform. For example, I think the /d/ in /di/ depends on something about the entire waveform. I believe that the /d/ in /du/ also depends (in the same way) on the entire waveform. So you don't hear /d/ until the entire waveform has occurred (because the output of the /d/ function doesn't go to max until the whole waveform has occurred). There is pretty good evidence that this is the case. If you just play what appears to be the /d/ segment of the waveform it just sounds like a click. The acoustic /d/ part of /di/ sounds like a different click than that in /du/ so there is no perceptual constancy of the "/d/" part of the acoutic waveform alone -- even though teh waveform is the result of the same articulation pattern (and even if the articulation is done by yourself; try it; just say the /d/ part of /di/ and /du/).

Speech recognition must be at the point where they have pretty successful phoneme detectors. A fairly speaker independent version of the hypothetical /d/ function must have been built already; one may be commercially available in voice recognition computer input systems. Maybe you could find a company that makes these and ask how they do it -- or go to the voice recognition literature. I betcha that a fairly good /d/ detector exists and can be used in your model of a person mimicing /d/ when they hear it.

Best   Rick

Date: Thu Jul 08, 1993 9:11 am PST
Subject: Re: Speech Model

[Martin Taylor 930708 11:45]
(Tom Bourbon various, 930703-07 and Rick Marken 930707.1900)

Taking the last first:

(Marken)

>Speech recognition must be at the point where they have pretty successful

>phoneme detectors.

I think it fairer to say that speech recognition is at the point where it is reasonably assumed that there never will be successful phoneme detectors that base their outputs on the current waveform alone. Most phoneme-based recognizers assert a set of possible phonemes that might be present during any particular frame. To make an assertion as the which phoneme is actually present, one typical method is to see which sequences among the possible phonemes correspond to reasonable strings; reasonable might mean "having a high sequential probability of occurrence" or "leading to a syntactically permissible set of words" or "something else the maker defines."

To check for yourself the likelihood of finding accurate ways of extracting phonemes from waveforms without recourse to higher-level constructs, listen carefully to almost any US broadcaster saying the word "President." I haven't actually done this recently myself, but I did when the President was Bush, and the only phonemes that seemed to be represented in the waveform were /p//r//e//s//n/. It was quite rare that there was anything corresponding to the stop that a /d/ should have. Nevertheless, if you didn't listen very closely, you heard the /i//d//e//n//t/. A "good" phoneme recognizer would have a hard time getting four of the last five from the speech waveform. They are apparently perceived from the expectation.

Which brings me to the main point, which is to describe what we are trying to get at with the "syntax recognizer" that I have mentioned from time to time. It is quite close to what Tom and Andy are talking about, but it finesses the issue of recognition by analysis versus analysis-by-synthesis. We are aiming at a system/device/hierarchy that understands, rather than recognizes speech, built on a strict PCT model. (I don't believe that you can get human-level speech recognition by itself, without understanding).

Phase 1, which is to begin construction when the Mac-based Control System Editor is in a suitable state, hopefully in a week or so, is the syntax recognizer. It is based on the Little Baby (LB--the self-teaching version of Bill Powers' Little Man). The Little Baby project in its original form was stopped when we could not get Chris Love's contract renewed, but it is being redeveloped by Allan Randall to become the syntax recognizer. Here's how we hope it will work.

The LB consists of a hierarchy of ECSs of the ordinary kind; no tricks such as models of the world or internal feedback (imagination loops). Each ECS has a perceptual input function (PIF) that contains a time- difference operator as well as a state vector. That is to say that its sensory inputs may be seen as the perceptual signals of the level below (or the direct signals from its sensors, at the bottom level) plus the time derivatives of those signals (actually time differences, since this is a sampled system).

The environment of the LB is a 3-space in which a target point moves, and in which the LB can move a cursor. We are, initially at least, not worrying about stereo vision and arm dynamics, as Bill Powers' LM does. We assume that the LB can apply some "force" to the cursor, which has some "mass" and which lives in an environment with friction (or viscosity, since it is in a 3-space).

Any control hierarchy must have some kind of top-level reference signal set. The LB has as top-level references the requirement that the cursor be on the target in each of the three dimensions (i.e. 3 degrees of freedom, represented in at least 3 top-level ECSs--there may be more, but they won't be orthogonal if there are). We assume that the actual deviation between target and cursor is knowable through some magic sensor system that does not need to bear any obvious relation to the sensory inputs to the low-level ECSs, which come from "non-magic" sensors such as eyes).

As the target moves through the space, the LB has to learn what to do to track it; it must reorganize. We recognize two forms of reorganization in this particular experiment, reconnection or sign-inversion of output links, and continuous modifications of PIFs. Since the PIFs are connected as a multilayer perceptron, we know that there exists an algorithm for allowing them to form any desired categorization of the input space, provided that there are at least 3 layers (I think 3 is right, but I would have to look it up to be sure). We do not intend to use any existing algorithm, but it is comforting to know that the problem is known to be soluble. What we intend to do is to modify the PIFs based on the ability of the individual ECS to control its perceptual signal (remember that output reorganization is happening at the same time, so we may run into all sorts of instabilities here).

The above is a first cut at how we intend the LB to learn. There are other possibilities, some of which have been discussed here in the past.

Now let's think about the behaviour of the LB if and when it has learned to control all its perceptual signals. What will we see in its PIFs? It seems clear that if the motion of the target is random (no correlation between its positions from one sample to the next at the Nyquist rate), the time-derivative aspect of the PIFs will be useless, and will have evolved to have zero weighting. But if there is correlation, then the time-derivatives will be useful, and will have some weight that relates to the perceptual signal in a way that tends to make the output go the right way to track where the target is likely to go.

Now consider what makes the target move--the disturbance in the environment. Firstly, we think of the 3-D environment as a "feature space" such as the feature space often associated with phonemes, but with only 3 features. In this space, we define the locations of the symbols of some alphabet, so that there are a finite number of points specified by being the "canonical" feature values of some symbol. Next, we define a formal syntax whose terminal elements are the symbols of this alphabet. By specifying the probabilities at each selection point in the grammar, we make a machine that emits grammatical strings with determined probability distributions for sequential patterns at each level of grammatical abstraction. So the target moves in ways that are non-random at several levels of abstraction (2 or 3, to begin with).

When the LB has learned, it should be able to provide from higher ECSs reference signals to lower ECSs that correspond to the probable movements of the target within the grammar (i.e. whichever abstraction is currently active). It's not the same as analysis by synthesis, even though the LB should be able to move its cursor in grammatical trajectories among the symbols, if it were to be driven by random top-level reference signals instead of the zero signals that correspond to "passive" listening.

The idea is that there is no such thing as "passive" listening to speech or to anything else you are tracking and might be controlling. The LB, and, I think, human listeners, do not perform analysis-by-synthesis in the sense of imagining what the articulators might have to do to form the incoming sounds. They can extract the various levels of abstraction through the PIFs. But they also can control and track, meaning that the PIF is used more to represent the deviation from expected sensory input than to represent the actual, raw input. In this system, "Pres'n" in the appropriate context, is a perfectly adequate sensory input for the listener to hear "President."

I'm not sure if this short summary is enough to make clear either what we are trying to do or why it finesses the issue of analysis versus analysis-by-synthesis. I hope it gives some clues. And I think it ties quite closely with what Tom and Andy are talking about.

Martin

Date: Thu Jul 08, 1993 10:16 am PST
Subject: Responding to CSGnet

[from Gary Cziko 930708.1750 UTC]

I have performed some magic so that all responses to new CSGnet messages (today and thereafter) will by default go back to CSGnet and not just to the individual who sent the original message.

This has always been the case for almost all CSGnetters except for a few like Tom Bourbon and Bill Silvert whose local systems generate their own "Reply-To" fields which will now be ignored.

Of course, you can always send a personal reply to any CSGnet message by changing the addressee field from CSGnet to the e-mail address of the individual sender.--Gary

Date: Thu Jul 08, 1993 10:31 am PST
Subject: Phoneme detection

[From Bill Powers (930708.1145 MDT)]

Tom Bourbon and Andy Papanicolaou (930707.1653) --
Rick Marken (930707.1900) --
Martin Taylor (930708 11:45) --

Some very good thinking about phoneme recognition. Martin's project is really quite exciting, because it's going to use reorganization with control systems -- something we've just barely touched. And Martin brings up a problem that has to be kept in mind -- sloppy speech, and the fact that we have to bring context into it in order to decide which of all the possible words might actually have been said. This is really a multilevelled problem, not just a problem at the level of phoneme detection.

However, work at the phoneme level will be worth while, because the more that can be accomplished at this level, the less work will be left for higher levels, or reorganization, to accomplish. If the phoneme detectors can at least report what WAS heard in a reasonably human fashion, at least that much ambiguity will have been removed.

The basic problem is the one Tom has stated: to what does the perceptual signal correspond? Rick pinned the problem down to the nature of the input function. What we want is an input function that will provide a signal indicating the degree to which (Martin would probably say "the probability that") a given feature of speech is present. However, the best approach may not be to try to get the phoneme perception in one step. What we need is to have a set of signals that indicate different aspects of the speech pattern, as independent from each other as possible. Then these signals can serve as inputs to higher-level systems that perceive functions of these signals. This is really what multilayer perceptrons do, except that they form the intermediate levels of signals at random, and using only a single principle, weighted summation. Martin is going to add transition information, which I think is the right next layer.

In furtherance of my own thinking on this subject, I've got my sonograph program running again on my new machine. I want to try some ideas for perceptual functions. But I am having a silly and frustrating problem with displaying the sonograms. When I did this before I had a monochrome VGA monitor, and producing grey scales was easy. Now I have a fancy SuperVGA system, and I'm going nuts trying to figure out how to create a variable-intensity black-and-white (or at least one-color) display. I know the answer is right here in front of me, scattered through three books, but I just can't get the old brain in gear enough to put it all together.

SO: if anyone can give me a hand with this I will be EXTREMELY grateful. I just want to write pixels that have variable intensity in one color or white, instead of variable colors. Preferably 256 levels, but 64 or even 16 levels will do. This is like working on your Nobel Prize speech and not being able to find a pencil anywhere in the house. turbo C 2.0, by the way.

By the way, Rick, the click when you just say the /d/ is only at the moment of release. Before that, the voice starts up with a continuous sound that precedes the release. If if doesn't, you get a /t/. I'm curious to see whether the microphone picks that up, because it's pretty much an inside noise. Tom's ascii sonogram seems to indicate that it is picked up.

Best to all -- Bill P.

Date: Thu Jul 08, 1993 1:35 pm PST
From Tom Bourbon and Andy Papanicolaou [930708.1421]

Rick, your reply contained some nice points. Some of them seem to reinforce and clarify ideas we tried to present in our posts and we indicate those in our reply below. But there was another issue lurking in the background in our original posts, one we were waiting to address until after we saw if there was any interest in the general topic. (So far, you are it. The reaction might be limited, but at least it is high quality.) That other issue has to do with whether PCT might be used, not just to model speech perception-production *after* someone else identifies the relevant parameters of the speech waveform and of the perceptual process, but as a source of testable ideas about how the perceptual process is organized and about how it functions. You spotted that possibility and addressed it. It seems to us that, if the possibility is real, then someone might be able to use PCT to design and build a better speech recognizer-producer, sell it, earn some money, and, as a side effect, finesse all of the academic debates.

>[From Rick Marken (930707.1900)]
>
>Tom Bourbon [930707.1653] and Andy Papanicolaou --
>
>>To duplicate my intention, the model must have a
>>specific value for the reference signal, p*. I decide to define
>>my perceptual experience of "just below the target" in terms of
>>pixels, which are units of vertical resolution on the computer
>>screen.
>
>I think you are descibing the cart before the horse. The first
>thing you did was define the perceptual variable, p, by defining
>the function that computes it:
>
>p = pixels(t-c)
>
>The perceptual variable, p, is the number of pixels between target
>(t) and cursor (c) -- as vertical distance. p is just a number;
>your function (which I am calling "pixels"), because of the way it is
>computed, defines that number as the signed vertical pixel distance
>between cursor and target. This was the first BIG step in your
>modelling; you are guessing that the subject is controlling a particular
>perceptual variable defined by the "pixels" function of the objective
>situation rather than some other one (for example, you could have defined
>p as (t+c) or sqrt(t-c), etc.-- pretty silly choices, but possibilities
>nevertheless).

We have no problems with anything you said there. Your presentation made, more clearly than ours, the point that in the modeling of tracking we (anyone) first "define the perceptual variable, p, by defining the function that computes it." We left that step implied in our post, a mistake if our goal is clear communication. And it is true that the adoption of [p := t - c], to use in the model, is a "guess," no matter how intuitively appealing that choice might seem. There was no way to know if the choice was a good one until after testing the model, with that choice of p (and the accompanying p*) installed, then comparing the results of the model's tracking performance with that of people. (Incidentally, as silly as some other choices for p might seem, we do know from basic psychophysics that certain transformations of the "objective" distance, c-t, would more closely approximate human perception of such distances, therefore, if such a transformation were included in the function that computes p, it might increase the level of agreement between the model and a person.)

> The problem with tracking tasks is that the most profound
>part of the modelling is the simplest -- defining the perceptual function.

Agreed. And of course, when they realize that to be true, many would-be supporters of PCT balk, then put the theory aside. They recognize that there is *much* work to be done on identifying and modeling perceptual functions.

>If you build models to control other variables -- like size or shape, which
>are functions of two independent variables, x and y -- then the
>"definition of p" problem becomes more interesting -- ie. p can
>plausibly be x*y, y/y or x+y, etc.

Agreed. And everything you are saying drives home more forcefully the problem we raised originally: for some things, speech for example, identifying, defining and computing p is "interesting" to the extreme.

>Now, to make the model work you have to select a value for p* -- a free
>parameter. This is just another number -- but it must be selected so that,
>when p=p*, the corresponding objective distance between t and c is
>as close as possible to what was observed in the experiment.

Yes.

>>Now I intend to detect /d/s in a stream of speech and to say /di/
>>every time I hear a /d/. I do so to my satisfaction.
>
>And, again, the big problem (as I see it) will be to design the
>perceptual function:
>
>p = /d/(acoustical waveform)
>
>where I have named the perceptual function /d/.

Yes. As we said in our posts, that is where we see the big problem, also. If behavior is about perception -- if it is all perception -- then the biggest task ahead for developing a "mature" PCT science is gaining a solid understanding of perception. Once you know the perceptual function, you can plug that into the loop, turn it on, and let it run.

>.. Where the "pixels" function
>transformed t-c into a time varying scaler representing vertical distance
>in pixels, /d/ transforms the acoustic waveforms into a time varying
>scaler representing degree of "/d/-ness" in the waveform in unknown
>units.

Yes, and that says a lot. For example, about this thing called, "a time varying scalar representing degree of "/d/-ness" in the waveform in unknown units," what is the "/d/-ness" that is to be transformed?

> To pick off the /d/s in continuous speech, I would create a new

>perceptual variable -- pd -- which is the difference between
>p and a perception that is a "model" of /d/ -- call it p'. p' is
>the "imagined" output of the /d/ function -- the result of creating
>the maximum output that that function can produce. p' is still
>just a number but it's a number that corresponds to the most /d/ like
>perceptual signal. I would then make pd the input to another control
>system -- the "decision" system. The output of this system is the
>word "/di/" whenever pd is greater than or equal to the reference
>input to this system. The reference for this system, pd*, is set
>so that the system responds "/di/" only to those parts of the waveform
>to which the subject responds "/di/". Of course, this decision making
>control system (that controls pd) is part of another control system
>that is trying to match the perceived output ("/di/") to pd; when
>pd is >= 0, then perceived output should be "/di/", other wise not.

This is the bigger issue we mentioned at the top. Is it possible to use PCT loops *inside* the perceptual function boxes in a model for speech perception? Does anyone out there know if such a thing has been done? (Bruce Nevin, Avery Andrews, Keith Deacon, anyone else?) We are not current on all of the literature on such models, but when we look for material on them, much that what we see looks disappointingly similar to what we find in textbooks from the 1960s and 1970s. On the other hand, if that is really the case, then there would be room for a PCT-driven program to develop a working model. For example, if PCT systems like the one you mention were in the perceptual function boxes of a model, then they would work more like *creators* of distinctive features than like detectors, and the overall loop could use its articulators to produce matches between its perceived and imagined-created /d/s. (A question: The way you described it, wouldn't your /d/ detector put out its biggest signal when there is not a /d/ in the input variable? If [pd := p - p'], then pd is 0 when p and p' match, and pd is maximum when p-p' are dissimilar. Are we reading it correctly?)

Does anyone else think there is merit in such an undertaking? Has it already been done?

The following remarks from you are about the horse that we put behind the cart -- they are about the perceptual signal which we mentioned after discussing the reference signal.

>> To which aspect of the world must our p* conform? That is the question we
>>encounter, and fail to answer, every time we try to reason our
>>way through this problem.
>
>I think you are really asking "to which aspect of the world does p
>conform?"; p* implicitly corresponds to the aspect of the world
>represented by p.

Agreed.

> This may be a fairly academic distinction but
>it might help focus your efforts -- the problem of modelling
>speech (in this case) is really a problem of determining the
>perceptual function, /d/, which may be a function of the outputs

>of several lower level perceptual funcitons -- equally difficult to model.

Agreed, as we have been saying from the start. (And it *is* an academic distinction, but still useful.)

>> The
>>perceptual signal for /d/ should be an analogue of ... of what?
>
>YES. That's the question.

YES.

>I have a suspicion that the /d/ function is based on more than just
>the frequency transition part of the syllabic waveform. For example,
>I think the /d/ in /di/ depends on something about the entire waveform.
>I believe that the /d/ in /du/ also depends (in the same way) on the
>entire waveform. So you don't hear /d/ until the entire waveform has
>occurred (because the output of the /d/ function doesn't go to max
>until the whole waveform has occurred).

Yes, and that was the idea Liberman and the other people at Haskins were dealing with, many years ago, when they began talking about "context conditioned variation." It seems that for the field of speech perception-production, as for behavioral science in general, people were able to feel, and spin theories about, isolated parts of the same elephant.

>There is pretty good evidence
>that this is the case. If you just play what appears to be the /d/
>segment of the waveform it just sounds like a click. The acoustic /d/
>part of /di/ sounds like a different click than that in /du/ so there
>is no perceptual constancy of the "/d/" part of the acoutic waveform
>alone -- even though teh waveform is the result of the same articulation
>pattern (and even if the articulation is done by yourself; try it; just say
>the /d/ part of /di/ and /du/).

All of that is certainly true and is behind many of our remarks in the first post on speech, the one with the fancy ASCII spectrograms.

>Speech recognition must be at the point where they have pretty successful
>phoneme detectors. A fairly speaker independent version of the hypothetical
>/d/ function must have been built already; one may be commercially available
>in voice recognition computer input systems. Maybe you could find a company
>that makes these and ask how they do it -- or go to the voice recognition
>literature. I betcha that a fairly good /d/ detector exists and can be
>used in your model of a person mimicing /d/ when they hear it.

So what is the story on this point? Are there already such devices? Exactly how good are the very best available voice recognition systems? Exactly how do they create their perceptual functions? Does anyone have any current information on this subject, and if you do, are the big problems already solved, or is there room for a PCT alternative?

Until later, Andy and Tom

Date: Thu Jul 08, 1993 2:17 pm PST
Subject: Synonyms for alignment

[From Kent McClelland (930708)]

Bill Powers, Greg Williams, Bill Cunningham, Gary Cziko, Bob Clark, Hank Folson, Tom Bourbon (Did I miss anyone?) (various posts). . .

Thank you for all the suggested synonyms for alignment, though I'm not quite ready to give up on the original term, which I see as having some useful connotations (aligning oneself with a group or a cause, etc.). But in any case, whether we call it equivalence or functional equivalence or harmonization or compatibility or operational equivalence or whatever, the analysis of how people often end up doing the "same thing" or seeing the world in the "same way" seems like a pretty important goal for a PCT-based sociology. Bill Powers has made a couple of interesting points in this discussion:

Bill Powers (930703.2100 mdt)

>The reason I want to say "equivalent" instead of "aligned" is
>that the concept of alignment begs a question that nobody is in a
>position to answer: are my perceptions like yours even when we
>agree we are perceiving the same thing? Equivalence is violated
>often enough that the most probable answer is "no." And as a
>practical matter, actual alignment of perceptions and reference
>levels isn't required; all that is required for cooperative
>action and effective communication -- within a limited domain --
>is equivalence. Equivalence can be achieved even when perceptions
>are radically different.

Bill Powers (930607.1800 MDT) [I think it's meant to be 930706.]

>Any talk of equivalence or alignment seems to imply a third party
>judging "objectively" whether two other people have the same
>goals or perceptions. We need some way to say that for ANY
>observer, goal equivalence is judged, basically, by the Test.

I guess I'm almost as interested sociologically in perceptions of alignment, perhaps mistaken ones, as in real equivalence, especially since true equivalence of control systems and reference levels is intrinsically so rare, as several of Bill P.'s and Tom's posts have clearly pointed out. I can think of a couple of instances in which perception might not match the "objective reality" but the mismatch would have important sociological consequences.

In one generic case, an individual imagines himself to be in alignment with another person and doesn't have the time or interest to do any extensive testing of that hypothesis. Operating on the mistaken assumption of alignment, this trusting individual carries on as if the equivalence were real, and only a very grievous failure of equivalence is enough to call for a reassessment. As you're no doubt getting tired of hearing me say, ethnomethodologists have analyzed such situations in some detail and make a pretty convincing claim that a good bit of social life is carried on in just that fashion, and that it works well enough for practical purposes in most cases. Only when we are confronted with incontrovertible evidence do we reject the common-sense assumption that others see the world just as we do and that any apparent discrepancies will eventually clear themselves up. (The other side of this coin shows a situation in which the individual encounters someone else, notes some obvious differences in race, sex, social class, or the like, and leaps immediately to the assumption that alignment of any kind with the other person is inconceivable.)

Another generic case is that in which the individual imagines himself to be acting entirely independently but ends up doing exactly what almost everybody else is doing. You look at the sky, decide that it would be a great day to take off from work and head for the beach, and find to your surprise that the freeway is jammed with others who had the same happy thought. I suggested in an earlier post that cultural alignments must be blamed for such predicaments.

In trying to fashion a PCT explanation of these and similar situations, I know I'm venturing beyond the normal bounds of PCT, which would be to focus on a single individual and look at the determinate way in which that person's behavior was an outcome of controlling certain perceptions. What seems paradoxical to me is that the higher levels of human control are achieved from experience in a social environment in which such mistaken assumptions of equivalence (or nonequivalence) apparently run rampant. Thus, I imagine, these ambiguities and illusions get built right into the individual's psychological makeup. And my conclusion is that we won't ever understand any given individual without getting a handle on the social milieu, as well. (Is my sociological chauvinism getting totally insufferable here, or what?)

Cheers, Kent

Date: Thu Jul 08, 1993 3:32 pm PST
Subject: Speech Model

[From Rick Marken (930708.1500)] Bill Powers (930708.1145 MDT) --

> work at the phoneme level will be worth while, because
> the more that can be accomplished at this level, the less work
> will be left for higher levels, or reorganization, to accomplish.

Yes. I was thinking that one way to identify possible acoustical correlates of the phonemes is in terms of the output variables that would have to be varied in order to produce acoustical variations that change the perception from one phoneme to another. I think the acoustical linguists have a pretty good idea of what these output variables are: I don't know what they are but I think they are variable aspects of the vocal tract that, changing over time, change the nature of the acoustical signal coming out of the mouth. I think there are vocal tract models that let you vary a few parameters (which would be the output variables of the PCT speech model) to create speech-like acoustical signals. Such a vocal tract model will eventually have to be part of a PCT model of speech -- either a mathematical version or actual tubes, reeds and pressure sources.

Each output variable that determines a parameter of the vocal tract model should probably be part of a separate control loop. So there will be a perceptual input associated with each output variable going to the vocal tract model. It seems to me that the perceptual function associated with each of these outputs should compute a perceptual signal that represents the aspect of the acoustical wave that is most strongly influenced by the vocal tract parameter with which it is associated. I think there might be a relatively easy way to determine what perceptual function to associate with each vocal tract parameter. It would be based on Martin's linear learning scheme. The acoustical outputs of the vocal tract could be played into a bank of octave band filters (the basilar membrane). The outputs of this bank of filters are the inputs to the perceptual functions associated with each vocal tract output variable. The weights on the inputs to each perceptual function are then adjusted (using the learning scheme) so that variation in the output of each perceptual function is maximum when the variation in the acoustical input is caused mainly by variation in the output (vocal tract parameter) associated with that perceptual function. So the weights are set so that the variance in perceptual signal 1 is greatest when the variance in vocal tract parameter 1 is greatest; the variance in perceptual signal 2 is greatest when the variance in vocal tract parameter 2 is greatest, etc. I think there are about eight vocal tract parameters so there would be eight perceptual functions (this means a minimum of eight inputs to each perceptual function -- an eight filter basilar membrane).

When learning is complete, you should have a set of nearly orthoganal control systems that control perceptions corresponding to nearly orthogonal effects of the outputs on the acoustical signal. Now these control systems can be used to produce more complex acoustical events. This is done by varying the reference signals to the control systems; the reference signals now specify the desired perceptual effects of vocal tract variations. One way to proceed at this point is to use your own perceptual system to determine how to vary the references to this bank of control systems in order to produce a particular sound; that is, convert the acoustical waveform that is generated by the vocal tract model into an audible signal. Then see what kinds of sounds result from manipulating the reference signals that you send to the bank of control systems that you have created. When you can produce a familiar sound by manipulating these references, then you have a time varying set of perceptual signals (the perceptual signals that were specified by your references) that must now, somehow, be combined in a second order perceptual function. And this is where I run out of ideas. I don't know how to create that second order perceptual function -- but I would at least know it exists because I am perceiving the results of such a function in my own brain.

I would love to try this approach. Does anyone have any vocal tract model code sitting around the house?

Best Rick

Date: Thu Jul 08, 1993 3:47 pm PST

Subject: Re: Phoneme detection, orthogonality

[Martin Taylor 930708 18:00] (Bill Powers 930708.1145)

I didn't mean to imply that it is useless to do as much as you can with phoneme recognition. Some phonemes, in some conditions, are relatively reliably recognized, and no doubt there will continue to be improvements in technique. All I wanted to say was that the concensus seems to be that the information needed to determine all phonemes is not in the speech waveform. Some of it resides in the language knowledge of the listener.

>However, work at the phoneme level will be worth while, because
>the more that can be accomplished at this level, the less work
>will be left for higher levels, or reorganization, to accomplish.

Yes, provided that the recognition output is not asserted to be more assured than it really is. One of the biggest problems in this kind of work is the tendency to take the most probable result as the correct result at too low a level of abstraction. It is better to say that a sound has 70% /a/ness and 30% /o/ness than to report to the next level that it is an /a/ if it isn't.

>If the phoneme detectors can at least report what WAS heard in a
>reasonably human fashion, at least that much ambiguity will have
>been removed.

It's not clear what would be meant by reporting what was heard in a human fashion, because what the human consciously hears is so dependent on expectation and linguistic knowledge, and, as Bruce Nevin keeps pointing out, the contrasts with other possibilities in the same context.

One kind of human experiment is to present people with little chunks of speech bounded by masking noise, and to try to get them to identify the chunk in some way (e.g. did this come from /dog/ or /log/). (You can't do this by clipping the speech chunk with silence at both ends, because the ear hears clearly that the chunk did NOT come from either, but from something that might, perhaps, be /op/, because an unreleased /p/ is close to a sudden silence.) If the phoneme level system produces results like the human on such tasks, that might be what is meant by "reporting what was heard by the human."

>The basic problem is the one Tom has stated: to what does the
>perceptual signal correspond? Rick pinned the problem down to the
>nature of the input function. What we want is an input function
>that will provide a signal indicating the degree to which (Martin
>would probably say "the probability that") a given feature of
>speech is present.

No I wouldn't. Not this time. The core concept of the system is that the symbolic nature of the syntax or of the speech is converted into a multidimensional set of degrees to which features are present. I have the feeling that control within such a space should be a much more robust way of recognizing the different levels of abstraction than the classical AI methods.

The normal approach to speech recognition, and one implicitly followed by all the contributors to this discussion, is to categorize elements of the speech. Some recognizers do this at the phoneme level, some at a syllable or diphone (phoneme-pair) level, some at a word level. But all of them make categories and then perform logic on the membership of the various categories, based on the continuous input waveform. I want to reverse this concept. The Phase 1 syntax recognizer takes a string that is known to be formally and symbolically defined, and turns it into a trajectory in a continuous feature space. Thereafter, we never extract symbols from the string at all. The control system tracks (predicts, we hope) the locations in the feature space, and at different levels of abstraction, we get the current (and projected) degree to which the speech has different features at that level of abstraction. Perhaps those features will be identifiable with phonemes, syllables, or words, but perhaps not.

It doesn't matter, within the PCT context, whether words are recognized or not, provided that the results provide the right kinds of disturbances within the larger hierarchy. The results of ANY level of abstraction are perceptual signals, which can serve as inputs to the perceptual input functions of other ECSs. It really shouldn't matter that one ECS, whose PIF happens to be a template for the concept "excellent" has a high output while another, whose PIF corresponds to "very good" also has a high output. The idea get across, even though the words may not be uniquely recognized.

=====================

I want to make a smooth transition here to another topic that often comes to the fore in my mind: orthogonality.

Bill says:

>What we need is to
>have a set of signals that indicate different aspects of the
>speech pattern, as independent from each other as possible. Then
>these signals can serve as inputs to higher-level systems that
>perceive functions of these signals.

I don't go along with "as indepent from each other as possible," though I agree with the rest. Let's go back a year or so to a posting that has stuck in my mind:

(Bill Powers 920722.0800)

>There is a way of turning control systems off in a nervous system that is
>very simple. It starts with the realization that neural functions are
>always one-way -- that is, neural signals can't go negative. In the
>standard diagram, we have error = reference - perception. This means that
>the perceptual signal is inhibitory, the reference signal excitatory. If
>the reference signal is simply set to zero, there is no excitation of the
>comparison neuron, and no amount of inhibitory perceptual signal will ever
>make it fire. So this comparator will produce zero error signal if the
>reference signal is zero, regardless of the amount of the perceptual
>signal. The control system is turned off.

>To get two-way control about zero, a pair of control systems is always
>required in the nervous system. The pair of systems treats opposite
>directions of change of the perception as positive, and the error signals
>in the paired systems have opposite effects on the controlled variable. The
>simplest example is a pair of opposed muscles and their associated spinal
>control systems for controlling force. If the arm exerts a leftward force,
>the left-controlling control systems sense and control a positive force to
>the left. If the force is to the right, the right-controlling control
>systems sense and control a positive force to the right. This much you'll
>find in BCP.

[Me, now] This means that we CANNOT be dealing with ECSs that are orthogonal. (I know, the pair ACTS like a two-way single-axis ECS, but it isn't built as one, and as I will argue below, that matters)

>To think of this pair of control systems as a single control system that
>can exert a continuum of forces passing through zero, we must think of the
>reference signals as a balanced pair. If the rightward reference signal is
>nonzero, there is a force to the right. As this reference signal declines
>toward zero, the rightward force declines toward zero. Then, just as the
>rightward reference signal reaches zero, the leftward one begins to rise,
>and the force begins to rise toward the left.
>
>If both the rightward and leftward reference signals are zero, this pair of
>systems is turned off. A disturbance may cause an inhibitory feedback
>signal to arise, but because there is no excitatory reference signal
>reaching either the leftward or rightward comparators, there will be no
>error signal to drive either of the pair of outputs. The system will not
>resist disturbances in either direction.
>
>In order to achieve control of an arm in the state of zero net force, it's
>necessary to add a common-mode signal to the pair of reference signals. Now
>the "resting" state is that in which both control systems contain error
>signals, causing the muscles to pull against each other. The net left-right
>force is zero, but any force disturbance will cause one error signal to
>decrease and the other to increase, so there is control in the vicinity of
>zero NET force. Both control systems in the pair are receiving nonzero
>reference signals now, with only the difference in magnitudes showing up as
>a net left or right force.

We could easily be talking about the 2-person rubber-band situation, here. Both people can pull, and they both have to, if the knot in the band is to be kept on the reference location.

>The common-mode force is, of course, called muscle tone. A control system
>that controls for muscle tone controls to keep the SUM of the two positive
>force signals at a constant level. A second control system can then control
>to keep the DIFFERENCE between the same two force-signals at another
>reference level, which sets the net sideward force to left or right. The
>difference-controlling system has to emit a pair of output signals that
>vary in a complementary way; in fact it must also have a balanced pair of
>comparators in order to handle positive and negative errors. The higher-

>level muscle-tone control system can be a one-way system, because the sum
>of the muscle tensions can never be less than zero.
>
>In our work on the arm model, Greg Williams found a reference that provided
>force-length curves for various muscles. These curves can be fitted quite
>closely with a second-power function over most of the force range (below
>the saturation level of tension). Muscle tension is produced by stretching
>the passive component of the spring, so muscle tension goes very nearly as
>the square of the driving signal and the amount of contraction.

Bill's posting continued with the equations that showed:

>If the rest of
>the system is linear, the loop gain of this force-control system is
>linearly proportional to muscle tone, and the differential force at
>constant muscle tone is a LINEAR function of the differential contraction
>in the two muscles (until one muscle or the other totally relaxes).
>
>This is why there is no control when you totally relax all your arm
>muscles, which means setting muscle tone to zero. An external disturbance
>will not produce any reaction; the arm will just give way and swing like a
>dead fish.

There's more, but what gets to me is the notion that everything has to be done by a vector balance between ECSs that have opposed outputs. This is normally called "conflict." Both one-way ECSs have sustained error. Such a situation is often discussed as a problem that the system will resolve by reorganization. But here, we see it as a central requirement for a properly functioning hierarchy. The degree of error in the individual ECS actually determines the gain of the two-way virtual ECS made from the pair. It can't be reorganized to avoid conflict, without destroying the operation of the hierarchy.

Now suppose that we are again talking about the rubber-band situation, but that the problem is moved out of the single dimension, by moving the reference point sideways. The subject cannot keep the knot over the reference point without changing the direction of pull, something a single one-way ECS cannot do. It just pulls. But if a third person has a rubber band connected at the same knot (or even nearby), the three of them can keep the knot over a reference position that moves in a plane, provided it stays within the triangle formed by the end points of the bands (allowing for the unstretched length, of course). Four people can together do it more easily, and a whole circle of people with rubber bands connected together at a central point can do so very easily (provided their references for the position of the knot are the same). Now, if we look at the forces on the circle of rubber bands, they look very like the outputs of motor neurons in the brain that are tuned to particular directions, when a move is (to be) made in a particular direction. Neurons tuned to nearly the move direction fire strongly, neurons tuned in widely different directions hardly fire at all. The vector sum is correct.

To come back to the speech problem, I assume that the PIFs will be parts of one-way ECSs, and that there may be opposed pairs, but more probably there will be ones that "pull" in all sorts of directions in the space at any particular level of the hierarchy (the different levels don't define the same space, because of the nonlinearity of the PIFs that generate the sensory input to the next higher level). The "identity" of the input at any particular level is defined here by the vector sum of all the directions in which the perceptual signals are "pulling." And that can be anywhere, not necessarily in the same direction as is defined by any particular PIF.

I do not assume that it is necessarily a good thing that the PIFs define orthogonal features, though I expect that to some extent they will come to do so. (Actually, I expect that they will cluster around principal component directions in the data space, but not to be limited to those directions). For different kinds of percept, different basis functions may be optimal--at the phoneme level, fricatives have a very different set of components than do vowels, for example. It would be (probably) more efficient to have sets of PIFs that are preferentially sensitive to directions that occur often in the data space than to force them into a set of orthogonal direction.
================

I have no idea whether this speech recognition project will work. It won't even be speech recognition within the term of the current contract on which Allan is working. The syntax recognizer and the associated experiments with the analogues of noise and coarticulation will be quite enough for the next year or so, and probably longer. What I'm looking for is a proof of concept. I think PCT shows the way to a robust way of developing perception of "nearly formal" systems such as speech, and if that works, it should scale well to the real-life problem of speech (I hope).

Martin

Date: Thu Jul 08, 1993 4:00 pm PST
Subject: Science

Dag Forssell (930708 1440)]

I have now read THE UNNATURAL NATURE OF SCIENCE: Why Science Does Not Make (Common) Sense by Lewis Wolpert Harvard University Press $19.95, and find it very reasonable and quite delightful. In my post (930630 0910), I shared the book review by LEE DEMBART, which Rick Marken (930630.1400) reacted to:

>Great find. Wolpert gets it right about science and then seems to fall
>apart when he gets to behavioral science.
>
>>"Science always relates to the outside world, and its success
>>depends on how well its theories correspond with reality,"
>
>Correct.
>
>>To the relativists and post-modernists, Wolpert issues a simple and
>>elegant challenge: What alternate theories do you have to explain the
>data?
>
>We know how they answer to that one: change the subject.

>
>>Moral and political problems are outside its purview, as are
>>justice, happiness, love and ultimate values. These matters are in
>>principle beyond the range of science.
>
>Wrong on that one. Why do scientists think that? Especially ones who
>have held brains in their hands. Is it just that human behavior is the
>last place where materialists can clutch at the remnents of a lost
>spirituality?
>
>>Furthermore, some systems, such as human behavior and society as a
>>whole, are so complex that "knowledge in these fields is barely at
>>the stage of a primitive science."
>
>Then why are they were treated as though they were simpler than a ball
>rolling down an inclined plane? Perhaps the people who are dealing with
>these purportedly complex phenomena just haven't got a clue about what
>they are looking at -- and none of the tools to understand it.

Wolpert argues that the role of science is to develop knowledge. The decision to use the knowledge is political, not scientific, and the best approach in the long run is to inform the public as well as possible.

Wolpert illustrates this by discussing the thorny issues of development of the atomic bomb, eugenics etc.

When it comes to the social sciences, Wolpert's view is clear. A quote from the chapter on Non-Science follows: P 134-135. I feel that the reviewer failed to convey the nuances here. You will note that Wolpert is in our corner by classifying the social sciences (which do not relate to anything else and cannot be falsified) as non-science:

....
It is also not possible, at the present time, to do any experiment at a lower level of organization - that is, at the level of brain function or neurophysiology - which would contradict psychoanalytic theory. Current explanations of dreaming couched in neurophysiological terms or computer analogies provide no explanation of the content of specific dreams. How, then, could one show that there is or is not an id or an Oedipus complex? At present it is not possible to relate the ideas of psychoanalysis to any other body of knowledge: they are entirely self-contained.

The current situation in psychoanalysis is in some way similar to the study of embryonic development in the eighteenth century. The claims of the rival theories of preformation and epigenesis could not have been resolved at that time because the state of biological knowledge and of technology were both inadequate. It would be hard to deny that the eighteenth-century embryologists were scientists: they designed experiments and made observations to the best of their ability - their science was simply premature and primitive. Both of the rival groups had enormous difficulty in accounting for the emergence in the embryo of highly organized patterns and forms, and invoked the idea of a 'building master' or 'vital force' or just assumed that everything was preformed. These were essentially ad hoc inventions, effectively having the same complexity as the phenomena they were meant to explain. In this sense they resemble the ego, id and super-ego and the emotions assigned to the unconscious.

Those engaged in psychoanalysis are dealing with a much more difficult problem than embryonic development. Not only is the subject-matter so much more complex, it is not easily accessible to experimental investigation in the way that embryos are. It is not known what equivalent to the 'cell' is required for understanding human behaviour, or even whether such an equivalent exists. Psychoanalysis is much worse off than eighteenth-century embryology.

One should be suspicious of ideas, like those of psychoanalysis, which have been so easily incorporated into our everyday thinking. If the rest of the physical world follows laws quite different from common sense, it would be surprising if the workings of that most complex of organs, the brain, could be so readily understood. For example, recent studies show just how unnatural the workings of the brain are with respect to language. Vowels are handled in a different way to consonants, and verbs and nouns are stored in different regions, as is shown by brain damage in specific regions. Even inanimate and animate nouns are categorized.

It can be argued that human behaviour and thought will never yield to the sort of explanations that are so successful in the physical and biological sciences. To try to reduce consciousness to physics or molecular biology, for example, is, it is claimed, simply impossible. This claim is without foundation, for we just do not know what we do not know and hence what the future will bring. No matter whether analogies between computers and the brain are correct, ideas about the problems of thinking and brain function have been greatly influenced by them. A characteristic feature of science is that one often cannot make progress in one field until there has been sufficient progress in a related area. The recent advances in understanding cancer were absolutely dependent on progress in molecular biology.

....
Best, Dag

Date: Thu Jul 08, 1993 6:56 pm PST
Subject: Replying to CSGnet

[from Gary Cziko 930709.0237 UTC]

Lee Dickey has brought to my attention that replying to CSGnet as a NetNews (Usenet) newsgroup (where we are "bit.listserv.csg-l") is a bit different than via the listserver.

For news groupies:

>A Reply goes back to the sender only.
>A Followup goes to the whole newsgroup [CSGnet].

I hope this is all straightened out now.--Gary

Date: Thu Jul 08, 1993 8:54 pm PST
Subject: Instructional feedback & higher level tests

[From Tom Hancock (930707)] Rick Marken,

I took you up on your suggestion: I performed the experimental task again myself and introspected as I did so. In addition, I reread my posts and each of your posts to me (930607.1100), (93061.1400), and (930702.1030). These tasks were quite helpful. Thanks for your input; my understanding of PCT continues to improve.

However, I feel quite uncertain of your having accurately understood my work. I wrote a lengthy reply and description of the perceptual variables involved in my experimental work, but since your driving concern in life is not understanding my research I believe I'll refrain from the lengthy post (unless you really want it) and just settle for a few briefer comments and other queries on your comments (930702.1030).

RM
>There's that "use of feedback" again. Control systems don't
>use feedback; they control it.

In the experiment I reported, instructional feedback is post response instructional information (Did you inadvertently confuse this with the feedback loop in a control system and thus misinterpret most of my study?). I propose that instructional feedback messages are used to stimulate two categories of perceptions which are controlled and are testable in the rough.

RM
>Did the independent variable account for 99% of the variance
>in the dependent variable for ANY subject? If not, I'd forget
>about publishing this until I got some real results.

Rick, your insistence on .99 correlations seems narrow and limiting for me. I respect that this is possible (and laudable) with the tracking studies that you have run. And I concur that this is necessary for laying the foundations of PCT. But this is not presently possible with a typical training task where a complex of controlled variables are operative at once and where the references are likely to be quite dynamic. I am interested in understanding higher level functioning in educational contexts, and funding agencies seem to be interested in applications. If I controlled for .99 to the exclusion of my other concerns then I might be out of work and also some consumers might not receive beneficial applications of PCT. And who knows, we might not make some of the potential progress needed in understanding control at the higher levels.

RM
>The group subject data is random noise.

Rick, I am familiar with Runkel's book, and I respect the need for precisely fit models; but please help me see how consistent trends for ALL subjects in a group (albeit low correlations) and also R squares of .57 or .31 can be random noise.

RM
>I would suggest that you work with ONE person, keep trying
>variants of your study until you can identify a variable that is
>unquestionably controlled by this subject, and then develop a
>working model of this control process.

That's a good recommendation and I hope to pursue it with SIMCON, but will I ever get unquestionable control when I deal with dynamic educational settings. Do you have any other suggestions along this line (in addition to forgetting about group statistics)? I have previously headed that way with a modeling software called STELLA. But recently I have been controlling more for publication: I confess my sins, but I feel I must persist this summer for survival sake.

RM
>You seem to be very attached to this kind of experiment --
>where a person rates his or her certainty that their response
>was or will be correct.

Partly I am interested in it because it was pivotal in my dissertation three years ago. Also, it is an easy way to monitor what is going on within the subject--the amount of sensed error; it seems to be an efficacious way to involve the subject in the human computer interaction--monitoring controlled variables for subsequent adaptive instructional feedback; and it is somewhat reliable (both in accuracy of predicting actual correctness and in associated response latency patterns--means).

RM
>...it's hard to test for control when one aspect of the
>controlled variable is an imagination. But it seems like that's
>what you are interested in --what subject's do about
>perceptions that have a large imagination component.
>"Confidence" about the correctness of a response is a
>perception, but it must be largely self generated -- an
>imagination.

Yes, that is what I am interested in. But I see an imagination component as part of the controlled perception and another as part of the signal for references, while the certitude estimate is an indicator of the error signals across several levels. The instructional feedback times are taken as evidence for the subject's studying the words and images in the instructional feedback message (creating vivid associations) in opposition to the error signals (indications of a lack of vivid associations and a lack of match between present time perceptions and the retrieved or imagined memory of previous trials which serve as the reference signal). That opposition comes in the form of prolonged perceiving of the associations between the stimuli and the correct name label (described in the feedback message) and/or the creation of more vivid/distinct perceptions of the associations. In both cases the imagination component could be involved. The point here is that there is opposition of some sort indicated by the feedback study times, and that opposition bears a systematic relationship to certainty estimates and outward correctness.

RM
>But, maybe it is based on perception of error signals, as you
>suggest. In that case, one can disturb this perception by
>giving false reports about the status of the subject's
>answers.

Yes, I ran a similar study (designed by my academic feudal lord) where false reports were randomly given to subjects. The subjects were told they were wrong but were actually correct and exhibited feedback time patterns similar to actually being wrong and vice-versa, being told they were right when they were not (grouped data).

And a comment from your post of 930614.1400.
RM
>How can I observe the degree of integration of prior
>perceptions into a distinct perception?

The cognitive psych lit (e.g J. R. Anderson) would indicate that a correct response to a higher level task which is rapid should be from a base of distinct or well integrated perceptions (no fan effect), while the same outward response if more delayed should indicate less integration and more interference from other prior perceptions. In the data I have gathered, certainty ratings of 100% are usually associated with the most rapid responses and the persistent trend (albeit much lower than .99) in several studies is an inverted U function across the range of certainty ratings. I take these 100% certainty responses to be indicative of distinct and well integrated perceptions.

And finally in that same post you gave some excellent examples of the typical control by students in classes. I have informally tested all of the ones you mention and have discussed the results frequently with my students. In addition we talk a lot about their goals and control when they are sitting down with a text book. They usually finish the term indicating that they have learned how to understand themselves better and that they are more efficient students and learners. It was the typical difference in goals of students that was one of the reasons for the experimental work I reported earlier: 1. some students really want to learn and understand, 2. some just want to be correct and do what I tell them to do , and 3. some want something unrelated to learning or outward performance. The way they handle feedback was my test of what they are controlling for. Goal 1 if they spend time with feedback if they have been correct but uncertain. Goal 2 if they spend time with feedback only when they have been incorrect but not when uncertain, and Goal 3 when the way they treat feedback does not relate to certainty or outward correctness.

Thomas E. Hancock

Date: Fri Jul 09, 1993 9:48 am PST
Subject: Re: what is "p" an analogue of?

[From Bill Powers (930709.0900 MDT)]

Tom Bourbon and Andy Papanicolaou (930708.1421) Martin Taylor (930708 18:00)

I have a crude sonograph running as of last night (intensity variations accomplished by varying dot spacing). It uses my A/D board and so can sample only 8000/sec (maximum of 4KHz bandwidth) which is barely adequate for a low voice. It takes about 10 min (on a 486/33) to show the spectrum for 1.5 sec of voice with enough frequency resolution to show the harmonics of the voice fundamental as separate bands. If progress warrants, I'll buy a Logitech sound board (about $150 now, mail-order) which can sample up to 41 KHz.

You asked whether there might be feedback loops within a perceptual function. One of the first things I found I had to do was to keep the envelope of the waveform at constant amplitude. The little control system that does this operates by varying gain. An unexpected development was that the lag of the control system can lead to emphasizing consonants. Even in this crude setup, the result is a clear (visual) differentiation between /bi/-/bu/, /ti/-/tu/, /di/-/du/, and /mi/-/mu/. Vowels are reduced to constant amplitude, but the loudness information is contained in the output signal from the gain control system.

The next step will be to build (program) tunable filters that will actively track individual frequency bands (the sonograph setup is just for getting information -- I don't think the brain uses that scheme). I think we will need about four or maybe five filters, one of them covering a band encompassing the first two harmonics. If the frequency filter for the fundamental and first harmonic can track these frequencies, the signal that does the frequency adjustment can also adjust the other filter frequencies. They will probably be broad-band filters, so the adjustment can be approximate. This will reduce the signal to constant frequency, too, with the adjustment signal containing the mean frequency information (inflection). So then we will have inflection and loudness information, with the remaining audio signal containing harmonics of constant vowel amplitude and frequency (relative to filter frequency) for further processing to work on.

At the moment, I'm visualizing a small number of active filters (which perhaps lock onto any frequencies in their vicinity), each filter providing a signal indicating amplitude within its band. These signals can then be combined by an array of higher-level input functions to generate parallel vowel-recognition signals. This is where Martin's perceptron approach might take over for a system that learns to perceive vowels.

The key to this approach is fairly simple. A filter is simply an oscillator with damping. Almost any kind of oscillator will serve; I use a sine-wave oscillator made of two integrators. The natural frequency of oscillation determines the center of the pass-band, and the damping determines the bandwidth. The C code is

```
void filter(int *input,    /* pointers to data arrays */
            int *output,
            float freq,    /* center frequency of filter */
            float damp,    /* bandwidth adjustment */
```

```
            int ndata)     /* length of data arrays */
{
float freqfact,y,x,avg;
float out2,out1,out;
int i,j,temp;
 freqfact = twopi*freq * dt;
 y = 0; x = 0.0; avg = 0.0; out = out1 = out2 = 0;
 for(i=0;i<ndata;++i)
  {
   temp = input[i];
   for(j=0; j < 4; ++j)
    {
      x += temp + freqfact*y - x*damp;   /* filter */
      y -= x*freqfact;                   /* filter */
      avg = fabs(y) + fabs(x);           /* rectifier */
      out2 += 0.01*(avg - out2);         /* output smoothing */
      out1 += 0.01*(out2 - out1);
      out  += 0.01*(out1 - out);
    }
   output[i] = 0.04*out* freq/400;       /* output array */
  }
}
```

The "input" array holds integer samples of the input waveform. The "output" array, of the same length, holds a smoothed and rectified signal indicating amplitude. The basic filter is the damped oscillator

```
      x += freqfact*y - x*damp;
      y -= x*freqfact;
```

If "damp" is small (like 0.01), and you initialize x to 0 and y to 100, you'll see a damped sine wave when you run this fragment.

To drive this damped oscillator, we add the value of the input signal, contained in "temp":

```
      x += temp + freqfact*y - x*damp;
      y -= x*freqfact;
```

Note that we're letting the oscillator go through four iterations for every sample of the input waveform. This keeps the oscillator working reasonable well up to half the sampling frequency. The value of "dt" is set elsewhere to be 1/4 of the sampling interval in seconds.

This particular oscillator contains a sine-wave and a cosine-wave (x and y). A full-wave rectification is achieved with the "fabs" (floating absolute) function; that flips the lower half of the sine wave to positive, so the sine wave becomes a series of bumps above zero. By adding the sine and cosine waves, rectified, we get a smoother output. The rectifier and filter that smoothes the bumpy signal to a smooth signal is

```
      avg = fabs(y) + fabs(x);
      out2 += 0.01*(avg - out2);
```

```
      out1 += 0.01*(out2 - out1);
      out  += 0.01*(out1 - out);
```

The signal "out" now contains a slowly-varying signal that indicates the magnitude of the filter output. Note that it is weighted by an amount that increases with frequency; this compensates for the falloff of the microphone signal with frequency.

The variable "freq" adjusts the center frequency of the filter. It adjusts the integration factors for the two integrators that make up the damped oscillator (filter). If "freq" is 1000, the center frequency is 1000 Hz on the time-scale of the real-world input. This won't be quite accurate for heavily-damped (broad- band) filters, but who cares? Everything's adjustable.

By adjusting the damping factor you can make the filter bandwidth whatever you like. I adjust it so the filter output can follow variations in amplitude reasonably well. Then I pick the slowing factor in the output section so the value of "out" closely follows the envelope of the amplitude variations without showing much of the individual oscillations.

To make the sonogram, this filter routine is called with the frequency set to a new value each time, and the output array is then plotted across one line of the screen at the appropriate vertical coordinate.

Oh, Tom, in case you want to try this: I've found that the RTD A/D board can be used this way:

```
for(i=0;i<MAXDATA;++i)
{
 do {b = inportb(0x241);} while (b & 0x80); /* wait not busy */
 audio[i] = ((inportb(0x241) << 8) & 0x0f00) +
                inportb(0x240);
 outportb(0x240, 0);   /* restart conversion, channel 0 */
}
```

This is all working up to trying to answer the question you raised a couple of days ago: how do we get "p" out of the audio input? I'm trying to figure out ONE way. From looking at the sonograms, we can see what the major frequency features are, and try to guess what kinds of input processing would be needed to pick them out. Then we just "build" what seem to be the required functions in simulation and see what they do. This isn't necessarily like the way the brain does it, but knowing one way that works (assuming it does), we can then look for other ways that also involve known features of the real neural networks.

I read somewhere that in the auditory nerve, the individual blips occur at audio frequencies. Signals representing frequency variations would be carried by much lower-frequency neural signals, in which the individual blips no longer have meaning, but only their mean frequency does. If a neuron received the audio-frequency blips, it would convert their mean frequency into post-synaptic potentials that varied slowly on the audio scale; they would produce output frequencies of a lower range, following the frequency envelope of the input blips. So this would work something like my rectifier and smoothing filter. A single neuron could serve as a tuned filter if it had positive feedback connections from its own output to its own input. While the actual mechanisms would be different from the computer program, the same principles could be working.

If the positive feedback mechanism (or the damping coefficient) occasionally went awry, the tuned filter could become a self- sustaining oscillator. That might account for the faint extremly- high-pitched whistle that aged ears like mine often provide their owners.

No time for other comments; back to work.

Best, Bill P.

Date: Fri Jul 09, 1993 9:50 am PST
Subject: Re: Tom Hancock's experiment

from Mary Powers [930709] Tom Hancock (93707)

Reading your post to Rick, I see this:

>Did you inadvertantly confuse this (instructional feedback)
>with the feedback loop in a control system and thus
>misinterpret most of my study? I propose that instructional
>feedback messages are used to stimulate two categories of
>perception which are controlled and testable in the rough.

I doubt if Rick confused anything inadvertantly. The real confusion lies in your use of the term "instructional feedback" as distinct from but all too easily confused with the feedback loop in a control system. This use of the term "feedback" is like bindweed in my flower beds - it's crawling all over the behavioral and social sciences and is just about impossible to uproot. The students are controlling for a certain input, but that input is not feedback - the feedack is the process by which they do the controlling.

All I can ask is that, if you are going to use the term "instructional feedback", to please NOT characterize your study as being about or using Perceptual Control Theory. PCT is a term used to distinguish the CSG brand of control theory from the bindweed kind. If you (quite reasonably) have to make a living cranking out a study to please funding agencies and full of group statistics and low correlations and stuff about instructional feedback, then by all means do so, but not under the aegis of PCT.

Mary P.

Date: Fri Jul 09, 1993 10:51 am PST
Subject: Instructional Feedback

[From Rick Marken (930709.0800)] Tom Hancock (930707) --

>I wrote a lengthy reply and description of

>the perceptual variables involved in my experimental work, but
>since your driving concern in life is not understanding my
>research I believe I'll refrain from the lengthy post (unless you
>really want it) and just settle for a few briefer comments and
>other queries on your comments (930702.1030).

I know it must seem like I am actively trying not to understand your research. I am not and I regret it if it seems that way; I think that the appearance of active resistence isa result of my shortcomings as a communicator and as an understander. I really do want to understand your research. But there are some fundemental obstacles for me -- that are my fault, not yours. The main obstacle for me is the data itself. For example, you say:

>Rick, your insistence on .99 correlations seems narrow and
>limiting for me. I respect that this is possible (and laudable)
>with the tracking studies that you have run. And I concur that
>this is necessary for laying the foundations of PCT. But this is
>not presently possible with a typical training task where a
>complex of controlled variables are operative at once and
>where the references are likely to be quite dynamic.

This is where we have a major disconnect. You accept the results you get as the best you can get. I think that if these are the best results you can get then they are not results at all. This is a basic philosophical difference. But you have the entire behavioral science on your side, as you suggest when you say:

>and funding agencies seem to be
>interested in applications. If I controlled for .99 to the
>exclusion of my other concerns then I might be out of work

Funding agencies want results and they will accept ANY result that can be summarized by a statistical number, even if that number shows plainly that the results are completely useless (from my perspective). You seem to agree with the funding agencies that a correlation of .43 is better than no correlation at all. I don't agree -- and, of course, I wouldn't get funded because I think no results are better than poor ones.

I also disagree with your rationale for accepting low quality research results:

>and also some consumers might not receive beneficial applications
>of PCT. And who knows, we might not make some of the
>potential progress needed in understanding control at the higher levels.

I don't believe that there will be benefits from or progress in PCT on the basis of low quality research results. Perhaps we could discuss this at the meeting. I know there are probably some in CSG who agree with your position. I'm willing to be swayed from mine, which is based on nothing more than intuition. I know that I can't do any meaningful modelling based on the kind of data you get in your experiments but I am only guessing that there is a way to study higher level control so that the results are as precise as those obtained in tracking tasks. Some people have taken some stabs at it -- and I would rate one or two as very successful. But, in fact, there are no "tracking type" results for control of the kind of higher level variables that you seem to be interested in.

>Rick, I am familiar with Runkel's book, and I respect the need
>for precisely fit models; but please help me see how
>consistent trends for ALL subjects in a group (albeit low
>correlations) and also R squares of .57 or .31 can be random noise.

I think Gary Cziko could explain this best. "Random noise" is hyperbole -- but when there is this level of inconsistency in data there is really no way of being sure of what any particular subject is controlling and why. If you have a copy of my mindreading program, why not do a few runs which show the disturbance-output correlations for each variable. Note that these correlations will occasionally reach up into the .4-.6 region for uncontrolled variables. If you were doing a study where you looked only at these correlations you would likely conclude that you have made an important discovery about a variable involved in behavior. In fact, the correlation is an irrelevant side-effect. This is why I worry about data that is not precise; it can look pretty good (and you might be able to make sense of it in you own mind) but it is very likely to be an irrelevant side effect. I don't think that it would be very useful to start telling people all these important conclusions based on PCT if there is the distinct possibility that these conclusions are based on irrelevant side effects. That's why poor data scares me; I think it's the door to misconception. If funding agencies want to fund that then there is nothing I can do about it -- but I don't have to participate.

>That's a good recommendation and I hope to pursue it with
>SIMCON, but will I ever get unquestionable control when I deal
>with dynamic educational settings. Do you have any other
>suggestions along this line (in addition to forgetting about
>group statistics)? I have previously headed that way with a
>modeling software called STELLA. But recently I have been
>controlling more for publication: I confess my sins, but I feel I
>must persist this summer for survival sake.

I want you to survive, for sure. So go ahead and do whatever they will pay you for. As long as it doesn't hurt anyone it's OK with me. But I suggest that, on the side, you try getting data the PCT way -- one person at a time. More importantly, keep playing around with your methods until you figure out a way to get perfect data from one subject. Dick Robertson has made some attempts at this; maybe he could mention them again on the net. I can't tell you how do do what you want -- all I can suggest is to get out of the mind set of thinking of an experiment and then doing it and just taking the results you get. If you find that a subject's ratings are not consistently related to their memory performance (or some other ratings) then try to figure out why not and CHANGE the way you do the experiment. Change the experiment daily or hourly if your must. But be willing to change in order to get good data; the goal is good data -- not maintaining preconceptions about the MEANS you should use to study something.

All this should be done, of course, in the background -- while you are collecting the "real" data to make the funding agencies happy. But if you can discover a reliable way to test for control of a higher level perceptual variable in just one person (so that you get a correlation of .999 between your disturbances and their outputs, say) then you WILL have a real PCT study to propose to the funcing agencies -- and you can confidently predict the results of your research and tell them EXACTLY what it means about the nature of human functioning.

I can't tell you how to do this because I'm happy study the properties of control systems controlling easy to quantify variables. But I think that the work on higher level control is EXTREMELY important -- in terms of theory, practicality and the visibility of PCT. I hope you can find time to pursue it. Maybe we can brainstorm on the net about possible ways to do this kind of research.

Best Rick

Date: Fri Jul 09, 1993 12:38 pm PST
Subject: Re: what is "p" an analogue of?

From Tom Bourbon and Andy Papanicolaou [930709.1355]

Bill, this is a quick reply, asking for some clarification of a couple of ideas in your post.

>[From Bill Powers (930709.0900 MDT)]
>
>Tom Bourbon and Andy Papanicolaou (930708.1421) --
>Martin Taylor (930708 18:00) --
>
>I have a crude sonograph running as of last night (intensity
>variations accomplished by varying dot spacing). It uses my A/D
>board and so can sample only 8000/sec (maximum of 4KHz bandwidth)
>which is barely adequate for a low voice. It takes about 10 min
>(on a 486/33) to show the spectrum for 1.5 sec of voice with
>enough frequency resolution to show the harmonics of the voice
>fundamental as separate bands. If progress warrants, I'll buy a
>Logitech sound board (about $150 now, mail-order) which can
>sample up to 41 KHz.

So, you got "the old brain" to work after all. Probably while burning midnight oil over the three books.

>You asked whether there might be feedback loops within a
>perceptual function. One of the first things I found I had to do
>was to keep the envelope of the waveform at constant amplitude.
>The little control system that does this operates by varying
>gain. An unexpected development was that the lag of the control
>system can lead to emphasizing consonants. Even in this crude
>setup, the result is a clear (visual) differentiation between
>/bi/-/bu/, /ti/-/tu/, /di/-/du/, and /mi/-/mu/. Vowels are

>reduced to constant amplitude, but the loudness information is
>contained in the output signal from the gain control system.

Here, we could use a little help so we can be sure we understand you correctly. First, the general method. We can imagine two different things you might be trying to do, and we bet you are doing only one of them.

Option 1: (a) you record 1.5 sec segments of speech, (b) you produce sonograms of the sounds, and (c) you run a control system to produce duplicates of the sonogram. In this option, the control loops are inside a modeled perceptual function and error signals from the loops are signals that are analogs of the acoustic waveform.

Option 2: (a) you record 1.5 sec segments of speech, (b) you use control loops to produce sonograms of the sounds -- sonograms that have specified features like constant amplitudes in distinct frequency bands, for example. In this option, once again, the control loops are inside a modeled perceptual function and error signals from the loops are signals that are analogs of the acoustic waveform.

Was one of these your method, or is there another option that we missed?

Whichever method you used, we are not clear about which distinctions you wanted to identify when you said:

"Even in this crude
>setup, the result is a clear (visual) differentiation between
>/bi/-/bu/, /ti/-/tu/, /di/-/du/, and /mi/-/mu/."

Is the differentiation you want to emphasize the one between members of each pair, for example, between /bi/ and /bu/? That would be a differentiation of vowels. Or is it between /bi/and/bu/, as one pair, and /di/and/du/, as another. That would be a differentiation between vowels within each pair, and between consonants across pairs.

Whichever way you answer either of our questions, you seem to jumping out to a strong start on showing how PCT can be used to work on a long-standing problem that interests many people. In fact, the flurry of articulate, imaginative and thoughtful posts from you, Martin Taylor and Rick Marken has outstripped our ability to produce a reply that is both immediate and coherent. We want to give that some deliberate thought over the weekend and will post something on Monday.

Until later, Andy and Tom

Date: Fri Jul 09, 1993 3:56 pm PST
Subject: To Mary, Rick, Bill, Dick

[From Tom Hancock 930709] Mary Powers (930709)

You are right I need to stop using the term instructional feedback. I had previously considered the phrase post-instructional information but I did not like the information part. Any suggestions?

Rick Marken (930709.0800)

I appreciate your helpful posture. And you have a point about having solid building blocks as the edifice of PCT is built.

Beside the fact that I am steeped in other methods I have been trying for a type of net casting to identify trends in previously uncharted waters and in addition, to test PCT predictions with the traditional analyses on a subject by subject basis. What's your feeling here?

I like your suggestion of changing the experiment daily if I must, until I can get more reliable results. Perhaps I'll plan for that this next school term.

No, I do not have a copy of the mindreading program--I did see it in Durango a couple years ago. I'd like to explore it as you suggest. Also how does one go about getting SIMCON?

Bill Powers

I just put in the mail (right away this time!) a copy of an article that will be printed in the August issue of Current Directions in Psychological Science, by Todd Nelson. I tried to ask his permission to send it to you, but couldn't reach him. I believe he wouldn't mind--he said he is open to (I was going to say feedback) comments. He told me he was on the net last fall for a while.

The title of the article is The Hierarchical Organization of Behavior: A Useful Feedback Model of Self-Regulation. I'll quote a few portions:

--(In the fourth paragraph of the introduction) Two decades ago, Powers postulated what he called a control theory of human behavior (BCP cited). Essentially, this theory holds that behavior is not the effect of antecedent causes in the environment, as many psychologists believe. Rather, behavior is the control of perception, in that humans seek to control via their behavior, the types of stimuli they perceive, such that the stimuli match internally held standards. Sensed deviations from a particular standard motivate a person to perform a behavior to change the environment, with the goal of causing the perceived environmental stimuli to match the standard.

A prevailing view in psychology has been that human behavior is a product of initial environmental influence. Behavior is thought to be the effect, and the environmental stimuli, the cause. However, many researchers have shown the utility of conceiving of human behavior as a cause of subsequent behavior in and of itself, in a closed-loop system (You are cited again here.)

--(Later in the paper) Carver and Schier's model adopts Power's proposal that behavior is organized hierarchically, in a cascading-loop structure...Empirical tests of the notion of negative feedback made it clear that there were some gaps in Power's ideas, and that, in order for the model to be complete, it had to take into account the effects of attention and outcome expectancy.

Then the rest of the paper describes Carver and Schier's model. Bill, perhaps a follow-up article or letter to that journal would be appropriate at this time for the reading audience.

Dick Robertson

Have you tried any higher level testing lately? I am familiar with the one reported by Runkel and what you reported at the annual meeting a few years ago.

Tom Hancock
14210 N. 56th Pl. Scottsdale, AZ 85254 (Home)

Date: Fri Jul 09, 1993 4:08 pm PST
Subject: Hancock research

[From Dag Forssell (930709 1255)] Rick (930709.0800) Tom Hancock (930707)

>I also disagree with your rationale for accepting low quality research
>results:
>
>>and also some consumers might not receive beneficial applications
>>of PCT. And who knows, we might not make some of the potential
>>progress needed in understanding control at the higher levels.
>
>I don't believe that there will be benefits from or progress in PCT on
>the basis of low quality research results. Perhaps we could discuss
>this at the meeting. I know there are probably some in CSG who agree
>with your position.

Just to proactively prevent any misunderstanding about where some of us who advocate the importance of social applications come from:

I agree completely with Rick that low expectations on results make for bad science. Bill has adressed this several times, as in the IV-DV discussion [Bill Powers (930428.0700)]. Tom, If you missed it, take a close look. Bill provided critique of examples of published studies, both good and bad.

To support my explanations of PCT, I use the simple tracking experiments, and computer demonstrations. I make it clear as best I can that extrapolations into the upper levels are conjecture at this stage of development of PCT and HPCT. But even so, this conjecture is much more sensible than anything offered in the social sciences today.

An understanding of the demonstrable basics of PCT and an individual extrapolation into uncertain territory of the upper reaches of HPCT allows application of PCT and generates many benefits, both personal and "bottom line", but is not research.

Best, Dag

Date: Fri Jul 09, 1993 5:47 pm PST
Subject: LIVING THINGS - RKC

[From Bob Clark (930709.0930 PM EDT)] Mary Powers (930704)

>But this is the point: the broad assertion that every living
>thing is a control system.

I don't think I understand this. Are you proposing:

>that every living thing is a control system.

This seems to be a generalization from observation, concluding that being a "living thing" correlates very highly with "being a control system." But this depends on the definition of "living thing," and I find several different versions of this definition. Such a generalization doesn't work very well when either item to be correlated is unclear. In any case, such a generalization can, in principle, be invalidated by a single exception.

On the other hand, this could be a suggestion for an alternative definition of "living thing." Thus: If it is not "a control system," it is not "living."

In my view, by any definition of "living" that I have found, a great many living things involve control systems.

In general, I conclude that any combination of parts that passes The Test with respect to at least one variable, includes at least one control system.

At the same time, a failure of The Test may merely be due to an incorrect selection of the variable to be tested. This is a one way logic. The object tested could, in fact, include one or more undetected control systems.

Further, you seem to suggest that the "point of it" is:

>to look at these creatures, and the cells that compose them from this new
>perspective.

I don't know what you mean by "this new perspective." Does this mean to apply The Test to "each creature" and to "each cell?" If so, fine and let's do it.

Back to "every living thing is a control system."

It is my understanding that plants, in general, are "living things." They come in many sizes and shapes from single cells (bacteria) to giant Sequoias. For convenience, consider everyday trees and shrubs, etc. I am unaware of any (physically existing) variable that such plants control. In addition, I don't find any mechanism by which the plant can affect a physically existinge variable. I also don't find any mechanism by which the plant can "perceive" changes in such variables. If you, or anyone, knows of such a plant and/or variable, I'd be very interested.

In a way, that is the reason I suggested the Venus Fly-Trap. It is classified as a plant, yet it catches and digests insects.

Rick Marken (930704.1400) adds:

>I don't think that one has to be an entymologist to see that the behavior
>of the flytrap involves control. All you have to see is that the apparent
>stimulus for flytrap behavior is affected by that very behavior -- there
>is a closed loop.

Yet the Fly-Trap is not an insect, it only "eats" them. Botanists who have studied this plant report its "catching" action (closing its leaf lobes) is initiated by the disturbance of at least two of three special hairs located on the leaf lobes. But the closing of the Trap and later digestion of the insect has no effect at all on the hairs that detected the presence of the insect. The reported details are rather more complete, but the original disturbance of the special hairs is never "opposed."

It is not a closed loop.

>>The operation of the Venus Fly Trap can be described without
>>need for control systems.

>So can the operation of human beings. But is it a good
>description? PCT says no, look again.

As far as I am concerned, "the operation of human beings" cannot be described without including control systems.

>How does apparent S-R behavior really work?

I am in no way advocating S-R interpretation of behavior. I do recognize that it is sometimes used as a convenient short-cut in certain communication situations.

I find it useful to distinguish between "S-R" interpretations and "cause-effect" interpretations. With the latter, one can trace the sequence of events through the interconnected elements in terms of recognized physical and chemical concepts. With "S-R" descriptions, the interconnections are at least omitted, if not entirely ignored. For some purposes, this may be a convenient abbreviation, but it is certainly an inadequate description of any feedback system.

>Ditto relaxation oscillators in coelenterates. I don't know what
>a relaxation oscillator is, but to say cells operate as though
>that is what they are begs the question of how they go about

>acting that way.

A "Relaxation Oscillator" (see your library) is any combination of components that act to accumulate some item until a test point is passed. When that point is passed, the contents are dumped, the test point is re-set and accumulation begins again. There are many examples of such oscillators, using many things: sand, water, electrical charges and heat. No control system is required, although control systems can demonstrate this behavior under some conditions.

I probably should not have referred to coelenterates, which are multi-cellular animals that have been studied extensively. (See your library.) With this example, I intended to point out that some multi-cellular animals do not have readily detected control systems, particularly regarding their externally oriented feeding/excreting systems. They are equipped with tentacles that bring food particles into the cavity, which is later emptied when digestion, etc is complete. These activities are readily accounted for without using control theory. In addiation, there may be some controlled variables somewhere, but they are not readily apparent.

I did not intend to apply the concept of relaxation "oscillators" directly to single cells, although I think it possible that some of them may have this characteristic.

>A kidney cell may look like it simply has an output
>function, in the context of the organ of which it is a part. But
>the cell itself lives an active metabolic life, maintaining
>itself as a kidney cell. Its output, along with all the other
>kidney cells, rids the organism of substances that would be
>poisonous, but that overall effect is a side effect of what the
>kidney cell itself is controlling for.

Can you use The Test to demonstrate some variable (or variables) controlled by the cell? Without such a demonstration, the assertion is empty. Assertions are easily made, but often hard to support.

>Since nobody on the net is an invertebrate physiologist, we'll
>be wondering about that for a while. But I think the question
>"suppose it's control systems all the way down?" is more
>interesting than saying "it's only a .... ".

An "interesting question," perhaps, but there is a large number of "interesting questions."

Regards, Bob Clark

Date: Fri Jul 09, 1993 8:17 pm PST
Subject: Phototropism; R-Squares

[from Gary Cziko 930710.0400 UTC] Bob Clark (930709.0930 PM EDT) said:

>For convenience, consider everyday trees and shrubs, etc. I am
>unaware of any (physically existing) variable that such plants control. In
>addition, I don't find any mechanism by which the plant can affect a
>physically existinge variable. I also don't find any mechanism by which the
>plant can "perceive" changes in such variables. If you, or anyone, knows of
>such a plant and/or variable, I'd be very interested.

Plants on my window turn toward the light. I turn them back so they face the room, but after a day or so they are again facing the window.

From my high school biology I seem to remember that the darker side of the stem grows faster than the side with more light which causes the plant to tilt toward the light, maximizing the the amount of light falling on its leaves (a side effect perhaps of phototropism, but not an irrelvant one).

I don't know where you would put the reference level and comparator in such a system, but since the pattern of light falling on the plant leads to a response which changes the pattern of light falling on the plant, there is a closed loop between input and output.
==

Tom Hancock replying to Rick Marken said:

>Rick, I am familiar with Runkel's book, and I respect the need
>for precisely fit models; but please help me see how
>consistent trends for ALL subjects in a group (albeit low
>correlations) and also R squares of .57 or .31 can be random
>noise.

An R square of .31 is 83% useless (coefficient of alienation is .66--square root of 1 minus R square). An R square of .57 is only 66% useless.

But "a consistent trend for ALL subjects in a group" promises to be a finding worth reporting. But I'm not sure what you mean by "trend" here.--Gary

Date: Sat Jul 10, 1993 8:40 am PST
Subject: Sonography

[From Bill Powers (930710.0930 MDT)]

Tom Bourbon and Andy P. (930709.1355) --

Your two guesses were

>Option 1: (a) you record 1.5 sec segments of speech, (b) you
>produce sonograms of the sounds, and (c) you run a control
>system to produce duplicates of the sonogram. In this option,
>the control loops are inside a modeled perceptual function and
>error signals from the loops are signals that are analogs of
>the acoustic waveform.

>Option 2: (a) you record 1.5 sec segments of speech, (b) you
>use control loops to produce sonograms of the sounds --

>sonograms that have specified features like constant amplitudes
>in distinct frequency bands, for example. In this option, once
>again, the control loops are inside a modeled perceptual
>function and error signals from the loops are signals that are
>analogs of the acoustic waveform.

You're way ahead of me. The control system for amplitude is applied to the raw acoustic wave (an "automatic volume control"):

```
      Raw wave ---> Gain ------->  output wave ---------->
                     ^                |
                     |                |-- amplitude perception
                      -- out funct --Comp
                                      |-- ref amplitude
```

I don't plan on using the sonograph as an input function -- it's just a way for me to look at the way frequency information is distributed in the voice signal.

What I'm doing here is busily re-inventing the wheel. Last night I discovered formants. My sonograph resolves individual harmonics of the fundamental voice frequency. When I do a glissando up and down through an octave, maintaining a single vowel sound, all the frequencies move slightly up and down the display (which covers over 6 octaves) and the harmonics spread farther apart -- but the clusters of bright patches on the sonogram remain in the same place! I thought that raising and lowering the pitch would move the positions of the patches, which is why I was thinking of frequency-tracking filters to eliminate inflection effects.

What the sonograph shows is not the frequency content of the voice sound generator, but the acoustic properties of the throat- mouth-nose cavity. When I hold my mouth to make an "O", the low- frequency patches alone show up, and remain in the same place when I run the voice pitch up and down. When I say "EEE", there is a low-frequency patch and two very much higher frequency patches with a big gap in the middle frequencies, and again they stay about the same as voice pitch goes up and down through a octave or so (I can display only up to 4 Khz, so I'm missing one patch which I can just see the start of).

I suppose that linguists have known this all the time, but it came as a big surprise to me. No wonder they talk about articulators! Obviously what is being controlled is the auditory signature of particular mouth configurations, excited by a harmonic-rich noise source. So the perceptual functions we want will be based on filters, with their output amplitudes combined to pick out the typical signatures of various cavity conformations. The relative positions of the formants (I may as well use the accepted term) are directly affected by the way the tongue, glottis, and lips are arranged. So the control loops will be quite direct; all we have to do is to set up perceptual functions that use weighted combinations of formant frequencies to produce signals that can be controlled by easily separable manipulations of the parts of the vocal apparatus. The frequency filters will have quite broad bandwidths, and locating them properly for a given person's formants only has to be done once.

The result will naturally look a lot like controlling the articulators, but that will be only because their positions are closely connected with the auditory signatures actually under control. It does mean that one can control kinesthetic sensations from articulators and get fairly close by imagining the sound that goes with them. But clearly, judging from the experience of deaf people, kinesthetic control is not discriminating enough to produce the clear sounds of speech. Deaf speech is like deafferented motor control: it can be done with a lot of learning, but the feedback makes all the difference for skilled performance. Maybe one thing that could come out of this would be a way to detect the signatures and use some other sense for the feedback channel. This is another reinvented wheel, but maybe viewed from a slightly different slant.

RE: the dscriminations:

I meant that the vowel in /ti/-/tu/ is the same as in /di/-/du/, respectively, but that the consonant in either /ti/ or /tu/ looks different from the consonant in either /di/ or /du/. But don't make too much of this. There's still a lot of fiddling to do, and my sonograms are none too crisp.

You asked for some of my Procomm-Plus scripts, and perhaps others might want them, too. These are ".ASP" files. They have to be compiled using "aspcomp.exe", to produce files with the extension ".ASX" . Just type "aspcomp <filename>" to compile.

GETMAIL.ASP

```
proc main
if not connected
 TRANSMIT "ATDT2477299^M"
endif
start:
if not connected
 pause 5
 goto start
endif
pause 1
transmit "^M"
waitfor "username>"
transmit "powers_w^M"
waitfor ">"
pause 1
transmit "c^M"
waitfor "username:"
pause 1
transmit "powers_w^M"
waitfor "password:"
pause 1
transmit "password^M"  ;(i.e., your password)
waitfor "$"
pause 1
```

```
transmit "MAIL^M"
when 0 "%MAIL-E-NOTEXIST" call logoff
waitfor "MAIL>"
call getdate
transmit "dir^M"
waitfor "MAIL>"
pause 1
transmit "EXTRACT/ALL "
transmit s1
transmit "^M"
waitfor "MAIL>" 60
transmit "del/all^M"
waitfor "MAIL>"
transmit "exit^M"
waitfor "$"
transmit "kermit^M"
waitfor "Kermit-32>"
pause 1
transmit "send "
transmit s1
transmit "^M"
getfile kermit
waitfor "kermit-32>" forever
pause 1
transmit "quit^M"
strcat s1 ";1^M"
waitfor "$"
pause 1
transmit "del "
transmit s1
waitfor "$"
pause 1
transmit "lo^M"
waitfor ">"
pause 1
transmit "lo^M"
quit
endproc

; This procedure gets YYMM.DDx where x ranges from A to Z,
; returns in s2

proc getdate
strcpy s9 "ABCDEFGHIJKLMNOPQRSTUVWXYZ"
Date s1
substr s2 s1 0 2
substr s3 s1 3 2
substr s4 s1 6 2
strcpy s1 "MAIL"
strcat s1 s4
strcat s1 s2
strcat s1 "."
strcat s1 s3
```

```
init n1 0
chdir "\bitnet"
dateloop:
strcpy s2 s1
substr s3 s9 n1 1
strcat s2 s3
isfile s2
if success
 inc n1
 goto dateloop
endif
strcpy s1 s2  ;s1 contains available file name
chdir "\pcp2"
endproc

; if error message or no mail, end mail, log out

proc logoff
transmit "quit^M"
waitfor "$"
transmit "lo^M"
waitfor ">"
transmit "lo^M"
hangup
quit
endproc
```

The other useful one just logs on:

```
proc main
if not connected
 TRANSMIT "ATDT2477299^M"
endif
start:
if not connected
 pause 5
 goto start
endif
pause 1
transmit "^M"
when 1 "username>" transmit "powers_w^M"
waitfor ">"
transmit "c^M"
waitfor "username"
transmit "powers_w^M"
waitfor "password"
transmit "yourpassword^M"   ; your password here
waitfor "$"
endproc
```

Your system will probably require different commands, but maybe this will get you started.

Best, Bill P.

Date: Sat Jul 10, 1993 9:49 am PST
Subject: Plant etc. control systems

From Bill Powers (930710.1045 MDT)] Bob Clark (930709.0930) --

>It is my understanding that plants, in general, are "living
>things." They come in many sizes and shapes from single cells
>(bacteria) to giant Sequoias. For convenience, consider
>everyday trees and shrubs, etc. I am unaware of any
>(physically existing) variable that such plants control.

There's some ambiguity here concerning how we identify control systems. At the level of organization at which we look at the nervous system, there are components which are not control systems (neurons, input functions, comparators, output functions) but which, when connected properly, constitute a control system. So whatever the level of analysis, there will be identifiable subsystems which are not control systems.

Also, at any level of analysis, there will be side-effects of control actions which are not themselves controlled variables. For example, one reliable outcome of visual-motor tracking is to produce lactic acid in the muscles. Another is to heat the bearing holding the control handle.

A human observer looking at an organism can't tell side-effects from controlled variables without investigating whatever control systems can be discovered. Science has so far looked at the behavior of organisms without considering the difference between intended and unintended effects of actions. Therefore the list of known behaviors is an unsorted jumble of control processes and byproducts of control processes: important effects and unimportant ones. Human observers pick the effects of most interest to them, not those necessarily pertinent from the point of view of the behaving system. What the human observer sees as the primary behavioral phenomenon may not exist at all in the world of the behaving system.

When a human observer looks at a phototropic plant (mentioned by Gary Cziko 930710.0400 UTC), he sees a movement in space related to the direction of the sun and to the surroundings. But it is unlikely that a plant knows anything of space or directions; the spatial orientation of the plant is clearly a side-effect of controlling some other variable. If there is a controlled variable, it must relate to something the plant can sense (sensing being simply the creation of an internal signal that depends on an external physical quantity). For the plant, this signal might be related to photosynthesis, or to chemical substances released by certain wavelengths in the solar spectrum, including radiant heat. Furthermore, there must be a differential effect that depends on the orientation of the plant relative to the direction of incidence of the solar flux. If the reference level for the difference in the sensed input in different parts of the plant is zero, no actual reference signal is required (in some plants the reference signal is variable: certain desert plants hold their leaves edge-on to the sun in the hottest parts of the year, but in winter hold them face-on). The error signal could differentially adjust the growth rate in parts of the stem, or in more active plants, the fluid content and thus the distention of certain cells that twist the stem.

Plants are sensitive to a great many physical quantities, and have a considerable range of behaviors that affect those same quantities. They are not, of course, quantities of much interest to a human being, but they are the important aspects of the world for the plant. Application of the Test would be quite possible. Of course this has never knowingly been done; instead, plant behavior is explained in terms of local cause-effect phenomena. But the same could be done with human beings. If one examines only the components of a control system, the overall organization can easily escape notice.

It's important to give the PCT hypothesis a thorough test with respect to all organisms, and with respect to organismic functioning at all levels. There have been many unrecognized instances of the Test. Injection of glucose directly into the bloodstream is known to shut down production of glucose by the liver; infusion of thyroxin immediately suppresses the Thyroid-Stimulating Hormone production by the pituitary, and if maintained will cause atrophy of the thyroid. Artificially reducing the cholesterol level in the bloodstream will result in a greatly increased rate of cholesterol synthesis by the liver. In fact, the major outputs of EVERY main organ system are suppressed when artificially elevated, and raised when the output is artificially depressed. Evidence of negative feedback exists everywhere in the body, even to enzyme-controlled reactions in which the activity of the enzyme is lowered by the final catalyzed product.

One of the great difficulties that PCT has comes from its very existence. The behavior of organisms has been explained over and over for 350 years or so: it is very difficult for either scientist or layman to accept the idea that a fundamental concept has been omitted from all these explanations. All explanations offered during the history of the life sciences have managed to do without the principles of control -- which means that a lot of persuasive garbage has been generated. It's often hard to spot what is wrong with an explanation until a better one has been found. But when the better one is found, those who have made a living by cloaking the wrong one in attractive and reasonable-sounding rhetoric face an embarrassing choice: reject the new explanation, or admit that they have been convincing people with false reasoning. A new idea wreaks a lot of damage to self-images.

Best, Bill P.

Date: Sat Jul 10, 1993 11:01 am PST
Subject: Closed Loop, Speech Model

[From Rick Marken (930710.1200)]

Bob Clark (930709.0930 PM EDT) --

>I am unaware of any (physically existing) variable that such plants control.

I think Gary Cziko (930710.0400 UTC) made a good point, saying:

>Plants on my window turn toward the light. I turn them back so they face
>the room, but after a day or so they are again facing the window.

I think these plants control a chemical signal that represents the amount of light (integrated over a considerable amount of time, of course) falling on the top of their leaves. It is pretty obvious that the physical correlate of this signal is controlled -- as Gary proved. I've always wondered why behavioral scientists left the behavior of plants out of their science. Plants generate all kinds of interesting purposeful behavior (like the phototropism) -- the time constants on these control loops are very long, but they are control loops nonetheless.

Bob Clark continues:

>Yet the Fly-Trap is not an insect, it only "eats" them.

Yes. The word "fly" was in my brain when I suggested asking "entymologists" about their sensors. I should have said "botanists".

> Botanists who have
>studied this plant report its "catching" action (closing its leaf lobes) is
>initiated by the disturbance of at least two of three special hairs located
>on the leaf lobes. But the closing of the Trap and later digestion of the
>insect has no effect at all on the hairs that detected the presence of the
>insect.

How could this be true? The hairs must transform the effects of the disturbing variable (fly mass?) into some kind of signal -- probably chemical -- that somehow causes trap closure and digestive secretion. So there is a "flytrap" equation that relates output, o, (trap closure, chemical digestive secretion) to the level of the chemical sensory signal, s:

o = f(s)

The sensory signal is caused by forces exerted on the hair cell. These forces are determined by the mass of the fly combined with the outputs of the flytrap. The mass of the fly decreases as it is digested and the trap itself, when it is closed, must exert some force on the fly, which effects the force exerted by the fly on the hair cells. So the chemical signal resulting from hair cell stimulation is (over time) a changing result of fly mass (the disturbing variable, d) and the effects on the fly of the outputs of the flytrap, o, so:

s = g(o) + h(d)

So we have two simultaneous equations describing a closed loop -- input effects output and output effects input. So the flytrap is unquestionably a closed loop system -- it exists in a closed loop relationship with its environment. If the product of the signs of the relationships around the loop is negative then this is a negative feedback loop -- and s (the sensory variable) is controlled. I think the feedback in the flytrap loop is negative because it is stable (over a long time scale) -- the long term value of s is kept near zero. This is a slow loop, as I said, but it is a closed negative feedback loop, nevertheless. The flytrap is controlling at least one (probably many) chemical - sensory variables.

>It is not a closed loop.

Is so! (I think?). The flytrap is a purposeful behaving system; it is controlling (sloooowly -- I had a pet flytrap so I know) some sensory consequences of its output effects on its own sensory input variables. Maybe behavioral scientists don't study plants because the are so slow. But I think they more than make up for it (compared to animals) in terms of docility, beauty and general cleanliness.

Bill Powers (930710.0930 MDT)--

> The relative positions of the formants (I may as
>well use the accepted term) are directly affected by the way the
>tongue, glottis, and lips are arranged. So the control loops will
>be quite direct; all we have to do is to set up perceptual
>functions that use weighted combinations of formant frequencies
>to produce signals that can be controlled by easily separable
>manipulations of the parts of the vocal apparatus.

Yes. But I'm wondering about what vocal tract variables are actually manipulated by the control systems. It seems like there are more than three output degrees of freedom (tongue, glottis and lips). The tongue has at least two, the glottis may just be one and the lips are probably two, at least. And maybe these degrees of freedom are not all physically independent -- I'm not sure I can do what I want with my tongue while I do what I want with glottis and lips. There must be models of the vocal tract that take into account real anatomical interactions. I think I remember drawings of vocal tract models in a book by G. Fant -- probably published c. 1968.

If we could make a plausible vocal track model then we could play the acoutical outputs of the model into your sonogram. By varying the parameters of this model one at a time we could see which perceptual (sonograph) patterns covary with these parameter variations and then try to build the perceptual functions (filters?) that should be hooked up as inputs to these output (vocal tract parameter varyinging) functions. This is an "eyeball" version of what I suggested doing in a previous post with what was essentially linear discriminant analysis. The linear disciminant approch could be based on the outputs of your sonograph filters and the perceptual functions would be linear weightings of these outputs -- making it unnecessary to build new filter functions. One problem with the linear discriminant apprach is that some of the features associated with vocal tract output variations are probably time dependent. But maybe we could then use a matrix of filter outputs, the columns being the value of the filter outputs during fixed time intervals.

>The result will naturally look a lot like controlling the
>articulators, but that will be only because their positions are
>closely connected with the auditory signatures actually under control.

I think this will be the main immediate result of this exercise -- showing that the coordinatoin of the articulation parameters is an expected "side effect" of controlling for the perceptual consequences (acoustic and kinesthetic) of these articlations. As I mentioned to Tom and Andy in a private post, speech recognition, though far from perfect, is a lot better off than we are likely to make it in the next couple of years. I have heard that there are speaker independent, continuous speech voice recognition systems for computer input (you talk into the computer and your speech is turned into text versions of what you are saying) that are reaching 90+% correct for certain kinds of material. I know that speech recognition improvements are asymptoting -- but there are apparently some awfully good candidates for the perceptual functions already in existence. Of course, part of the success of these systems depends on understanding "context" (as Martin noted) -- but the speech pattern recognizers alone do pretty well; good enough for the control of perception model of speech production, I bet.

Best Rick

Date: Sat Jul 10, 1993 8:54 pm PST
From: Dag Forssell / MCI ID: 474-2580
TO: Hortideas Publishing / MCI ID: 497-2767
Subject: Science fiction
[From Dag (930710 2200) - direct]

Greg, Thanks for your comment on myths type 2. I dropped the project.

I will appreciate your comment on the following. Don't want to post if this is suitable for reading in Durango. I may use it somehow in letters to explain and separate myself from the crowd.

SCIENCE, not FICTION:

The next behavioral science illustrated in a thousand words.
Are you limited by prescriptions of an obsolete science?

It is very exciting to participate in the development of a true science of life, based on engineering principles. It does require clear thinking, because both the explanations and the conclusions are different from what is widely understood (but incomplete or false) today.

With a clear science of behavior that lends itself to testing and validation, the path is clear to unprecedented progress in social, educational, managerial and leadership practices.

To get a feel for changes in science and our vision of a different future, let us go....

Back to the sixteenth century:

Imagine that we were born here and now study the science and practice of alchemy -- named for the art of making gold and silver. It is based on practical chemistry know-how, developed by trial an error over many centuries, and incorporates astrology, philosophy and mysticism. As a science it offers empirical, descriptive theories, prescriptions and recipes passed down from past generations of scientists. Alchemy works, and the accomplishments are undeniable. Just look at the great variety of products it it has given us: metals, metal plating, medicines and much more.

In the 1500's, we live in a society accepting and dependent on alchemy, where our scientists know what they know, are proud of it, respected, and authorities on their specialty. They write the textbooks (Gutenberg's printing press is a blessing) used in alchemy school, referee and edit the scientific journals. We cannot imagine a different science with different ground rules, different explanations and much better results, so naturally the vast majority of us using alchemy's teachings are proud of what we know and satisfied with the results we get.

Fast forward to the late twentieth century:

The science and practice of chemistry is now based on clear engineering principles - what we call generative theoretical models. Scientists can predict results and design new compounds even before they mix chemicals, because they have a carefully tested and validated theory that explains what goes on as the elements interact. When we think of alchemy, we recognize that the scientists who knew what to do in the 1500's, even though they offered explanations, had no clear or valid understanding of the underlying processes. We understand now that they could not know in detail why and how their chemistry worked -- when it did. Their descriptive theories have been forgotten and we smile a knowing smile when we hear about their quest to turn lead into gold. We recognize that it would take more than just a few minutes to explain about generative theoretical models, atoms and the periodic table of the elements to scientists who were not used to think that way and had never heard of them. -- No that is not right, they knew all about atoms, but not in the way we do now. That prior knowledge would only have made it harder for them to understand properly.

As a byproduct of the scientific revolution in chemistry, historians studying the 16th century approaches to metal smelting, alloying etc., can understand why they were successful with some processes but unreliable or failing with others.

Here in the late twentieth century, the sciences and practices of psychology are based on practical know-how, developed by trial and error (including statistical studies) over centuries. As a science it offers empirical, descriptive theories and prescriptions (where practice often has nothing to do with the theories) passed down from past generations of scientists. Psychology works after a fashion, and the accomplishments seem undeniable. Just look at the dozens of "treatment modalities" used by counselors, scores of (incompatible) leadership programs taught in industry and "common sense" acceptance in our culture.

In the 1990's, we live in a society accepting and dependent on our behavioral sciences, where our scientists know what they know, are proud of it, respected, and authorities in their specialties. They write the textbooks used in schools of psychology and sociology, referee and edit the scientific journals. We cannot imagine different behavioral sciences with different ground rules, different basic explanations and much better results, so naturally the vast majority of us using the teachings of psychology as we work with people are proud of what we know and satisfied with the results we get. Besides, promotions for "new and better" training programs are so plentiful we have learned to avoid them.

Fast forward to the twentyfirst century:

The behavioral sciences are now based on clear engineering principles - what we call generative theoretical models, including recognition of and an accurate explanation of the phenomenon of control. Parents, spouses, educators, managers, leaders, workers - all are better able to develop satisfying, productive lives, because all have learned a carefully tested and validated theory called Perceptual Control Theory (PCT) that explains what goes on. PCT is so basic to human growth and development, and so easy to understand, that the basics are taught in elementary school. When we think back, we recognize that the scientists who prescribed how to deal with people in the 1990's, even though they offered explanations, had no clear or valid understanding of the underlying functional relationships - how a mind works. We understand now that they could not know in detail why and how human relations (courting, parenting, education, supervising, cooperation) worked -- when they did. Their descriptive theories have been forgotten and we smile a knowing smile when we hear about their quest to shape the behavior of others. We recognize that it took more than just a few minutes to explain about generative theoretical models, control of perception and feedback loops to scientists who were not used to think that way and had never heard of them. -- No that is not right, they knew all about control and feedback, but not in the way we do now. That prior knowledge only made it harder for them to understand properly.

As a byproduct of the PCT revolution, historians studying the variety of 20th century approaches to education, leadership, and quality management can understand why some seemed successful, why others failed and what the human costs were.

Back again to the late twentieth century:

You have glimpsed the future. Are you satisfied in the present? You can take advantage of PCT without waiting for the whole world to adopt it.

Thanks, Dag

Date: Sun Jul 11, 1993 2:13 am PST
Subject: Re: Speech Model

[Martin Taylor 930709 11:10] (Rick Marken 930708.1500)

> I think there are vocal tract models

>that let you vary a few parameters (which would be the output
>variables of the PCT speech model) to create speech-like acoustical
>signals. Such a vocal tract model will eventually have to be part of
>a PCT model of speech -- either a mathematical version or actual
>tubes, reeds and pressure sources.

You are following an honoured tradition here. I think it was Alexander Graham Bell's father who built such a vocal tract model to demonstrate to his elocution students/patients what was happening when they made certain sounds. My friend Louis Pols in Holland was given a picture of himself fitted with a rudimentary mechanical vocal tract to celebrate his thesis defence. But nowadays, the models are more commonly software, with appropriate hardware means for exciting the air vibrations. One way of looking at LPC is called a log-area model, which in effect reduces the vocal tract to a series of parallel-walled cylinders of different lengths and diameters.

For speech synthesis of good quality, the tendency is to ignore the actual vocal track and work directly with the formants, using either a parallel or a series set of filter banks. A vocal impulse (which must be carefully shaped) is used to excite the formant filters, or else a noise source is used to produce the fricative elements (sometimes both together, as in /z/).

> I think there
>are about eight vocal tract parameters so there would be eight
>perceptual functions (this means a minimum of eight inputs to each
>perceptual function -- an eight filter basilar membrane).

It's not clear what a vocal tract parameter is, in the human. In the software, it's clear. Different models have different numbers. The order of magnitude is right, but I think 8 is a bit low (generally speaking there is a centre frequency and bandwidth for each of the first three formants, and a centre frequency for the next 2 (though these are often held constant). Then there may be separate fricative formant control, since the time course of the band-shape of the bursts is quite important in distinguishing some phonemes. And one must have a control for nasality, and then there are some parameters for the impulse shape, which may be ignored in some synthesis systems. Let's say 8 to 12 parameters in most cases.

> The acoustical outputs of the vocal tract could be
>played into a bank of octave band filters (the basilar membrane).
>The outputs of this bank of filters are the inputs to the perceptual
>functions associated with each vocal tract output variable.

Maybe. But one has to realize that the basilar membrane filters are not octave-band, and may well be active (i.e. under direct perceptual control, like muscle fibres--that's controversial and I expect most acoustic physiologists would disagree). Considered as passive filters, there is a well defined curve representing the critical band as a function of frequency. It is about constant up to around 500 Hz, and is proportional (more or less) to frequency above something like 1000 Hz. I seem to remember it being 1/6 octave, but it might be 1/3. The reason I forget is that I had an (unresolved) argument about the interpretation of these critical bands a few years ago. I think that they represent something like double the actual filter width, but the community as a whole disagrees. So I don't remember whether 1/6 was my estimate or the accepted one (I could look it up, but for this argument it doesn't matter).

A more important caveat is to consider the possible preprocessing that happens to these filter outputs before the results become part of a PIF. There is both temporal and frequency masking, and lots of non-linear effects such as the generation of sum and difference tones and the like. Beyond this, there is the question as to whether one is at all interested in the individual temporal variations of the different outputs.

A former colleague, Roy Patterson, has generated a software model of what he thinks goes on in the very early stages of hearing, the result of this processing being what he calls a "Stabilized Auditory Image," which doesn't change so long as the sound you perceive doesn't change. I like it, because in the image there seems to be some element that corresponds to anything you can hear, and what you can't hear is suppressed. It is the SAI that I would like to use as the input to any recognition engine, but it is VERY computationally intensive to produce (it can work on 2 TMS 320C30 chips, but not on much less). The parallel computation in the nervous system could produce it quite easily, though.

The SAI has many more instantaneous degrees of freedom than is apparent in the waveform, but it changes slowly unless the perceived signal changes fast, so it is not really adding spurious degrees of freedom. It is just that there is a lot to analyze in any image. I would think that this large number would have to be refined down to the 8 or 12 degrees of freedom for vocal tract control in the PIFs for speech recognition. (The auditory system didn't evolve to understand speech. That came a lot later, and there are lots of potentially controllable perceptual degrees of freedom out there in the auditory envirnment.)

>Does anyone have any vocal tract model code sitting around the house?

Not me, but I think you could probably pick one up from one of the companies specializing in software for speech research.

Martin

Date: Mon Jul 12, 1993 4:42 am PST
Subject: "All"?

From Greg Williams (930712) Bill Powers (930710.1045 MDT)]

>One of the great difficulties that PCT has comes from its very
>existence. The behavior of organisms has been explained over and
>over for 350 years or so: it is very difficult for either
>scientist or layman to accept the idea that a fundamental concept
>has been omitted from all these explanations. All explanations
>offered during the history of the life sciences have managed to
>do without the principles of control -- which means that a lot of
>persuasive garbage has been generated.

I am curious about why you persist with such extreme statements. They DO grab one's attention, true; and they DO make PCT sound absolutely revolutionary and a total break from all previous explanations in biology. But they also make you (and other PCTers, by implication) sound ignorant of the abundant recognition of control processes (yes, with reference signals internal to organisms!) in the biological literature. Take a look at Prosser's COMPARATIVE ANIMAL PHYSIOLOGY, for example, or some of the human physiology texts which recognize living control systems (Tom Bourbon has remarked on these on the net), or many articles in BIOLOGICAL CYBERNETICS and PHYSIOLOGICAL REVIEWS.

Not ALL nonPCT thinkers in biology have completely shunned living control models since Wiener and his colleagues originally proposed reference signals internal to organisms back in the 1940s. Many biologists (especially physiologists) make models which are "orthodox" in terms of the PCT "revolution." For those folks, the "revolution" is actually The Establishment -- they make their livings constructing and testing the same kinds of models as PCTers make -- in respiratory physiology, in acid-base regulation, in temperature regulation, regulation of food intake, etc., etc. Refusal to recognize the existence of PCT-type models outside of the CSG is a very big step toward being unable to learn from the rich array of those sometimes very detailed and well-tested models. I am sure that those who have been making such models outside of the CSG could learn much from you (who have certainly thought more deeply about living control systems than anyone else) and other PCTers -- but first the possibility of some common ground and interest must be appreciated on both sides.

PCTers certainly aren't obliged to rummage through the garbage of nonPCTers. But must all of the PCT-type models of non-PCTers be counted as "garbage"?

Sometimes embarrassed by your hyperbole,

Greg

Date: Mon Jul 12, 1993 2:17 pm PST
From: Hortideas Publishing / MCI ID: 497-2767

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Beep... Beep... Beep... Beep

Hi Dag --

Yes, the line is busy right now. My computer disk drive card went out and I'm frantically trying to get things working right again, plus I've got HI coming back from the printers to get folded and mailed, and I'm supposed (ha!) to be working on CL. In short, all hell is breaking loose here. How about next week or at the meeting for my comments on your past, present, future essay?

Going zonkers, Greg

Date: Mon Jul 12, 1993 2:50 pm PST

Subject: Control of perception

[From Rick Marken (930712.1030)]

Greg Williams (930712) to Bill Powers (930710.1045 MDT) --

>I am curious about why you persist with such extreme statements.

> they also make you (and other PCTers, by implication)
>sound ignorant of the abundant recognition of control processes (yes,
>with reference signals internal to organisms!) in the biological literature.

Speaking as an implied ignoramus (and I am, indeed, ignorant of much of the biological literature) I wonder whether, in all the "abundant recognition of control processes" that exists in this literature there is any explicit recognition of the fact that it is perceptual variables that are controlled. This is the essense of PCT -- and the main point of the application of control theory to the purposeful behavior of living systems: what is controlled by a control system is a perceptual representation of some aspect of an objective state of affairs. It seems to me that, if there were "abundant recognition of control processes" in biological science then biologists would be referring to this fact about control (and to Powers' PCT model based on this fact) like crazy.

I have run into a few applications of control theory in biology. In every case there was absolutely no recognition that the variable controlled by the control system was a perceptual input variable.

I get the distinct impression that the non-PCT applications of control theory in biology are like the non-PCT applications of control theory to manual control in psychology: control is looked at in terms of input-output transfer functions rather than as the control of perception. The reference signals that are internal to the organism in these models are seen as offsets in the I-O tranfer functions. These applications of control theory reveal some interesting characteristics of the control process but they miss one little point -- namely, that what is controlled by a control system is a perceptual signal. PCT is about why this little point is not so little.

Of course, if you know of references in the biological literature where a big deal is made of the fact that it is perception (not output) that is controlled, then I think that it would be VERY important for us to know about it (and whoever said it would DEFINITELY be a person with whom even a PCT fanatic hyperbolyte like myself could have a cordial conversation).

Best Rick

Date: Mon Jul 12, 1993 3:43 pm PST
Subject: Williams on Prosser

[From Gary Cziko 930712.2000 UTC] Greg Williams (930712) said to Bill Powers:

>I am curious about why you persist with such extreme statements. They

>DO grab one's attention, true; and they DO make PCT sound absolutely
>revolutionary and a total break from all previous explanations in
>biology. But they also make you (and other PCTers, by implication)
>sound ignorant of the abundant recognition of control processes (yes,
>with reference signals internal to organisms!) in the biological
>literature. Take a look at Prosser's COMPARATIVE ANIMAL PHYSIOLOGY,
>for example . . .

I had the pleasure of chatting with Professor Prosser a few months ago about the evolution the nervous system. At no time did he mention that nervous systems functioned as control systems.

I have also consulted his his book _Adaptational Biology_ where in his chapter on behavior he says nothing about the control of perception and devotes a lot of ink to "fixed action patterns," which he says completely accounts for the behavior of certain organisms.

I will take a look at his _Comparative Animal Physiology_. If you can direct me to certain pages where he reveals knowledge of how control systems operate in organisms, I would much appreciate it. Since he is on this campus, I would be pleased to have another control theorist to talk to.--Gary

Date: Mon Jul 12, 1993 4:10 pm PST
Subject: Applied PCT, Little Talking Baby

[From Rick Marken (930712.0800)] Dag Forssell (930709 1255) --

>To support my explanations of PCT, I use the simple tracking experiments,
>and computer demonstrations. I make it clear as best I can that
>extrapolations into the upper levels are conjecture at this stage of
>development of PCT and HPCT. But even so, this conjecture is much more
>sensible than anything offered in the social sciences today.

I know. You do an excellent job in your presentations. In fact, your way of showing how the science of control relates to "real world" applications of PCT is just what I think of as the right way to do applied PCT. I think we are on much more solid ground (in terms of applications) when we "advertise" PCT in terms of the scientific demonstration of characteristics of control rather than in terms of promised results of applying PCT. I think that your conflict demo and your description of how this relates to management practices is far more valuable than all the promises you could make about greater productivity (or whatever) that might result from applying PCT. This is all I was trying to say in my discussion of applied PCT.

Martin Taylor (930709 11:10) --

>It's not clear what a vocal tract parameter is, in the human. In the
>software, it's clear.

That's what I thought.

>A more important caveat is to consider the possible preprocessing that
>happens to these filter outputs before the results become part of a PIF.

I think of the filter outputs themselves as the output of a PIF (perceptual input function). Anything that transforms external or neural variables into another neural variable is a PIF from my point of view.

Since auditory perception is one of your baileywicks, Martin, why not turn your "little baby" project into one where the model baby learns to talk -- a "little talking baby". I think that a baby doing "crib talk" is doing something like what I've suggested in the last couple posts -- learning what perceptual variables it can control by varying the parameters of its vocal tract. Once the kid has built up the control systems that can control perceptual aspects of the acoustic consequences of "vocalizing" it can start using these control systems (varying the reference signals sent to them) to produce higher order perceptual variables -- phonemes and words. I think this process is modelable and would be an excellent demonstration of the power of PCT.

What do you think?

Best Rick:

Date: Mon Jul 12, 1993 5:08 pm PST
Subject: Example of speech recognition

[Martin Taylor 930712 14:15]

A colleague forwarded this to me. I think it might amuse and inform people who want to know the current state of the art in speech recognition (which, contrary to Rick's statement, does not seem to be asymptoting, since the best performance of systems tested by DARPA (now ARPA) has been improving on a more or less straight-line bases over the last several years.)

If I guess correctly, the system being used in the following is Dragon Dictate. The words must be spoken with a distinct pause between them, but the vocabulary is virtually unlimited. It is probably as near state-of-the-art as is available commercially, despite being a couple of years old.

==========================

Al Sicherman
Minneapolis Star Tribune
Sunday, June 13, 1993

I am writing this down for you Gentle Reader, even as I speak

As it tends to do, technology marches on. And it seems to be marching over me. I am dictating today's column into a device that changes my spoken words into typing on my computer.

Yes, that's right, I am sitting in my chair, with my hands folded in front of me; I am speaking into a little headset microphone and words are appearing on a screen. Ain't science grand?

At this point in today's column, I am correcting the frequent misunderstandings that arise between me and the machine so that what you are reading looks just fine.

In fact, however, the rather darling computer program that is interpreting my deathless words is even now making a zillion incorrect guesses about what I am saying, most of which aren't even close. I should acknowledge that its second guess is quite often correct, but we aren't playing horseshoes here.

The only reason you can make anything out of this is that I am correcting the machine as we go. To be fair, it is still in the process of learning my voice. It has only been listening to me for a solid month. Presumably after a lengthy exposure to my dulcet tones - say, 10 or 15 years - it would unerringly transcribe my every utterance. In the meantime, it's a little dicey.

I should be gracious enough to say that the reason I am pulling up a microphone - instead of a keyboard or a typewriter or a linotype - is that my hands (not unlike my feet, my back, my knees, my esophagus and my head) are failing to perform up to minimal expectations, and my doctor has recommended that I wear strange-looking wrist bands and do what I can to minimize wrist strain from typing.

All right, my choices are: Abandon what I laughingly call my profession in favor of something that doesn't use the hands, such as bubble-blowing or grape-stomping; ignore the doctor and go through the day with my wrists on fire, or spend my time dictating to a computer that thinks that when I say "require" I mean "retire."

It's an easy decision. The company has brought in this dictation computer on a trial basis; five of us are trying it. (The worst of it is that chewing sounds confuse it, so I can no longer eat while I type.)

OK, enough Mr. Nice Guy. Here, unedited, is how this device heard me recite a few familiar passages. I will correct the titles, but that's all:

The Raven
Once upon a midnight jury, well I powder, week and very,
Over many a right and serious volume of forthcoming more -
While I not, clearly next, suddenly their game a having,
As of some one gently wrapping, rapid at my chamber your.
"Kiss some Mr.," I mother, "having at my chamber or:
Only this and nothing more."
Coast the Reagan: "Everywhere."

Lincoln's Gettysburg Address
For store and 7 years ago our fathers wrote fourth on this content a
new nation, embassy in liberty and education to the protozoan that all
them are created people.

Annabel Lee
It was many and many the year uncle,
In a keynote by the see,
That a maiden there lived when you may no
By the name of animal Lee.

Preamble to the Constitution
We, the people of the united space, In order to form a more perfect union, establish justice, injuring most family, provide for the , defense, problem the general Walter, and severe the lessons of liberty to ourselves and or', to morning and establish this consideration for the united states of America.

Eleanor Rigby
Eleanor really picks up the race in the church
Where a wedding as in
Lives in a tree
With at the window
Wearing the face that she teeth in a jar by the your
Who is it for?
All the only people, where to they all, from?
All a only people, where to they all, from?

The Arrow and the Song
I shot an bureau into the hair,
It tell to earth, I new not where.

Paul Revere's Ride
Listen, by children, and you shall here
Of the midnight by of call radiator.
...
Want, if by land, and to, if by see;
And I on the opposite shore will be,
Ready to wind and sound the along
Through every Nelson says village and from.

The Star Spangled Banner
Old say can you see by the tongs early late,
What so probably we pale at the college last cleaning,
Whose broad strikes and great stores, through the parallels five
Or the reference we watch were so talented string?
And the rockets read letter, the follows bursting in air,
A group through the night that our flight was still their;
Go say does that star scheduled manner yet wave,
Or the land of the free, and the call of the great?

Though there are many more works of Enemy Lobster Although (whom you many know as Henry Wadsworth Longfellow), including The Religious Watchman (the Village Blacksmith), I think we should stop.

Maybe another time I'll read aloud some complete garbage (passages from the Congressional Record; the lyrics of "Louie Louie," or the fine print on my credit-card bill) and see whether the computer turns it into Shakespeare.
===================
End quote.

Martin

PS. Our internet connection is very uncertain. I understand that the Mississippi floods have something to do with the problem. I am making no serious attempt to contribute to CSG-L until I hear that the linkage has been, shall we say, solidified. And I am away for a month, starting July 16.

Date: Mon Jul 12, 1993 5:39 pm PST
Subject: More "All"?

From Greg Williams (930712 - 2) Rick Marken (930712.1030)

>Speaking as an implied ignoramus (and I am, indeed, ignorant of much
>of the biological literature) I wonder whether, in all the "abundant
>recognition of control processes" that exists in this literature there
>is any explicit recognition of the fact that it is perceptual variables
>that are controlled.

Well, you'd better look for yourself. You might be surprised. Be that as it will, be aware that it isn't common practice in biology to actually name, for example, input signals to biochemical sensors outside the nervous system as "perceptual" signals. (They aren't psychologists!) And reference signals usually get called "set points," instead. Nevertheless, I think that at least some (as opposed to NONE!) of the biologists who deal with negative-feedback control models are bright enough to realize that such models act to try to bring the difference between the sensed signals and the respective set points (INSIDE the organisms, according to their models) to small values. That is your definition of "controlling perceptual variables," isn't it? At least some do make a pretty big deal about the sensed signals corresponding with but not being an "objective state of affairs." But maybe not a big enough deal for you? If you have bones to pick with how much importance some accord to this "main point," why not pick them with the biologists themselves, rather than via my passing along what I see them saying?

>I have run into a few applications of control theory in biology. In every
>case there was absolutely no recognition that the variable controlled
>by the control system was a perceptual input variable.

>I get the distinct impression that the non-PCT applications of control
>theory in biology are like the non-PCT applications of control theory
>to manual control in psychology: control is looked at in terms of
>input-output transfer functions rather than as the control of perception.

Even if the models are the same as PCT models, to Hell (as Devils) with them, eh?

Some references, besides those I already mentioned in my post to Bill:

L.E. Bayliss, LIVING CONTROL SYSTEMS (catchy title!), Freeman, 1966

John H. Milsum, BIOLOGICAL CONTROL SYSTEMS ANALYSIS, McGraw-Hill, 1966

H. Kalmus, ed., REGULATION AND CONTROL IN LIVING SYSTEMS, Wiley, 1966

I expect a full report showing the utter lack of "control principles" in any of these books on my desk tomorrow morning. That will show you've done you're homework and are prepared with solid arguments, so you won't need to lapse into bald assertions again.

>The reference signals that are internal to the organism in these models
>are seen as offsets in the I-O tranfer functions.

That isn't my impression in at least some of the models in the above books. Can you guess which models? If you don't see any, I'll look more deeply in the piles in my office and come up with the book that has part of the same model used by Bill in the Little Man (yes, credit was given in the manuscript Bill and I sent to SCIENCE).

I have yet to be convinced that not making a big deal about "control of perception" equals "doing without the principles of control," as Bill Powers claimed. Nevertheless, it does appear to me that at least some biologists who have never heard of PCT DO make something of a big deal of "control of sensed signals." Still, I suppose that even a little difference bigness-of-deal and/or terminology is enough for you to declare PCT totally and revolutionarily different than what they are doing. Sigh.

>Gary Cziko 930712.2000 UTC

In English, I see! [Private joke.]

>I had the pleasure of chatting with Professor Prosser a few months ago
>about the evolution the nervous system. At no time did he mention that
>nervous systems functioned as control systems.

>I have also consulted his his book _Adaptational Biology_ where in his
>chapter on behavior he says nothing about the control of perception and
>devotes a lot of ink to "fixed action patterns," which he says completely
>accounts for the behavior of certain organisms.

>I will take a look at his _Comparative Animal Physiology_. If you can
>direct me to certain pages where he reveals knowledge of how control
>systems operate in organisms, I would much appreciate it. Since he is on
>this campus, I would be pleased to have another control theorist to talk
>to.

You misunderstood. I was not claiming that Prosser is "another control theorist" of your ilk. I mentioned his textbook simply because it covers more than just human physiology, but most any text on human physiology would do as well as C.A.P. to illustrate that, in Bill Power's terms, "the principles of control" are NOT "done without" by physiologists working in areas such as temperature regulation, control of feeding, and so forth (to which chapters of Prosser's book I would refer you; my own copy is buried somewhere at least temporarily unfindable). Consider also looking at some of the references I gave in reply to Rick, above.

As ever,

Greg

P.S. on Rick's homework assignment: I meant in "all" of the books, not in "any" of them. I'm not THAT nice.

Date: Mon Jul 12, 1993 5:47 pm PST
Subject: Hyperbole, machine recognition, CSG meeting

[From Bill Powers (930712.1910 MDT)] Greg Williams (930712) --

ME:
>>All explanations offered during the history of the life
>>sciences have managed to do without the principles of control
>>-- which means that a lot of persuasive garbage has been generated.

YOU:
>I am curious about why you persist with such extreme
>statements. They DO grab one's attention, true; and they DO
>make PCT sound absolutely revolutionary and a total break from
>all previous explanations in biology. But they also make you
>(and other PCTers, by implication) sound ignorant of the
>abundant recognition of control processes (yes, with reference
>signals internal to organisms!) in the biological literature.

>Sometimes embarrassed by your hyperbole ...

In the present (which I count as the time since World War II but which others may measure differently), more and more scientists have begun exploring control theory as it applies to various organismic subsystems. Prior to the development of control engineering in the mid-1930s, however, ALL of the life sciences developed without any way of understanding the principles of control. This means that all the great historical thinkers found and promulgated explanations of behavior in which behavior was seen not as part of a control process, but as output. A large body of logical explanation, backed up by experimental evidence and both scientific and philosophical justification, was accumulated, taught, and written down as fundamental to the life sciences. These foundational ideas are still with us. Control theory (in many hands) is slowly chipping away at these foundations, but the mainstreams of the life sciences (even when they allow for certain restricted uses of control models) still strongly resist the implications with respect to basic concepts of what organisms are and what behavior is.

I am not in the least embarrassed to say that control theory in general and PCT in particular amount to a complete upheaval in basic scientific ideas about living systems.

Martin Taylor (930712.1415) --

Loved the speech recognition example. I am coming to appreciate what the difficulties are. Clearly, no matter how good a phoneme recognition system gets to be, it can only serve as the front end of a complete recognition system.

Best to all, Bill P.

Date: Mon Jul 12, 1993 8:02 pm PST
Subject: CSGnet High and Dry

[from Gary Cziko 930713.0311 UTC]

Martin Taylor 930712 14:15 mentioned:

>PS. Our internet connection is very uncertain. I understand that the
>Mississippi floods have something to do with the problem. I am making no
>serious attempt to contribute to CSG-L until I hear that the linkage has
>been, shall we say, solidified. And I am away for a month, starting July 16.

To reassure Martin and other CSGnetters, although CSGnet is based in Illinois, we are in the east central part of the state, a three or four hour's drive from the mightly and swollen Mississippi. Although we have had a rainy summer as well, there has no been flooding hear to menace CSGnet.--Gary

Date: Mon Jul 12, 1993 8:53 pm PST
Subject: Say ya wanna revolution

[From Rick Marken (930712.2100)] Greg Williams (930712 - 2)

>Well, you'd better look for yourself.

Why? I was just asking. I'm not a biologist and I thought maybe you could tell me about the biological stuff that had been done along PCT lines (testing for controlled variables, for example). I didn't dispute the fact that there are useful applications of control theory in biology (as there are in manual control). I was just wondering if they had picked up on the control of perception aspect of control or whether, like the manual control people, they were more interested in the dynamics of transfer functions. Nothing wrong with the latter -- just not fundemental, at least from my perspective. It seems like one should know what is being controlled before measuring how well it's being done.

>be aware that it isn't common practice in biology to
>actually name, for example, input signals to biochemical sensors
> outside the nervous system as "perceptual" signals.

And well they shouldn't. But they might want to call the outputs of these biochemical sensors "perceptual" or "sensory" signals.

>At least some do make a pretty big deal about the sensed
>signals corresponding with but not being an "objective state of
>affairs." But maybe not a big enough deal for you?

The only "deal" I would like to see made is "the test for the controlled variable".

>If you have bones to pick with how much importance some accord to
>this "main point," why not pick them with the biologists themselves,
>rather than via my passing along what I see them saying?

I was curious -- not meaning to pick bones. I'm not a biologist. I thought you knew something about it because you posted about it. I mainly pick psychological bones.

>Even if the models are the same as PCT models, to Hell (as Devils)
>with them, eh?

Not really. I've tried to discuss PCT with manual tracking people who were willing to listen. They usually aren't.

>I expect a full report showing the utter lack of "control principles" in
>any of these books on my desk tomorrow morning.

Again, I didn't mean to say that there was an utter lack of control principles in these articles. I just thought they might (MIGHT) have missed some points that are central to PCT and that would motivate a different sort of research (testing for controlled variables) than the kind usually associated with non-PCT applications of control theory (which mainly involves the observation of the dynamics of disturbance-output relationships).

>That isn't my impression in at least some of the models in the above
>books. Can you guess which models?

If you could point them out it would be great. If there are articles which describe research that involes what is essentially the test for controlled variables then I would be thrilled to find them. I don't know if I would be competent to do that with a biology article -- that's why it would be nice if you could describe one where they do "the test". All I know is that I have never run across such an article in the manual control literature.

>Nevertheless, it does appear to me that at least
>some biologists who have never heard of PCT DO make something of a big
>deal of "control of sensed signals." Still, I suppose that even a
>little difference bigness-of-deal and/or terminology is enough for you
>to declare PCT totally and revolutionarily different than what they
>are doing. Sigh.

Ah, I see. You think it's rude for me to be crying "PCT revolution" when what is essentially PCT has been done in biology for years. OK. Maybe PCT has been done in biology for years -- I am not that familiar with the field. Fine. PCT is not revolutionary. It's just mainstream biological cybernetics, or whatever. Then why is the apparent response of the biological community to PCT about the same as that of the psychological community? Why is PCT ignored or rejected so strongly (is Randall Beer using PCT now)?.

The only reason we ever say that PCT is revolutionary is because people respond to it as though it is. Even people who like PCT are unwilling to accept some basic facts about how the model behaves.

I'm willing to stop saying that PCT is revolutionary. But do I also have to stop saying that people are wrong when they say that 1) perceptual information guides the outputs of a control system 2) that feedback is too slow to be involved in the control of most normal behavior 3) that most behavior occurs without feedback 4) that current IV-DV methodology is appropriate for the study of living control systems 5) that reinforcement selects behavior, 6) that control systems control their outputs, etc? Is everyone ignoring or vigorously rejecting PCT because it's NOT revolutionary; because everyone already knows PCT and only the few who don't know it say things like 1-6 above?

I don't mean to take away anything from the brilliant researchers who are applying control theory appropriately in biology, psychology, etc. But if they are also promulgating fundemental misconceptions about how control works (and virtually all manual control is based on misconception 1 above) then I think this is a disservice to the control theoretic perspective on living systems. Some of the misconceptions about control theory in the applications I know of (to manual control) are so fundemental that they are not really even applications of control theory. I would be surprised if this were less of a problem in biology -- but who knows?

Maybe you think that when PCTers say "revolution" it turns off all those serious biological scientists who would otherwise see the value of PCT because that's what they are doing anyway. Well, how about we don't say "revolution" for a month and see how many biologists start doing PCT science in that period?

Best Rick

Date: Tue Jul 13, 1993 4:40 am PST
Subject: Still more "All"?

From Greg Williams (930712) Bill Powers (930712.1910 MDT)

>In the present (which I count as the time since World War II but
>which others may measure differently), more and more scientists
>have begun exploring control theory as it applies to various
>organismic subsystems. Prior to the development of control
>engineering in the mid-1930s, however, ALL of the life sciences
>developed without any way of understanding the principles of

>control.

Thanks for clarifying what you count as history. I believe that many people's notions of "history," including my own, include at least some events which occurred more recently than 50 years ago, and thus your original unqualified statements sounded hyperbolic to me (and would have sounded so to them). The above is certainly not hyperbolic, and it clears up the ambiguity about what you count as "historical."

>I am not in the least embarrassed to say that control theory in
>general and PCT in particular amount to a complete upheaval in
>basic scientific ideas about living systems.

Nor am I. But I am still embarrassed by the potential for misunderstanding by nonPCTers in your original extreme-sounding statement about ALL biologists not using control principles. Again, without the qualification provided by explicitly pointing out your idiosyncratic idea of when history ends relative to the present, the potential is high for such statements to be treated as the ravings of a ignorant crank. By being careful to clarify, you can avoid such labels.

>Rick Marken (930712.2100)]

GW>>Well, you'd better look for yourself.

>Why? I was just asking.

So that (as I said in my previous post) you'll have a basis for conducting reasoned arguments, instead of just asserting, as is your wont.

>I'm not a biologist and I thought maybe you
>could tell me about the biological stuff that had been done along PCT
>lines (testing for controlled variables, for example).

You've got some references now. There are more where they came from. If you need some guidance, you might start with the glosses in a couple of those references of the work by Benzinger, et al., on regulation of body temperature (they did experiments to attempt to determine some controlled variables).

>I didn't dispute the
>fact that there are useful applications of control theory in biology (as
>there are in manual control). I was just wondering if they had picked
>up on the control of perception aspect of control or whether, like the
>manual control people, they were more interested in the dynamics of
>transfer functions.

I realize that. And I answered you in my previous post.

>Nothing wrong with the latter -- just not fundemental,
>at least from my perspective. It seems like one should know what is being
>controlled before measuring how well it's being done.

Exactly the sentiments of Benzinger (among others).

GW>>At least some do make a pretty big deal about the sensed
GW>>signals corresponding with but not being an "objective state of
GW>>affairs." But maybe not a big enough deal for you?

>The only "deal" I would like to see made is "the test for the controlled
>variable".

That is what some of the biological modelers apply (though they don't canonize it!), to provide a basis for constructing their control models. They are smart enough to not want to waste their time building models with untested controlled variables.

GW>>I expect a full report showing the utter lack of "control principles" in
GW>>any of these books on my desk tomorrow morning.

>Again, I didn't mean to say that there was an utter lack of control
>principles in these articles.

OK, you're off the hook. It was Bill (with my interpretation of "history") whom I THOUGHT believed this.

>I just thought they might (MIGHT) have
>missed some points that are central to PCT and that would motivate
>a different sort of research (testing for controlled variables) than
>the kind usually associated with non-PCT applications of control theory
>(which mainly involves the observation of the dynamics of disturbance-
>output relationships).

And I replied that I think at least some didn't miss the points. But you'll have to look at their work if you want to dispute my interpretation.

GW>>That isn't my impression in at least some of the models in the above
GW>>books. Can you guess which models?

>If you could point them out it would be great. If there are articles
>which describe research that involes what is essentially the test for
>controlled variables then I would be thrilled to find them.

See above.

>Ah, I see. You think it's rude for me to be crying "PCT revolution" when
>what is essentially PCT has been done in biology for years.

It doesn't so much sound rude as ignorant of nonPCT work in biology. (I quickly add that I have the highest respect for your knowledge of psychology.)

>OK. Maybe
>PCT has been done in biology for years -- I am not that familiar with the
>field. Fine. PCT is not revolutionary. It's just mainstream biological
>cybernetics, or whatever. Then why is the apparent response of the biological
>community to PCT about the same as that of the psychological community?
>Why is PCT ignored or rejected so strongly (is Randall Beer using PCT now)?.

I think it is mainly because of the self-isolating aura promulgated by some PCTers -- the aura that PCT is revolutionarily different from everything else and that everybody who says "but..." has the wrong view of what is important and what isn't. The bottom line: if the models are the same as PCTers' models, those who care mainly about modeling (and testing those models) in biology and who don't care much about meta-issues of significance of various parts of models are simply not going to NEED certain PCT ideas to do what they want to do. They already are using many of the same (control theory) principles as PCTers use in their models, and they simply don't care about the trappings. If the rejection is strong, it is probably because a whiff of cultish clinging to the extreme importance of ideas (nothing wrong with them in themselves) which they (the nonPCT control-theory-using biologists) see as either peripheral or obvious in terms of contributing to tested models.

>The only reason we ever say that PCT is revolutionary is because people
>respond to it as though it is.

Well, I doubt that is true for some PCTers. The notion of living control systems IS revolutionary with respect to pre-WWII ideas in biology. Some folks (but not ALL nonPCTers) haven't progressed beyond that era. But some nonPCTers, like Randall Beer, have. Beer certainly didn't act like PCT is revolutionary. (He DID act like some of its hypotheses are little-tested, which is certainly true.)

>I'm willing to stop saying that PCT is revolutionary.

You need to stop saying it is revolutionary with respect to at least some of current nonPCT biology.

>But do I also have to stop saying that people are wrong when they say that
>1) perceptual information guides the outputs of a control system

No. (Some nonPCT biologists do not say this.)

2) that feedback is too slow to be involved in the control of most normal behavior

No. (Some nonPCT biologists do not say this.)

3) that most behavior occurs without feedback

No. (Some nonPCT biologists do not say this.)

4) that current IV-DV methodology is appropriate for the study of living control systems

No. (Some nonPCT biologists do not say this.)

5) that reinforcement selects behavior,

No. (Some nonPCT biologists do not say this.)

6) that control systems control their outputs, etc?

No. (some nonPCT biologists do not say this.)

>Is everyone ignoring or vigorously rejecting PCT because it's NOT
>revolutionary; because everyone already knows PCT and only the few
>who don't know it say things like 1-6 above?

I think not EVERYONE, but some, yes. They don't "know PCT" as such, but they work with the same principles of control as do PCTers. And they simply don't see what the big deal is about PCT FROM THE STANDPOINT OF WHAT THEY THEMSELVES WANT TO ACCOMPLISH.

>I don't mean to take away anything from the brilliant researchers who
>are applying control theory appropriately in biology, psychology, etc.
>But if they are also promulgating fundemental misconceptions about how
>control works (and virtually all manual control is based on misconception
>1 above) then I think this is a disservice to the control theoretic
>perspective on living systems. Some of the misconceptions about control
>theory in the applications I know of (to manual control) are so fundemental
>that they are not really even applications of control theory. I would
>be surprised if this were less of a problem in biology -- but who knows?

You would know, if you looked at the literature.

>Maybe you think that when PCTers say "revolution" it turns off all those
>serious biological scientists who would otherwise see the value of PCT
>because that's what they are doing anyway. Well, how about we don't
>say "revolution" for a month and see how many biologists start doing
>PCT science in that period?

Several are already modeling living control systems without pledging to make a big deal out of what you think they should. Maybe instead of stopping saying "revolution," you should investigate how THEIR revolutionary (w.r.t. pre-WWII biology) work meshes with PCT notions.

As ever, Greg

Date: Tue Jul 13, 1993 7:28 am PST
Subject: PCT revolution

[From Bill Powers (930713.0700 MDT)] Greg Williams (930712) --

Greg, I think you're missing my meaning in one respect. When I say that control theory is not being used or used properly, I always specify the behavioral -- BEHAVIORAL -- sciences. In that category I do not include studies of regulation of respiration, body temperature, etc.. I have said throughout (for years) that others have successfully applied control theory to certain subsystems in the organism. There have been researchers doing this from the very beginning. I am conscious of, if not very familiar with, such efforts, and by specifying that I am talking about the behavioral sciences, try to take care that my blanket statements are not unfair to them.

However, look at the dates on the three books you mentioned yesterday: 1966! You also could have mentioned Ashby, MacFarland, Jones, Toates and probably others. Control theory reached its peak of popularity before 1966, and then was largely abandoned, having failed to develop into anything of general interest. Most of the books on control theory and organisms have been about isolated subsystems; most of them go over the same old ground again and again, using slightly different examples but all covering the basic literature on Laplace transforms, frequency domain, stability, and a few other topics. The basic model remains the engineering model that the engineering psychologists adopted in the late 1940s and early 1950s. That model is suitable for single control systems, although it leaves a wrong impression about what the whole system is for and how it relates to its surroundings, and it doesn't account for the sources of reference signals. But successful modeling has been done and I've never said it hasn't.

However, with respect to the BEHAVIORAL sciences, I don't think you can show me any examples (outside those informed by PCT) of the correct approach to control processes. By the behavioral sciences I mean those branches of any science concerned with the interaction of whole organisms and their environments. Biology, as a behavioral science, still uses fundamentally a cause-effect approach. Probably the closest independent approach to a PCT-like model is in sports psychology or sports physiology, which recognize the influence of mental images, psyching-up, and so forth, but without any formal model behind these empirical observations. PCT-like ideas seem to be most prevalent in fringe aspects of the behavioral sciences where the phenomena are recognized but without explanation.

For biology and biochemistry, the greatest stumbling block as far as I can see is still the concept of a reference level. For some reason never explicitly stated, biologists seem to think that this idea lets metaphysics in by the back door, or something. There seems to be a principle involved, something about biological organisms just behaving as their physical construction ends up making them behave -- as one obesity researcher said, "Settling points, not set points." I think there's a glimmer of recognition that control systems bring a spark of directing intelligence into biological processes. Despite the teleological language that biologists use in describing the behavior of whole organisms, any whiff of real teleology is instantly rejected, sort of like an anti-pheromone. You can say that wolf urinates on a tree to mark its territory, but not that it WANTS to mark its territory. One result of this general attitude is that even while biochemists investigate feedback loops in biochemical systems, they insist that there is no reference signal for such systems. I know of only one exception, in Hayashi and Sakamoto. I have written to several biochemists on this subject and have received explicit denials that reference levels play any part in such negative feedback loops. There is obviously some religious reason for excluding them.

The 1950s engineering model allows modelers to use reference signals without recognizing their true role. This is because the engineering model calls the reference signal an "input," implying that it is part of a normal causal input-output process. In many models, this input is labelled something like "required state." This allows the reference signal to be used in a model without ever answering the question of what "requires" that state, and for what purpose, and how. The result is that individual control processes are left isolated from each other and from the operations of higher systems (in biology the very concept of "higher" systems is, by and large, rejected because causation is generally conceived of as bottom-up).

This, in turn, means that the concept of hierarchical control can't develop: the idea of superordinate control systems that act by adjusting reference signals for lower systems. Without that concept, the isolated control systems can't form a whole system. There can be nothing to coordinate them. The behavior of the whole simply arises from what all the parts, individually, are doing. Technically this is true, but all the parts are not at the same level. I think the basic biological idea is that they ARE all at the same level.

So even where control theory is being used to model behavior, I think its real message is lost. The concept of a control system as determining what the environment is to do to it is simply missing. This is the central revolutionary concept of PCT. It changes the whole traditional mind-set toward behavior. Instead of showing behavior as a reaction to external influences, PCT shows it as as active process with the organism, not the environment, at the center.

This is really the concept that meets with the most direct resistance from conventional disciplines. 350 years of science has taught us the exact opposite. You have called this view, in the past, an extreme kind of organocentrism. It is. You have called for a compromise between extreme organicentrism and extreme envirocentrism. I do not think this compromise is feasible, or even if feasible, desirable. It would be like seeking a compromise between the phlogiston explanation of combustion and the oxygen explanation. If all behavior is part of a closed-loop process, there is no way to compromise by saying that sometimes it is simply a reaction to stimuli. Even where specific reactions to stimuli can be shown, they remain unexplained (under PCT) until the organism's reason for setting itself up this way is found -- and that reason, under PCT, will always be in order to produce some specific effect ON THE ORGANISM and not on the environment.

Best, Bill P.

Date: Tue Jul 13, 1993 9:09 am PST
Subject: Higher level investigation

[From Tom Hancock (930713.0735) Gary Cziko (930710.0400)

>An R square of .57 is only 66% useless.

I appreciate that an R square of .57 is partially useless. (For precision sake, isn't it a .43 factor of uselessness rather than .66 as you said. I believe that the coeffecient of alienation does not demand a square root as you indicated-what would be the point?)

But Gary, isn't it also partially useful-57%-which is much better than no evidence? No, it is not a precisely fit model. But it is a modicum of evidence of higher level PCT-based predictions, in an area where not much exists?

On the other hand, I appreciate that I should not report these results as a CSG sanctioned study, so that the high PCT standards should not be lowered.

>But I'm not sure what you mean by a trend here.

The trend for the individual subjects was as follows: Every subject (16) in the third of the subjects who had the highest achievement showed significant effects (alpha = .0001) for the PCT generated prediction.

P.S. Was my mail working OK? I received 2 messages over the
net on the weekend.

Dag (930709 1255),

I am encouraged to use more PCT demos in my classes--thanks. I got a copy of the IV-DV discussion by Bill.

Tom Hancock

Date: Tue Jul 13, 1993 9:57 am PST
Subject: Help

[From Tom Hancock (930713.1000)

Rick,
I understood that you are interested in helping me with PCT (930709.0800). I am interested in a more tightly conceived testing of a person's control at higher levels. Would you be willing to check out my recent logic in testing a subject's control while performing a computer-based drill task? If so, should I just send it to you directly, to save net space?

Tom

Date: Tue Jul 13, 1993 10:29 am PST
Subject: doing one's du di

[From: Bruce Nevin (Tue 930713 12:32:13 EDT)] Tom Bourbon [930702.1655]

> SAYING /di/ and /du/

I was cut off for a week because the listserver was sending mail to me as bnevin@ccb.bbn.com instead of as bn@bbn.com, and ccb got taken away. Am just catching up.

Thank you, Tom and Andy, for this (as usual) very clear presentation.

Liberman et al at Haskins do interpret their findings as validating the invocation of plans. However, that is because that is all they know about, or specifically because they don't know about PCT. They say (in a paper that I recently cited) that the elements of phonology are the intended gestures of speakers. For speech production, this is straightforward. For speech perception, the hearer must perceive in imagination what gestures might have been intended to produce the actually perceived acoustic outcomes. They interpret this in terms of plans.

I have argued (apparently unsuccessfully) that there is a simpler and more appealing interpretation in terms of controlling perceptions of what they identify as gestures.

I have argued (apparently unsuccessfully) that the gestures and their acoustic correlates are defined by their systematically mutual contrastiveness.

I have argued that this accounts for how we so readily accommodate variability, though the point was probably unintelligible without the first two being accepted, or perhaps merely unmotivated because others here are not yet ready to consider the problems of variation and diversity in language, which are perhaps analogous to recognizing the "same" face as an invariant.

The problem of "context-conditioned variation" is not really the same problem. In this case, the occasion for variation is always identifiable within the same utterance, and can be understood as the speaker's control of perception of one part of the utterance interfering with control of the perception of another part of the same utterance. Consider your example of /di/ and /du/. Your drawings follow:

===

Figure 1. ASCII drawings of idealized spectrograms for /di/ and /du/ follow:

```
                  /di/                /du/

          |   **************
     high |  *
          |_*                    _
          |                       *
          |                        *
frequency |                         **************
          |
          |   **************        **************
          |  *                     *
     low  | *                     *
          |__________________    __________________
            0             300     0             300
                        time in msec
```

==

What can be taken as invariant for the two occurrences of /d/ is a projected point at which the transition to formant 2 of the vowel begins. I have indicated this point by inserting the character _ into your diagram. The transition reflects the movement of the tongue from its position closing the oral cavity for /d/ to its position shaping the vowel formants for either /i/ or /u/, respectively. As you say:

>It appears
>to many people as if, depending on what steady-state formants the
>articulators *are going* to produce, they [the articulators]
> modify the transitions
>that they *do* produce, and the acoustic signal that results in
>either case is perceived as /d/.

But this is not a result of planning so much as a consequence of physics. The tongue in moving from closure at the alveolar ridge to the configuration that raises F2 for /i/ (or lowers it for /u/) must pass through intervening space in the oral cavity, and the transition is the acoustic consequence. The projected place whence the transition originated is the hearer's cue as to the location of the tongue blade prior to the vowel. Because it is the hearer's only cue, it presumably is among the perceptions that the speaker controls also, that is, among the perceptual inputs that may satisfy the input function for perceiving /d/. The hearer must be able reliably to divine the speaker's intended pronunciation of words or there is no language. (Problems of the hearer immediately forgetting the exact words, instead substituting her own words out of her process of controlling meanings, should not obscure this essential truth: it has to start someplace, she didn't pick those meanings entirely out of thin air.)

Now we come to why I say that contrast is fundamental. As you say:

> (1) There are no known acoustic invariances that our model can be
> made to control for.
>
> (2) There are articulatory invariances - i.e. particular
> sequences of movements of articulators. These are unacceptable,
> or appear to be, for reference signals. One reason they are
> unacceptable is because we may refuse to believe that movement
> invariances exist at all in spite of the evidence.

The articulatory "targets" are invariant not as absolute movements, but as movements made so as to be distinct from one another.

I will stop there. If that is not clear, more words won't help now. And I have to stop for familiar other reasons.

Bruce bn@bbn.com

Date: Tue Jul 13, 1993 10:47 am PST
Subject: PCT revolution

[From Rick Marken (930713.1000)] Greg Williams (930712) --

>So that (as I said in my previous post) you'll have a basis for
>conducting reasoned arguments, instead of just asserting, as is your wont.

Well, one man's assertion is another man's reasoned argument. Could you give me an example of where I "just asserted" something as a way of arguing for my point of view? The only assertion I recall making in this discussion is that I have found no evidence, in applications of non-PCT control theory in psychology, of any effort to determine what an organism might be trying to perceive. This assertion is based on a lot of experience with applications of control models to behavior. I suggested (I didn't assert) that this lack of interest in determining controlled perceptual varables might also exist in applications of non-PCT control theory in biology. You say that they do use "the test" in biology. That's good enough for me, but I would like it if you could describe an example, not because I don't believe you, but because I am really INTERESTED. I would have a tough time getting hold of the references you suggest -- we deal with unmanned space systems here so we don't have a lot of biology books and journals in the library.

>The bottom line: if the
>models are the same as PCTers' models, those who care mainly about
>modeling (and testing those models) in biology and who don't care much
>about meta-issues of significance of various parts of models are
>simply not going to NEED certain PCT ideas to do what they want to do.

I guess the question then is "what's a meta-issue"? If the biological modellers don't need PCT ideas then, of course, they won't be interested in learning them. The behavioral science control systems modellers don't seem to need them either. Since I am familiar with the behavioral science issues and models, it is my opinion that these modellers are making a BIG mistake by ignoring PCT "meta issues" like control of perception, hierarchical control, etc (if those are the meta issues). But maybe it's not a mistake for the biological modellers to ignore these issues. I don't know.

> If the rejection is strong, it is probably because a whiff
>of cultish clinging to the extreme importance of ideas (nothing wrong
>with them in themselves) which they (the nonPCT control-theory-using
>biologists) see as either peripheral or obvious in terms of
>contributing to tested models.

Again, my only experience with rejection of PCT by non-PCT-control theory-using people is with psychogists. These psychologists may be rejecting PCT becuase it is too cultish for them but, from my point of view, they seem to be rejecting it because they don't understand the proper application of control theory to living control systems.
>You need to stop saying it is revolutionary with respect to at least
>some of current nonPCT biology.

>they work with the same principles of control as do PCTers. And
>they simply don't see what the big deal is about PCT FROM THE
>STANDPOINT OF WHAT THEY THEMSELVES WANT TO ACCOMPLISH.

The manual control theorists in psychology are using the same principles of control as PCTers. They are just not applying these principles correctly. They don't know this -- nor do they care. They measure all kinds of fancy mathematical relationships between disturbance and output and assume that these relationships tell them something about the nature of the controller. They are basically wrong -- for reasons spelled out in painful detail on this net and in many PCT books and articles.

You said in an earlier part of the post that I make assertions instead of engaging in reasoned discussion. It seems to me that you keep asserting that many biological control modellers are contributing to our understanding of living systems even though they don't use (and don't want to use) PCT. Well, how about backing up this assertion. How about describing the work of biological modellers that helps us understand one or another of the aspects of control (the "meta issues") that are of concern to PCTers, eg. what variable(s) are controlled, how do you determine what variable is cotnrolled, what determines the level at which variables are controlled, what is the relationship (hierarchical, heterarchical, random) between controlled variables, how do you determine the relationship between controlled variables, how does the control organization develop, how does it change, what are the irrelevant side effects of control, etc. I don't doubt that such work exists; it would just help if you could describe it instead of posting what, for me, are difficult to get references.

Best Rick

Date: Tue Jul 13, 1993 1:21 pm PST
Subject: "All"? or Nothing?

From Greg Williams (930713 - 2) Bill Powers (930713.0700 MDT)

>Greg, I think you're missing my meaning in one respect. When I
>say that control theory is not being used or used properly, I
>always specify the behavioral -- BEHAVIORAL -- sciences.

Indeed. I stand corrected regarding both your meaning of "history" and now your meaning of "behavior," which you do not count as including such (oftentimes referred to as "vegetative") functions as temperature regulation, respiration, etc. Just as (I believe) many persons (including myself) count "history" as extending closer to the present than you do, I believe that many persons (including myself) count, for example, breathing, as behavior. I hope you can excuse my indiscretion in wrongly interpreting your concepts of "history" and "behavior." I can assure you that I won't make the same mistake again.

>However, look at the dates on the three books you mentioned yesterday: 1966!

And look at what I mentioned in my post before that: BIOLOGICAL CYBERNETICS, which continues to come out year after year. So does PHYSIOLOGICAL REVIEWS. Just because you don't see personally much work in a field doesn't mean it has been "largely abandoned."

>However, with respect to the BEHAVIORAL sciences, I don't think
>you can show me any examples (outside those informed by PCT) of
>the correct approach to control processes.

Well, I'm not going to photocopy a bunch of articles out of BIOLOGICAL CYBERNETICS (where there are many control-based models published of BOTH vegetative and non-vegetative functions) and send them to you -- get them yourself. After you've at least looked through several recent volumes, you can join Rick (?) in making evidence-based arguments about what you both currently only assert without (it appears to me) sufficient evidence one way or the other. If you contend that ALL of the contributors to BC miss the PCT point, you should be prepared to back that contention up with evidence. You might be right -- but I'm not about to be convinced by poorly supported hypotheses. Show me. (On the other hand, if you really mean "all the nonPCT work you've ACTUALLY SEEN," rather than truly "all," then I stand corrected once again -- but I believe that you must admit you've seen only a small sample of what's in recent BCs and PRs.)

As ever, Greg

Date: Tue Jul 13, 1993 2:07 pm PST
Subject: Help, du di

[From Rick Marken (930713.1300)] Tom Hancock (930713.1000) --

>Would you be
>willing to check out my recent logic in testing a subject's
>control while performing a computer-based drill task? If so,
>shou d I just send it to you directly, to save net space?

Sure. But why not post it to the net? I don't know if I can help you but there may be someone else on the net who can -- which would help us both.

Bruce Nevin (Tue 930713 12:32:13 EDT) --

>I have argued (apparently unsuccessfully) that the gestures and their
>acoustic correlates are defined by their systematically mutual
>contrastiveness.

I don't understand what this means. I guess I wasn't paying real close attention to the early go-round on this. Is there a way to define "systematically mutual contrastiveness" so that it can be used in a model of speech production/recognition?

>I have argued that this accounts for how we so readily accommodate
>variability

Does this mean that it accounts for the ability to recognize /d/ despite the variability of the acoustic correlates of /d/? What's wrong with recognizing it on the basis of what you mention below:

>a projected point at which the transition to formant 2 of the
> vowel begins.

So here is an acoustic invariant that can serve as the basis for recognizing /d/ in /di/ and /du/. I don't see why the recognition system would have to go through the process of generating the imagined articulations that produce this invariant -- not that they don't, I just don't see why it would be necessary.

By the way, if articulation is so important to recognition, why is it that people can recognize speeded speech (when properly compen- sated for frequency shift) which they could not possibly articulate?

Subjectively, I am aware of imagined articulation that occurs when I listen to some -- but not to all -- speech. I don't doubt that there are imagined articulations; I just don't believe that they are funde- mental to the process of recognizing speech (real articulations are, of course, fundemental to the process of producing it).

>Now we come to why I say that contrast is fundamental.

>The articulatory "targets" are invariant not as absolute movements, but
>as movements made so as to be distinct from one another.

So an articulatory target is a reference state of a controlled perception? A reference state that is defined as "distinct from other reference states for the same perception"? Why not just say that an articulatory target is defined by the value of a reference signal, r, that specifies the value of a perceptual variable, p, that represents the state of some articulator. The distinctness of articulatory targets is then taken care of by the fact that r can only have one value at a time so that articulatory targets are as distinct from one another as real numbers can be.

Best Rick

Date: Tue Jul 13, 1993 4:39 pm PST
From: Control Systems GrouSubject: Whaddya want?

From Greg Williams (930713 - 3) Rick Marken (930713.1000)

GW>>So that (as I said in my previous post) you'll have a basis for
GW>>conducting reasoned arguments, instead of just asserting, as is your wont.

>Well, one man's assertion is another man's reasoned argument.

Dream on. I'm talking about arguing from evidence, rather than sheer imagination.

>Could you give me an example of where I "just asserted" something as a way
>of arguing for my point of view?

Certainly. In recent posts regarding "environmental causation." You even ASSERTED -- with no supporting argument (and not very convincingly) -- at the last that I had convinced you to see it my way.

>I guess the question then is "what's a meta-issue"?

Issues like the relative "importance" of reference signals and environmental disturbances in behaviors (and other organismic functions, Bill). Like the "importance" of whatever you think nonPCT control modelers don't "get," despite having the same kinds of models as PCTers. Like what should be "The" definition of "control."

>If the biological modellers don't need PCT ideas then, of course,
>they won't be interested in learning them.

Some of them already use many of the same ideas PCTers use.

>The manual control theorists in psychology are using the same principles of
>control as PCTers. They are just not applying these principles correctly.

OK, what counts as applying the principles of control "correctly"? I'm at a loss to understand what could be incorrect about models which are the same as those used by some PCTers, so I'm REALLY interested in why you think the manual control theorists who use such models are applying the principles of control "incorrectly." For example, in one of those hard-to-find references I mentioned in an earlier post, John Milsum (an engineer at McGill -- maybe he doesn't count as a manual control theorist, but I'll risk your wheedling out of it) shows what I would consider to be a perfectly PCT-orthodox diagram of a model for speed control by an auto driver. He shows the "desired speed" as a reference signal INSIDE the driver's nervous system, coming into a comparator INSIDE the nervous system, and a perceptual signal INSIDE the nervous system derived from the car's speed getting subtracted from the reference signal at the comparator. He even (on the next page) says "we will in future reserve the term CONTROL SYSTEM (or CONTROL) for the active type alone, namely, incorporating a purposeful closed-loop action with negative feedback." Where has he gone wrong? Later in his book, he talks about how to determine controlled variables and describes interrelations among controlled variables in certain living control systems.

>You said in an earlier part of the post that I make assertions
>instead of engaging in reasoned discussion. It seems to me that
>you keep asserting that many biological control modellers are
>contributing to our understanding of living systems even though they
>don't use (and don't want to use) PCT.

No, I'm saying that they use the same control principles as are used by PCTers, and that they apparently don't see anything else in PCT (I haven't addressed differences from HPCT, which are definitely there), as necessary to THEIR PERCEPTIONS of "contributing to our understanding of living systems." So tell me what YOU count as "understanding of living systems." I won't assert that ALL PCTers could learn from some nonPCT models, but I'll say that I think some PCTers could.

>How about describing the work of biological modellers that
>helps us understand one or another of the aspects of control (the
>"meta issues") that are of concern to PCTers, eg. what variable(s) are
>controlled, how do you determine what variable is cotnrolled, what
>determines the level at which variables are controlled, what is the
>relationship (hierarchical, heterarchical, random) between controlled
>variables, how do you determine the relationship between controlled
>variables, how does the control organization develop, how does it
>change, what are the irrelevant side effects of control, etc.

I already described Milsum using the principles of control. Tell me how he doesn't meet your desiderata for REAL/CORRECT/PRIMO/WHATEVER living control system modelers, and we'll go from there. I'm tired of trying to hit a moving target. Tell me EXPLICITLY what you want to see (and I DON'T mean "nobody doing the right thing except PCTers"!), and I'll try to satisfy. Unless you want them to say, "Bill Powers says..."!!

As ever, Greg

Date: Tue Jul 13, 1993 8:58 pm PST
Subject: Greg's comments

[From Bill Powers (930713.2000 MDT)] Greg Williams (930713)

RE: meta-issues

>Issues like the relative "importance" of reference signals and
>environmental disturbances in behaviors (and other organismic
>functions, Bill). Like the "importance" of whatever you think
>nonPCT control modelers don't "get," despite having the same
>kinds of models as PCTers. Like what should be "The" definition
>of "control."

Importance: Reference signals specify inputs, disturbances result in outputs opposing the tendency of disturbances to alter the input from the specified level. I have remarked that behavioral outputs are relatively unimportant as objects of study (if this is what you're referring to), because all they do is match up against disturbances, and disturbances can arrive in any old random sizes and directions. Studying outputs thus amounts mainly to studying disturbances. Something wrong with saying that? Of course in really analyzing any system you don't ignore any part of the loop. But neither do you focus on one part of it and ignore the rest.

"Control modelers .. having the same kinds of models as PCTers."

Some do, most don't. Of course control theory is control theory, but there are many ways to use it. PCT specifically proposes that the variables controlled by human behavioral systems are perceptual signals, with the external counterparts of these signals not necessarily having any unitary existence (the taste of lemonade). Few models I have seen in the literature acknowledge the fact that it is the perceptual signal, not the external variable, that is controlled. Furthermore, engineers, who are interested in the external counterpart of the perception, often design bizzarre functions into the feedback part of the loop because that is a convenient place to insert compensations of various sorts. As long as THEY can see what the system is controlling, they don't worry about the perception resembling the external controlled variable. But PCT operates under an additional constraint. Not only must the behavior of the observed variables be modeled correctly, but the subjective experience of controlling must also be accounted for. This allows for much less freedom in selecting input functions.

"... what should be "the" definition of control."

It's not so much "the" definition of control, but "the most general" definition of control for which I argue. Other definitions of control -- meaning mainly open-loop definitions -- refer to systems that can achieve control only under special conditions. All components of an open-loop "control" system must retain accurate calibrations, and of course the controlled outcome must be protected from independent disturbances. While such conditions can be generated artificially, they are not the natural conditions under which living control systems must be able to work. And organisms must be designed to work under worst-case conditions, which can arise in an instant.

To me, the most telling criticism of other usages of the word control is that, in my limited experience, the people who employ those other meanings are actually thinking of true control, but are not aware of everything that is required to produce the control processes they see. I once had an argument with a philosopher who refused to believe that when he drove a car around a curve, he did not simply set the steering wheel to a fixed position and then restore it to normal after the end of the curve. More commonly, when people design open-loop systems they overlook their own part in the operation of the system: watching the outcome and adjusting the system so it responds correctly to its inputs. No open-loop control system can actually be made to operate correctly unless it's embedded in a closed-loop system that can perceive the actual operation and knows what the correct operation is.

I'm not saying that there is a best or only meaning for "control." I'm arguing in favor of reserving that term to mean closed-loop control, because the word already has that meaning to most people (even when they don't realize it), and inventing new words is an inferior solution.

>Some of them [biological modelers] already use many of the same ideas
>PCTers use.

The term "already" is a little misleading. That sounds as though these biological modelers arrived at their ideas, and then PCT came along and imitated them. Is that how you see the sequence of events?

>OK, what counts as applying the principles of control
>"correctly"? I'm at a loss to understand what could be
>incorrect about models which are the same as those used by some PCTers...

Incorrect uses of the principles of control include (1) putting the error signal or the comparator or the reference signal or all three in the environment outside the living control system (the math still works but the model is wrong), (2) designing a control model that controls its output, which is then supposed to lead to the desired remote effect (which can work only in a constant environment), (3) setting up a control model in which the reference signal gets into the organism via a sensory receptor (making an S-R system out of the organism), (4) setting up a control model in which the organism part is simply an input-output function with no reference signal at all (very common in manual control models), (5) modeling a low-level neurological control system as if it were a computer program (the TOTE unit), (6) analyzing a control process as if each component in the loop operated in sequence, being quiescent while the other components operate, and (7) analyzing a control process as a limit cycle, which is actually a symptom of poor design. To this I can add all the mistakes made by those who don't even think quantitatively, such as Carver and Scheier who assume that feedback exists only when a person is being self-aware, or Bandura and Locke who describe control processes in terms of goals, discrepancies, and error correction, and then claim that a control system can't exhibit this sort of behavior.

>... so I'm REALLY interested in why you think the manual control theorists
>who use such models are applying the principles of control "incorrectly."

Rick can answer this better than I. Many of the above mistakes have been made by prominent manual control theorists.

> For example... John Milsum (an engineer at McGill -- maybe he
>doesn't count as a manual control theorist, but I'll risk your
>wheedling out of it) shows what I would consider to be a
>perfectly PCT-orthodox diagram of a model for speed control by
>an auto driver. He shows the "desired speed" as a reference
>signal INSIDE the driver's nervous system, coming into a
>comparator INSIDE the nervous system, and a perceptual signal
>INSIDE the nervous system derived from the car's speed getting
>subtracted from the reference signal at the comparator. He even
>(on the next page) says "we will in future reserve the term
>CONTROL SYSTEM (or CONTROL) for the active type alone, namely,
>incorporating a purposeful closed-loop action with negative
>feedback." Where has he gone wrong?

Well, I'm glad he agreed with the usage of "control" that I had adopted ten years previously to 1966 and with the model Clark, McFarland, and I had published 6 years previously (reprinted in General Systems). How did he account for the setting of the reference signal? I'd say John was on the right track; if we had known about each other then, he might well be a member of our group now. Whatever happened to Milsun's ideas? Did he realize he was talking about control of perceptions? Did he see what this meant about behavior in general? Actually, he might well have had all these insights, and met with the same reception as PCT.

>Later in his book, he talks about how to determine controlled
>variables and describes interrelations among controlled
>variables in certain living control systems.

I suppose he presented these ideas as if nobody had ever published them before. But then I didn't read his stuff either.

Greg, you can find a piece of PCT here and another piece there, just as you can find references in the literature since the 19th Century or before claiming that behavior is goal-directed and purposive. But where can you find all the pieces in one place except in the literature of PCT? I have looked into countless papers on control theory and behavior, and have yet to see any of the significant insights of PCT even recognized. I have seen nothing on the subject that couldn't have been said in the 1950s, except perhaps for some mathematical techniques that weren't known then. Maybe you've seen things that

I haven't; that's very likely, considering the relative merits of libraries to which we have access. But as you have withdrawn support in that regard, that's just the way it will have to be.

Best, Bill P.

Date: Tue Jul 13, 1993 9:19 pm PST
Subject: Milsum control theory

[From Rick Marken (930713.2200)] Greg Williams (930713 - 3)

I said:

>The manual control theorists in psychology are using the same principles of
>control as PCTers. They are just not applying these principles correctly.

You say:

>OK, what counts as applying the principles of control "correctly"?

An excellent question, though I would add "to living systems" to the end of it. That is what PCT is really about. Control theory was around for quite some time before PCT came along. Powers did not invent (or improve) control theory; he realized how control theory applied to behaving systems -- no mean achievement.

>I'm at a loss to understand what could be incorrect about models which are
>the same as those used by some PCTers, so I'm REALLY interested in
>why you think the manual control theorists who use such models are
>applying the principles of control "incorrectly."

Well, then it's about time you found out. This is really the heart of PCT. PCT is not just control theory; it's how it's applied. It makes a BIG difference whether you apply control theory to living systems haphazardly (as it's currently done) or carefully (as it's done in PCT). The same tool can yield good or bad results depending on how it's used.

>For example, in one
>of those hard-to-find references I mentioned in an earlier post, John
>Milsum (an engineer at McGill -- maybe he doesn't count as a manual
>control theorist, but I'll risk your wheedling out of it) shows what I
>would consider to be a perfectly PCT-orthodox diagram of a model for
>speed control by an auto driver. He shows the "desired speed" as a
>reference signal INSIDE the driver's nervous system, coming into a
>comparator INSIDE the nervous system, and a perceptual signal INSIDE
>the nervous system derived from the car's speed getting subtracted
>from the reference signal at the comparator. He even (on the next
>page) says "we will in future reserve the term CONTROL SYSTEM (or
>CONTROL) for the active type alone, namely, incorporating a purposeful
>closed-loop action with negative feedback." Where has he gone wrong?

It sounds good to me! Sounds like Milsum drew the right diagram. That is impressive; I have never seen the control model diagram applied correctly to a simple manual control task. The functional equivalent of the reference signal (usually called the "command" signal) always comes in via the sensory input. The comparator is also sitting in the space outside the controller and if it's a compensatory tracking task the signal coming into the controller (via the eyes) is the error signal. But it sounds like Milsum got the diagram right; kudos. Is he still around? Maybe he would be interested in PCT.

>Later in his book, he talks about how to determine controlled
>variables and describes interrelations among controlled variables in
>certain living control systems.

Excellent! Does he describe any research aimed at determining controlled variables? I have never seen this book; it definitely seems like an example (based on your description) of a correct application of control theory to behavior. I have never seen a manual control book or article (and I have seen quite a few) that got the basic control diagram correct OR that talked about determining what variable was being controlled. Milsum seems to have gotten the application of control theory right. I'll believe it when I see the book (maybe at the meeting) but if the book says what you say it says then it is definitely an example of the appropriate application of control theory -- without the benefit of PCT.

> I'm tired of
>trying to hit a moving target. Tell me EXPLICITLY what you want to see
>(and I DON'T mean "nobody doing the right thing except PCTers"!), and
>I'll try to satisfy. Unless you want them to say, "Bill Powers says..."!!

I want to see just what you say Milsum did -- give a diagram which labels the variables in a behavioral control loop properly; give a description of the test for controlled perceptual variables; describe some research showing the application of the test. Milsum seems to have provided it. I ain't movin'; ya got me. Milsum is a non-PCT control theorist who seems to have it right. Bravo. The people who many consider leaders in this field (of manual control) have obviously not read Milsum. I'm talking about Sheridan and Ferrell, Wickins, Poulton, McRuer, Jex, Jagacinski, Pew, Schmidt, etc etc. These people have all got the control loop variables incorrectly identified; they have all done studies of "input - output" transfer functions that are really disturbance-output transfer functions that tell you nothing about the actual system transfer function. So can that be the end of the argument? Milsum is a non-PCT control theorist who gives the correct application of control theory to behavior. Even if he is the only non-PCT control theorist in the world who got it right at least he proves that there is one -- and I recognize his achievement. You have provided an example of what I thought there were no examples of -- a non-PCT control theorist who knows how to apply control theory to behavior. Congratulations. Now, can I get a free copy of the book?

Thanks Rick

Date: Wed Jul 14, 1993 3:52 am PST
Subject: PCT correctness

From Greg Williams (930714) Bill Powers (930713.2000 MDT)

GW>>"... what should be "the" definition of control."

>It's not so much "the" definition of control, but "the most
>general" definition of control for which I argue. Other
>definitions of control -- meaning mainly open-loop definitions --
>refer to systems that can achieve control only under special
>conditions.

I wasn't thinking about open-loop definitions; rather, "the" PCT definition of a "controlled variable" (as that which is maintained in the face of disturbances -- but really, more specific than that, since it must be the perceptual signal representing something outside the organism which an observer sees as maintained in the face of disturbances).

GW>>Some of them [biological modelers] already use many of the same
GW>>ideas PCTers use.

>The term "already" is a little misleading. That sounds as though
>these biological modelers arrived at their ideas, and then PCT
>came along and imitated them. Is that how you see the sequence of events?

Is "independently" better? I was speaking of folks with PCT-correct models who (apparently) never heard of your work or work by other PCTers. It would seem that the notion of using control ideas in living organism models was adopted in parallel by several persons, following Wiener et al.'s seminal paper.

>Incorrect uses of the principles of control include (1) putting
>the error signal or the comparator or the reference signal or all
>three in the environment outside the living control system (the
>math still works but the model is wrong), (2) designing a control
>model that controls its output, which is then supposed to lead to
>the desired remote effect (which can work only in a constant
>environment), (3) setting up a control model in which the
>reference signal gets into the organism via a sensory receptor
>(making an S-R system out of the organism), (4) setting up a
>control model in which the organism part is simply an input-
>output function with no reference signal at all (very common in
>manual control models), (5) modeling a low-level neurological
>control system as if it were a computer program (the TOTE unit),
>(6) analyzing a control process as if each component in the loop
>operated in sequence, being quiescent while the other components
>operate, and (7) analyzing a control process as a limit cycle,
>which is actually a symptom of poor design. To this I can add all
>the mistakes made by those who don't even think quantitatively,
>such as Carver and Scheier who assume that feedback exists only
>when a person is being self-aware, or Bandura and Locke who
>describe control processes in terms of goals, discrepancies, and
>error correction, and then claim that a control system can't
>exhibit this sort of behavior.

Thank you very much for the list. It appears to me that Milsum's example which I posted to Rick meets all of these requirements (although his terminology is different from that used in PCT, which is to be expected when two applications of the same principles develop in parallel). Twice bitten, thrice shy, I must ask explicitly whether you will count the Milsum example as a correct application of control principles in the behavioral sciences, which is what you said you didn't think I could supply. Because he is an engineer, maybe you think it doesn't count?

>Well, I'm glad he agreed with the usage of "control" that I had
>adopted ten years previously to 1966 and with the model Clark,
>McFarland, and I had published 6 years previously (reprinted in
>General Systems). How did he account for the setting of the
>reference signal?

I quote (p. 13): "Active closed-loop control of car speed arises in the general case where the driver decides that he will maintain a certain desired speed V*; this becomes the reference input to the system..." He doesn't go into how the decision happens to be made -- that would require something like HPCT, which he doesn't get into.

>I'd say John was on the right track; if we had known about each other then,
>he might well be a member of our group now.

Perhaps there are more kindred spirits (even now) than you have realized.

>Whatever happened to Milsun's ideas?

I don't know.

>Did he realize he was talking about control of perceptions?

You'd have to ask him.

>Did he see what this meant about behavior in general?

Ditto.

>Actually, he might well have had all these insights, and met with the
>same reception as PCT.

I believe that his book was a widely adopted text in "bioengineering." I remember seeing it being used at MIT for at least one course. I bought my copy at a used bookstore in Urbana, IL.

GW>>Later in his book, he talks about how to determine controlled
GW>>variables and describes interrelations among controlled
GW>>variables in certain living control systems.

>I suppose he presented these ideas as if nobody had ever published them before.

Well, no. In talking about this stuff, he references a lot of people who made control models of living systems. True, he didn't reference you. But it might be unreasonable to expect a control engineer to read PERCEPTUAL AND MOTOR CONTROL or GENERAL SYSTEMS YEARBOOK.

>But then I didn't read his stuff either.

Just as it might be unreasonable to expect you to read ONR documents and the various engineering journals referenced in Milsum's bibliography.

>Greg, you can find a piece of PCT here and another piece there,
>just as you can find references in the literature since the 19th
>Century or before claiming that behavior is goal-directed and
>purposive. But where can you find all the pieces in one place
>except in the literature of PCT?

NOWHERE else (at least in my fairly wide experience) can you find such comprehensive and fully thought-out applications of control principles to living systems. My argument here has been to show that at least some nonPCTers have applied such principles "correctly." So what? Given that there ARE "correct" models out there beyond the CSG, I suspect that some of the extra-CSG modelers could learn much from PCTers, and vice-versa. The first step to this potential outcome is abandoning assertions that everybody is incorrect except us.

>I have looked into countless papers on control theory and behavior, and
>have yet to see any of the significant insights of PCT even recognized.

OK, so if they're not "incorrect," they're unappreciative. You have to consider what they're trying to do -- what they want. Perhaps if they were treated a bit fairer with respect to their "correctness," some of their horizons might expand at least a bit, and they could appreciate the "significant insights of PCT" much better.

>I have seen nothing on the subject that couldn't have been said in the 1950s,
>except perhaps for some mathematical techniques that weren't
>known then. Maybe you've seen things that I haven't; that's very
>likely, considering the relative merits of libraries to which we
>have access. But as you have withdrawn support in that regard,
>that's just the way it will have to be.

I hereby restore support. I think further exploration of PCT-correctness via specific examples from the non-PCT literature could be enlightening. I've got some papers from HUMAN FACTORS in mind. But first tell me whether you think the Milsum example counts as "behavioral science."

>Rick Marken (930713.2200)

GW>>OK, what counts as applying the principles of control "correctly"?

>An excellent question, though I would add "to living systems" to the end
>of it. That is what PCT is really about. Control theory was around for
>quite some time before PCT came along. Powers did not invent (or

>improve) control theory; he realized how control theory applied
>to behaving systems -- no mean achievement.

So answer the question, already.

GW>>I'm at a loss to understand what could be incorrect about models which are
GW>>the same as those used by some PCTers, so I'm REALLY interested in
GW>>why you think the manual control theorists who use such models are
GW>>applying the principles of control "incorrectly."

>Well, then it's about time you found out. This is really the heart of PCT.
>PCT is not just control theory; it's how it's applied. It makes a BIG
>difference whether you apply control theory to living systems haphazardly
>(as it's currently done) or carefully (as it's done in PCT). The same
>tool can yield good or bad results depending on how it's used.

So answer the question, already.

>But it sounds like Milsum got the diagram right; kudos. Is
>he still around? Maybe he would be interested in PCT.

Maybe so.

>Excellent! Does he describe any research aimed at determining controlled
>variables?

I already said yes to this; for example, in his precis of the Benzinger work, and elsewhere as he describes control models developed by others, several of whom did (or used others') experiments to try to determine controlled variables of particular systems.

>I have never seen this book; it definitely seems like an example (based on
>your description) of a correct application of control theory to behavior.

Interesting. Milsum's book was on top of a pile in my office; the example I presented was at the front of the book. Just happenstance that I came upon a rare "correct" application of control principles to modeling living systems, I guess.

>I have never seen a manual control book or article
>(and I have seen quite a few) that got the basic control diagram correct
>OR that talked about determining what variable was being controlled.

Wel, I think I have. More on this after Bill tells me what counts as "behavioral science."

>Milsum seems to have gotten the application of control theory right.
>I'll believe it when I see the book (maybe at the meeting)...

I'll be sure to bring it. In the meantime, you might try ILL.

>So can that be the end of the argument?

I'd prefer that it be the beginning of a tearing-down of artificial walls around independent groups exploring models of living control systems, so that those in each group can learn from each other.

As ever, Greg

Date: Wed Jul 14, 1993 7:16 am PST
Subject: Re: PCT correctness

From Tom Bourbon [930714.0855]

After four days out of town, I discovered a stack of csg-l mail that includes the discussion of "PCT correctness." Skimming through it, I have the impression this is a game of "gotcha." Some players possess detailed facts, in this case, specific publications in which writers who are not affiliated with PCT have identified clear examples of the phenomenon of control and, further, have used workable, but non-PCT, versions of control theory to explain their examples. When other players, unaware of those writers, say there are none, those in the know say, "gotcha," then insist that their previously uninformed "competitors" declare aceptance or rejection of writers whose work they still have not read in the original. I haven't finished reading the accumulated mail, but I almost expect to come across a score sheet.

Isn't this a forum where we share information that might be useful to a common effort? Does anyone else think it might be more to our collective advantage if we were to post and keep a record of any "correct" applications we find? We could also try to locate the authors, if that is possible, and engage them in the discussion. I have tried that a few times myself. I know Bill Powers has done it many times, even when the authors were only remotely "on the right track" concerning control. I think the usual, if not universal, result of our overtures has been a big silence, a "no reply" or "no acknowledgement" from the other people, which I believe says something about their level of interest in helping develop a general science of control by living systems.

Many of my files are still disorganized after recent moves, but I will try to locate and post some of the examples I have found of correct and nearly-correct descriptions and modeling or both. Maybe some others on the net can do the same. I certainly would appreciate any additional examples I could get.

Until later, Tom Bourbon

Date: Wed Jul 14, 1993 8:19 am PST
Subject: Re: Higher level investigation

From Tom Bourbon [930714.1036] Tom Hancock (930713.0735)
>
>Gary Cziko (930710.0400)
>
>>An R square of .57 is only 66% useless.
>
>I appreciate that an R square of .57 is partially useless. (For

>precision sake, isn't it a .43 factor of uselessness rather than
>..66 as you said. I believe that the coeffecient of alienation does
>not demand a square root as you indicated-what would be the
>point?)

Gary had it right, Tom. The coefficient of alienation (uselessness) for R-squared = .57 is .66 -- and that is pretty useless. I will make a point of bringing a reference or two on this topic to CSG later in the month.

>But Gary, isn't it also partially useful-57%-which is much
>better than no evidence? No, it is not a precisely fit model. But
>it is a modicum of evidence of higher level PCT-based
> predictions, in an area where not much exists?

Also at CSG, you can watch some of the good old tracking tasks like those Bill and I used in "Models and Their Worlds." You will see that the S-R and plan-driven models produce obviously wrong predictions that correlate as highly with human performance as the correlations you reported -- some correlate even more highly. And, true to the spirit of the coefficient of uselessness, you can see, immediately, that it would be impossible to use those models to predict what will happen in future runs by a person. As hard as it is for some of us to abandon the incorrect things we were taught in school, poor correlations are poor correlations; what most of us were *not* taught is that, in fact, those correlations are useless when it comes to saying anything at all about how a particular person will act.

>The trend for the individual subjects was as follows: Every
>subject (16) in the third of the subjects who had the highest
>achievement showed significant effects (alpha = .0001) for the
>PCT generated prediction.

Did you do a mathematical "trend analysis" on your data, or are you following the convention (trend?) in the behavioral sciences and saying there is a trend when what you mean is that some people had scores like the ones you expected. There is a big difference between the two usages.

Until later, Tom Bourbon

Date: Wed Jul 14, 1993 8:29 am PST
Subject: games

[From: Bruce Nevin (Wed 930714 11:15:51 EDT)]

Eric Berne made some astute observations about things people do together. Although his explanations may be of limited value (e.g. his reconstruction of Freudian superego, ego, and id as "parent", "adult" and "child" components of personality), his observations were drawn from a good deal of clinical experience and can be checked against current experience and extended (or grounded) with tests for controlled perceptions.

For example, in a "game" called Why Don't You Yes But" a person poses a problem or asks a question, others respond with "why don't you" proposals, and the questioner parries the responses with reasons why the proposal wouldn't work. The reasons for rejecting a proposal sometimes suggest that something more than the obvious is going on. For an unsubtle example, the person may say "I tried that and it didn't work" when he in fact hadn't tried it and may go on actually to try it later, alone. (Rich Janda tells me that Noam Chomsky does a version of this, rejecting an argument that he will later incorporate into a different discussion as his own, without attribution.)

Berne's description suggests that the person who is "it" is controlling a perception of being in authority in a discussion, loosely analogous to the comparator in a control system, and a perception of others thought to be powerful and competent being shown up as powerless to help him with his problem. It may be that the person who is "it" seldom otherwise has such perceptions, and enjoys them. In any case, the perception of being "one-up" or in the "comparator" seat seems to be the motivation.

There are many other patterns of social transactions in Berne's '60s pop-psych book _Games People Play_ (and in the later books _Scripts People Live_ and _Beyond Games and Scripts_) which may bear reconsidering. Perhaps someone here has done this.

Bruce bn@bbn.com

Date: Wed Jul 14, 1993 1:43 pm PST
Subject: GENERATIVE MODELS - RKC

FROM: Bob Clark (930714.04:15 PM EDT) Bill Powers (930706.1800 MDT)

Following my usual procedure, I've reviewed your discussion of this topic, expecting to select portions for further discussion.

In doing this, I "am reminded" of items from my memory for which I find a recording -- or portion of a recording. I also search my memories for related matters, words, concepts, phrases, etc that may be helpful. In addition I try to imagine extensions and/or combinations of these memories for help in finding relations between your presentations and my memories.

When this doesn't seem to work, I turn to my "supplemental memories" in the form of dictionaries, encyclopedias, etc. I search these in a similar manner.

Throughout this process, I am comparing the material I find with what I have from you. I am seeking a resemblance that can be used for further discussion.

This description summarizes my view of the DME seeking recordings (memories) that can be used for operating my output functions in producing more words.

In the present situation I have found a good many interesting and related ideas, concepts, etc. But they tend to involve such matters as the history of ethics, philosophy of natural law, history of development of natural science, development of physical science, etc. These are interesting and important topics in their own right, but they would require much more time and effort than I care for at this time.

Perhaps suitable definitions and explanations of your phrase, "generative model" are included somewhere, but I failed to discover them.

I can only report that I failed to receive your message -- that I do not "understand" your discussion.

From your associated remarks I conclude that this topic is very important to you. I would very much like to find a way to include it within my working memory.

Regards, Bob Clark

Date: Wed Jul 14, 1993 1:47 pm PST
Subject: MEMRY/TEMPLTES ETC - RKC

FROM: Bob Clark (930714.04:00 pm EDT) Bill Powers (930706.1800 MDT)

I am surprised by your comments in terms of "templates." In my post, Bob Clark (930705.4:55 PM EDT), I was merely giving a couple of examples of devices that are used for "pattern recognition." The general concept of "pattern recognition" includes a great variety of systems, devices, procedures etc. You know a lot more about this subject than I do.

You describe perceptions in terms of "scalar values" of lower-level signals as follows:

>The model I settled on some time ago, while it has some
>drawbacks, treats perceptions simply as scalar values of
>functions of lower-level signals, each of which is also a scalar
>representation of still lower-level perceptions. By scalar I mean
>that any perception is indicated strictly in terms of amount or
>magnitude, by a single one-dimensional neural signal.

>Of course what we normally think of as perception is composed of
>hundreds or thousands of perceptual signals. In the PCT model,
>each signal represents just one variable attribute of experience
>which can only be present to a greater or lesser degree. When we
>abstract patterns from such collections of signals, we do it with
>a higher-level perceptual function that receives sets of the
>lower signals, applies a transformation typical of that function,
>and reports the result, once again, as a single one-dimensional
>signal.

I have no trouble, in principle, with such a description of perceptions, recognizing that the description of remembered perceptions and reference signals must have a similar structure. As you point out, it ends up with a "single one-dimensional signal" that can result from various different combinations of lower order signals. In some situations, however, distinctions among these different combinations may be important.

In addition, I wonder how your description of perceptions answers the following questions:

1. What is it that "perceives" this final "single signal?"
2. And where is it?
3. Is "attention" included somewhere?
4. How about "consciousness?"
5. How can "anticipation" and "planning" be included in your structure/

It seems to me that your description is a form of "Engineer's Viewpoint," and we do need this viewpoint. However, this system seems to be completely automatic -- given a complete statement of its top-most reference levels, its performance would be completely predictable! Perhaps the seeming variability among individuals is "merely" the result of our lacking the "complete statement."

However, there are viewpoints other than the Engineer's. For some purposes, the "User's Viewpoint" is useful. It seems to me that the User can have access to several levels within the hierarchy, each with its own perceptual characteristics and significance. In addition, the User seems to be capable of perceiving more than one "memory track" "simultaneously," perhaps with rapid switching back and forth among them.

Regarding memory, it can be conceived in terms of a complex, multi-dimensional recording of perceptual signals, as you describe them, representing on-going events as they were occurring.

Any form of recording must, at least, include control of play-back rate, cycle time or the like, resembling ordinary video tapes in this respect. With this feature, play-back of "rotating cubes" even including changing backgrounds seems quite feasible.

"Recording" is suggested because it is a familiar term that includes the essence of memory. Such a "recording" may have limitations of accessibility and/or accuracy, but it generally is:

a) accessible for review (imagination),
b) capable of comparison with other perceptual variables (whether current, or remembered),
c) may be available to serve as a source of reference levels within the hierarchy

The anatomical/neurological nature of memories is another challenging topic.

Nevertheless, quite a bit is known about how memories are formed, their properties, availability and application. For one who is interested in using them, these are the more immediate, but not the only, considerations.

This Post touches on several topics -- perhaps they should be considered separately.

Regards, Bob Clark

Date: Wed Jul 14, 1993 3:06 pm PST
Subject: Re: Sonography;machines;du di

From Tom Bourbon and Andy Papanicolaou [930714.1404]

Last Friday, we said that on Monday we would reply to some of last week's posts about speech. Tom was called out of town for Monday and Tuesday and Andy does not use devices with keyboards, preferring instead devices with a graphite core and one rubber tip. While we play catch up, we will reply to a few posts during the past four days from Bill, Martin, Bruce and Rick.

>[From Bill Powers (930710.0930 MDT)] Tom Bourbon and Andy P. (930709.1355)
--

>Your two guesses were
>
>>Option 1: (a) you record 1.5 sec segments of speech, (b) you
>>produce sonograms of the sounds, and (c) you run a control
>>system to produce duplicates of the sonogram. In this option,
>>the control loops are inside a modeled perceptual function and
>>error signals from the loops are signals that are analogs of
>>the acoustic waveform.
>
>>Option 2: (a) you record 1.5 sec segments of speech, (b) you
>>use control loops to produce sonograms of the sounds --
>>sonograms that have specified features like constant amplitudes
>>in distinct frequency bands, for example. In this option, once
>>again, the control loops are inside a modeled perceptual
>>function and error signals from the loops are signals that are
>>analogs of the acoustic waveform.
>
>You're way ahead of me. The control system for amplitude is
>applied to the raw acoustic wave (an "automatic volume control"):

```
>
>
>      Raw wave ---> Gain ------->  output wave ---------->
>                     ^               |
>                     |               |-- amplitude perception
>                      -- out funct --Comp
>                                     |-- ref amplitude
>
>
```

>I don't plan on using the sonograph as an input function -- it's
>just a way for me to look at the way frequency information is
>distributed in the voice signal.

We don't think we are ahead of you at all. The system you diagrammed is like what we envisioned in our Option 2. You did not think of your PCT gain -control device as an input function controlling a feature in an analogue of the acoustic signal, but that is what it is. And the PCT tuning filters would play a similar function for a different aspect of the acoustic event, frequency. The outputs (error signals) from your control devices would control several degrees of freedom in a sonogram that is an analogue of the acoustic events in speech.

>What I'm doing here is busily re-inventing the wheel. Last night
>I discovered formants. My sonograph resolves individual harmonics
>of the fundamental voice frequency. When I do a glissando up and
>down through an octave, maintaining a single vowel sound, all the
>frequencies move slightly up and down the display (which covers
>over 6 octaves) and the harmonics spread farther apart -- but the
>clusters of bright patches on the sonogram remain in the same
>place! I thought that raising and lowering the pitch would move
>the positions of the patches, which is why I was thinking of
>frequency-tracking filters to eliminate inflection effects.

Would a "high-quality" sonograph do the same things? We don't have one, so we don't know. (By the way, videos of those sessions, with you sliding up and down the scale, would bring a pretty penny!)

>What the sonograph shows is not the frequency content of the
>voice sound generator, but the acoustic properties of the throat-
>mouth-nose cavity. When I hold my mouth to make an "O", the low-
>frequency patches alone show up, and remain in the same place
>when I run the voice pitch up and down. When I say "EEE", there
>is a low-frequency patch and two very much higher frequency
>patches with a big gap in the middle frequencies, and again they
>stay about the same as voice pitch goes up and down through a
>octave or so (I can display only up to 4 Khz, so I'm missing one
>patch which I can just see the start of).
>
>I suppose that linguists have known this all the time, but it
>came as a big surprise to me. No wonder they talk about
>articulators! Obviously what is being controlled is the auditory
>signature of particular mouth configurations, excited by a
>harmonic-rich noise source.

That sounds promising, but is it *specifically* what the linguists had in mind when they began to talk about articulatory plans? Our impression is that Liberman and others were not addressing that particular possibility. Had they been, it seems that they would have spoken of the invariances you described -- did they ever notice such things? (These specific questions are directed more to the linguists on the net than to you, Bill.)

> So the perceptual functions we want
>will be based on filters, with their output amplitudes combined
>to pick out the typical signatures of various cavity
>conformations. The relative positions of the formants (I may as
>well use the accepted term) are directly affected by the way the
>tongue, glottis, and lips are arranged. So the control loops will
>be quite direct; all we have to do is to set up perceptual
>functions that use weighted combinations of formant frequencies
>to produce signals that can be controlled by easily separable
>manipulations of the parts of the vocal apparatus. The frequency
>filters will have quite broad bandwidths, and locating them

>properly for a given person's formants only has to be done once.

Again that sounds reasonable and promising.

>The result will naturally look a lot like controlling the
>articulators, but that will be only because their positions are
>closely connected with the auditory signatures actually under
>control. It does mean that one can control kinesthetic sensations
>from articulators and get fairly close by imagining the sound
>that goes with them. But clearly, judging from the experience of
>deaf people, kinesthetic control is not discriminating enough to
>produce the clear sounds of speech. Deaf speech is like
>deafferented motor control: it can be done with a lot of
>learning, but the feedback makes all the difference for skilled
>performance. Maybe one thing that could come out of this would be
>a way to detect the signatures and use some other sense for the
>feedback channel.

Has this possibility been explored before? In our ignorance, we ask the linguists. (Don't read that the wrong way.)
==============================
Subject: Re: Speech Model

[Martin Taylor 930709 11:10] (Rick Marken 930708.1500)
..
[Rick]
>> I think there
>>are about eight vocal tract parameters so there would be eight
>>perceptual functions (this means a minimum of eight inputs to each
>>perceptual function -- an eight filter basilar membrane).

[Martin]
>It's not clear what a vocal tract parameter is, in the human. In the
>software, it's clear. Different models have different numbers. The order
>of magnitude is right, but I think 8 is a bit low (generally speaking
>there is a centre frequency and bandwidth for each of the first three
>formants, and a centre frequency for the next 2 (though these are often
>held constant). Then there may be separate fricative formant control,
>since the time course of the band-shape of the bursts is quite important
>in distinguishing some phonemes. And one must have a control for nasality,
>and then there are some parameters for the impulse shape, which may be
>ignored in some synthesis systems. Let's say 8 to 12 parameters in most
>cases.

Eight to twelve parameters for production. That means 8 to 12 control loops in a simplified Little Baby Articulator, right, assuming that you allow each loop to both push and pull on its degree of freedom in the model for the articulators.

Martin, your example of the state of the art in voice recognition [930712.1415], from Sicharman's column, was a jewel. It certainly helps bring all of our speculations and questions into perspective.

=================================

Subject: doing one's du di

Bruce Nevin (Tue 930713 12:32:13 EDT)] Tom Bourbon [930702.1655]

> SAYING /di/ and /du/

[Bruce]
>I was cut off for a week because the listserver was sending mail to me as
>bnevin@ccb.bbn.com instead of as bn@bbn.com, and ccb got taken away. Am
>just catching up.

We are pleased that you are back on the net. One of our motivations in starting the discussion of du di was our desire to engage you in more discussion on the subject of speech. (Will you be at CSG in Durango?)

>Liberman et al at Haskins do interpret their findings as validating the
>invocation of plans. However, that is because that is all they know
>about, or specifically because they don't know about PCT.

We assumed that was the case. A significant question is, is there a way we can engage them in this kind of discussion, and if so, how? Our first guess is that it cannot be done, given that many people out of the Haskins group have been so vocal in their opposition to PCT. (In this case, the opposition is real, not imagined. For many years, Turvey and Fowler and their associates have actively dismissed and rejected PCT.)

>I have argued (apparently unsuccessfully) that the gestures and their
>acoustic correlates are defined by their systematically mutual
>contrastiveness.

When we first read this post, we began framing questions to ask you in reply, but we discovered that Rick Marken [939713.1300] anticipated most of them. For one,

[Rick]
>I don't understand what this means. I guess I wasn't paying real close
>attention to the early go-round on this. Is there a way to define
>"systematically mutual contrastiveness" so that it can be used in
>a model of speech production/recognition?

We share Rick's uncertainty concerning the meaning of your phrase, "systematically mutual contrastiveness" and concerning how that concept would be implemented in a PCT model for speech.

>I have argued that this accounts for how we so readily accommodate
>variability, though the point was probably unintelligible without the
>first two being accepted, or perhaps merely unmotivated because others
>here are not yet ready to consider the problems of variation and
>diversity in language, which are perhaps analogous to recognizing the
>"same" face as an invariant.

Again, we are not certain what you mean. It is our impression that variation and diversity are treated in nearly all applications of the PCT model -- variable means (actions) to specified unvarying ends (perceptions), achieved in a variable environment, and variable (adjustable) reference signals that specify those ends. Apparently our interpretation of the words variation and diversity, as they might apply to speech, are different from yours. Can you help us identify the differences?

>The problem of "context-conditioned variation" is not really the same
>problem. In this case, the occasion for variation is always identifiable
>within the same utterance, and can be understood as the speaker's control
>of perception of one part of the utterance interfering with control of
>the perception of another part of the same utterance. Consider your
>example of /di/ and /du/. Your drawings follow: (Drawings omitted in this reply -- Andy and Tom)

We did not think of context-conditioned variation as a problem, rather, we reported that Liberman and others at Haskins originally said it was a problem that could be resolved by assuming articulatory plans. As an alternative to their position, you say that CCV can be better explained, with PCT, as:

>...the speaker's control of perception of one part of the utterance
>interfering with control of the perception of another part of the same utterance.

We agree that PCT might provide a more satisfactory explanation for what they called CCV, but we are curious about why you introduce the idea that in CCV the control of one perception is interfering with that of another.

>What can be taken as invariant for the two occurrences of /d/ is a
>projected point at which the transition to formant 2 of the vowel begins.
>I have indicated this point by inserting the character _ into your
>diagram. The transition reflects the movement of the tongue from its
>position closing the oral cavity for /d/ to its position shaping the
>vowel formants for either /i/ or /u/, respectively.

Now this is an interesting possibility. Rick picked up on it and suggested that such an invariant in the speech waveform, associated with an invariant in perceived speech, could be used as the reference signal in a simple PCT model which would then drive the articulators to produce the intended perceptual signal. If such a model were successful, there would be no need to assume an additional step, in whch the model (or a person) imagined the articulatory geatures needed to produce the perceptions. Rick already covered that topic.

Our remaining question, which is the major one we have for you, is whether this transition has been identified in speech waveforms and, if so, whether linguists have taken it into account in their models? If the transition is an acoustic and perceptual invariant, then there is no longer any need for any variety of plan model and it should be possible to develop a PCT model for detection and production of such perceptions. ..

>Now we come to why I say that contrast is fundamental. As you say:

>> (1) There are no known acoustic invariances that our model can be
>> made to control for.
>>
>> (2) There are articulatory invariances - i.e. particular
>> sequences of movements of articulators. These are unacceptable,
>> or appear to be, for reference signals. One reason they are
>> unacceptable is because we may refuse to believe that movement
>> invariances exist at all in spite of the evidence.

>The articulatory "targets" are invariant not as absolute movements, but
>as movements made so as to be distinct from one another.

>I will stop there. If that is not clear, more words won't help now.

It still is not clear. If, as you say, there *are* invariances in the waveform, then a PCT model can be made that will control for perceptions that are analogues of those invariances. It would then follow that the (historically assumed) invariances in the movements of the articulators would "fall out" in the operation of the model as matters of physics -- of changes in the configurations of the devices that produce the specified perceptions. The specified perceptions are distinct from one another; the acoustic events associated with each distinct perception are distinct; and the movements that produce each waveform are distinct. Is it necessary to assume that the *movements* are controlled so as to be distinct -- that there are "articulatory targets?"

Until later,

Andy and Tom

Date: Wed Jul 14, 1993 9:58 pm PST
Subject: Conversion of the PCTers

[From Rick Marken (930714.2200)] Greg Williams (930714)

>Perhaps there are more kindred spirits (even now) than you have realized.

The kindred spirits find PCT (or we find them); most are probably in the CSG. If someone is really a kindred spirit, I am sure s/he will not be driven away, even if we occasionally make hyperbolic claims like "control theory has never been applied correctly to living systems". I mean, if the person "gets it" (control of perception) then s/he probably had the same thought themselves (like "hey, all my colleagues think I'm nuts; they think I don't understand how to apply control theory to living organisms. I wish there were a group like CSG around"). Such a kindred spirit would be thrilled to find some other kindred spirits. As Tom pointed out in his reply to your post, most of these apparent "kindred spirits" are not.

GW>>>OK, what counts as applying the principles of control "correctly"?

RM>>An excellent question,

>So answer the question, already.

OK. Hang on. Heeeere's the answer. The correct application of the principles of control to living systems identifies the variable that is the functional equivalent of the controlled variable (whatever it is called) as a perceptual variable; the output of a sensor. Milsum seems to have gotten this point; I have NEVER seen another non-PCT application of control theory to the behavior of a living system that made this identification correctly.

>I'd prefer that it be the beginning of a tearing-down of artificial
>walls around independent groups exploring models of living control
>systems, so that those in each group can learn from each other.

Who's been building the walls? Bill Powers sure hasn't. I have never seen anyone be more encouraging and accepting of ideas that were even IN THE BALLPARK in an effort to find common ground with current researchers. Short of saying "OK, living systems don't really control perception; and if they do, it's not important anyway" Bill has tried to build BRIDGES to anyone who might have ANY possible interest in PCT -- while maintaining his intellectual honesty.

I certainly haven't been building artificial walls -- I spent 12 years trying to introduce control theory to people who might be interested. I never said to anyone "your approach is wrong" -- I just said "hey, take a look at what this model can do". I did a lot of bridge work and so did Tom Bourbon. Every one of these bridges has been met by a wall.

I am sure that it will be impossible to convince you that the walls that you see around PCT are created by the people you imagine to be the "kindred spirits" -- but they are. Did you notice all the interest we got from Beer? And how about Brooks? Or Bizzi? The reason I've tried to publish PCT papers is to "tease out" the kindred spirits. There ain't none (who aren't already with us). I look through recent issues of all the relevent Psych journals, The Systems, Machines and Cybernetics journal of IEEE and recent books on motor control in the HOPE of finding anyone who seems in ANY way to be on the right track. I have found a couple of friendly reviewers and some promising seeming article writers -- but my comunications have ALWAYS resulted in a dead end -- no interest on THEIR part.

I think you really believe that people would be swarming to PCT if we PCTers weren't so "cultish" and "wall building". I feel like a jew who has been forced by bias into a ghetto with other jews and then told that nobody likes jews because they're clannish. I find your attitude par- ticularly surprising because you experienced PCT ghettoization personally (with Beer, the Science article, etc.). I think the only way to get out of the PCT ghetto (and have all those kindred spirits like us) is to convert. Is that what you are suggestiong? The conversion of the PCTers?

I think you are insultingly wrong about this "wall building" accusation against us PCTers. But why not test this yourself; why not write or call or e-mail one of your kindred spirits and start building some bridges. Let's see if you can do better.

Or keep kvetching about the wall building. Whatever you like.

Frankly, my dear, I don't give a damn.

Best Rick

Date: Thu Jul 15, 1993 4:05 am PST
Subject: No need to read this, Rick

From Greg Williams (930715) Tom Bourbon [930714.0855]

>Isn't this a forum where we share information that might be useful to a
>common effort?

That's what I've been assuming all along. That includes providing evidence to back up one's contentions and one's contentions that others are incorrect in their assertions.

>Does anyone else think it might be more to our collective advantage
>if we were to post and keep a record of any "correct" applications we find?

I certainly do. (I hadn't looked at Milsum's book in a long time -- sometimes it takes a pretty big error signal to motivate sufficiently.)

>We could also try to locate the authors, if that is possible,
>and engage them in the discussion. I have tried that a few times
>myself. I know Bill Powers has done it many times, even when the
>authors were only remotely "on the right track" concerning control. I
>think the usual, if not universal, result of our overtures has been a
>big silence, a "no reply" or "no acknowledgement" from the other
>people, which I believe says something about their level of interest
>in helping develop a general science of control by living systems.

But a few have participated in discussions on the net. And I have seen more than one leave (in my humble, not-cared-about-by-Rhett-Marken opinion) because of the incredible flak they got from some PCTers (especially Bill) not on their models per se, but on their not "getting it" (meta-issues). I myself have brought a few individuals "remotely 'on the right track'" to the net; few of them remain even semi-active now. I believe (Rick, you've got a delete key) that an aura of "we're right and you're wrong" on the net is at least partly to blame. If others dispute this, well, we see things differently. At least (Rick, if you're still there!) I am trying to understand and care about other views on this.

>Many of my files are still disorganized after recent moves, but I will
>try to locate and post some of the examples I have found of correct and
>nearly-correct descriptions and modeling or both. Maybe some others on the
>net can do the same. I certainly would appreciate any additional examples
>I could get.

Sounds great to me. My original point in the "PCT-correctness" discussion was that overblown (in fact or interpretation) statements by PCTers are not necessarily very helpful in spreading the word about PCT, even though they might boost PCTers' morale. The key to limiting the tendency to overblow is more knowledge of what nonPCTers are doing as "close" to PCT work.

>Rick Marken (930714.2200)

>The kindred spirits find PCT (or we find them); most are probably
>in the CSG. If someone is really a kindred spirit, I am sure s/he
>will not be driven away, even if we occasionally make hyperbolic
>claims like "control theory has never been applied correctly to
>living systems".

Speaking personally as a veteran PCTer, sometimes I am close to being driven away. I can imagine (and I could be wrong, certainly) how much more difficult to keep on keeping on it must be for at least some newcomers. The fact is that some once-active participants on the net have dropped out after arguments about meta-issues. And if a prospective newcomer hits hyperbole before or in the first few times he or she logs on, I think the chances of not connecting or a disconnect ("These people are cranks!") is high.

GW>>I'd prefer that it be the beginning of a tearing-down of artificial
GW>>walls around independent groups exploring models of living control
GW>>systems, so that those in each group can learn from each other.

>Who's been building the walls?

Some PCTers AND some nonPCTers, in my opinion.

>Bill Powers sure hasn't. I have never seen anyone be more encouraging
>and accepting of ideas that were even IN THE BALLPARK in an effort to
>find common ground with current researchers.

I heartily encourage such efforts.

>I certainly haven't been building artificial walls -- I spent 12 years
>trying to introduce control theory to people who might be interested.
>I never said to anyone "your approach is wrong" -- I just said "hey,
>take a look at what this model can do". I did a lot of bridge work and
so did Tom Bourbon. Every one of these bridges has been met by a wall.

And I heartily encourage you to continue such bridge-building.

>I am sure that it will be impossible to convince you that the walls that
>you see around PCT are created by the people you imagine to be the "kindred
>spirits" -- but they are.

No, I do see artificial walls around nonPCT groups, too. That doesn't mean that we should give them excuses to build them even higher.

>I think you really believe that people would be swarming to PCT if
>we PCTers weren't so "cultish" and "wall building".

"To PCT" is part of the problem, as I see it. I'm talking about seeing PCT ideas as integrated into a larger effort of modeling.

>I feel like a jew who has been forced by bias into a ghetto with other
>jews and then told that nobody likes jews because they're clannish.

The problem with this metaphor is "feeling like a jew." If you "felt like a scientist who uses control models to study organisms," there wouldn't be the same bias in the first place. A more apt metaphor is the green-haired punk who complains that the cops are always singling him out as a suspect drug-runner. The punk generated his own singularity on purpose -- what does he EXPECT the cops to do when they have discovered that many drug-runners are green-haired punks? He has decided to set himself apart, and the cop treatment is one of the manifestations (of course, some cops actually are bigots about punks, just as some nonPCTers are in love with their own ideas and jealous of others' ideas, so there IS some bigotry there, too). Clearly, the ghettoization of the jews in Europe was largely due to heinous bigotry, and I wouldn't suggest otherwise.

>I find your attitude particularly surprising because you experienced
>PCT ghettoization personally (with Beer, the Science article, etc.). I
>think the only way to get out of the PCT ghetto (and have all those
>kindred spirits like us) is to convert. Is that what you are
>suggestiong? The conversion of the PCTers?

Certainly not. I would rather point to the possibility of becoming integrated with a broader study of control models of organisms. PCTers needn't compromise their intellectual integrity to do that, I think. But they also must not cling to hyperbole about what "no" noPCTers have done.

>I think you are insultingly wrong about this "wall building" accusation
>against us PCTers. But why not test this yourself; why not write or
>call or e-mail one of your kindred spirits and start building some
>bridges. Let's see if you can do better.

Sometimes it is better to speak the truth as you see it and be perceived as insulting. I accept that you are doing that, and I hope you can come to understand that I am, too.

I have already (in fact, just two nights ago, again, in a phone conversation with a user of our Beerbug program) done a lot of recruiting for the net. As I said above, much of it has come to nought, and NOT, I think, just because of the recruitees. If that is insulting, so be it. Regardless, I will continue to recruit. Bill Powers said not long ago that "perhaps groups are temporary expedients." I am SURE that individuals are temporary expedients. The wall-building will pass, sooner or later. No need to be in a rush to glory, is there?

>Or keep kvetching about the wall building. Whatever you like.

>Frankly, my dear, I don't give a damn.

That didn't stop Atlanta from burning, did it?

As ever, Greg

Date: Thu Jul 15, 1993 8:24 am PST
Subject: Re: Sonography;machines;du di

[Martin Taylor 930715 09:30]
(Bill Powers 930710.0930, Tom Bourbon and Andy Papanicolaou 930714.1404)

>>What I'm doing here is busily re-inventing the wheel. Last night
>>I discovered formants. My sonograph resolves individual harmonics
>>of the fundamental voice frequency. When I do a glissando up and
>>down through an octave, maintaining a single vowel sound, all the
>>frequencies move slightly up and down the display (which covers
>>over 6 octaves) and the harmonics spread farther apart -- but the
>>clusters of bright patches on the sonogram remain in the same
>>place! I thought that raising and lowering the pitch would move
>>the positions of the patches, which is why I was thinking of
>>frequency-tracking filters to eliminate inflection effects.
>
>Would a "high-quality" sonograph do the same things? We don't have one, so
>we don't know.

Yes, it would. The auditory system doesn't, quite. It resolves individual harmonics at the lower-frequency end, and individual glottal pulses at higher frequencies. The "bright patches" represent the various resonances of the vocal tract, and are commonly known as "formants". (Incidentally, the term is used not only of speech; it applies to musical instruments as well).

What you see in a sonogram depends on the bandwidths of your filters. If they are wide enough so that 1/W is small compared to the interpulse time, you see no harmonics, but you see the pulses. If they are narrow enough that 1/W is large compared to the interpulse interval, you see the harmonics. Between, you see a resolution cell that may be dominated by harmonics or by pulses, but shows a blurry notion of both. The ear has wider filters at high frequencies than at low.

>>What the sonograph shows is not the frequency content of the
>>voice sound generator, but the acoustic properties of the throat-
>>mouth-nose cavity.

Both, actually.

>>When I hold my mouth to make an "O", the low-
>>frequency patches alone show up, and remain in the same place
>>when I run the voice pitch up and down. When I say "EEE", there
>>is a low-frequency patch and two very much higher frequency
>>patches with a big gap in the middle frequencies, and again they
>>stay about the same as voice pitch goes up and down through a
>>octave or so (I can display only up to 4 Khz, so I'm missing one
>>patch which I can just see the start of).

The "patches" depend mostly on the sizes of the different cavities--a cavity being roughly a region in which the diameter of the vocal tract isn't changing rapidly. Those regions stay the same if you don't move your lips or tongue (more or less), so changing the pitch of your voice by tightening or relaxing the vocal cords has not effect on them.

>>I suppose that linguists have known this all the time, but it
>>came as a big surprise to me. No wonder they talk about
>>articulators! Obviously what is being controlled is the auditory
>>signature of particular mouth configurations, excited by a
>>harmonic-rich noise source.

Yes, exactly. And I don't know how long linguists have known these facts, but I wouldn't be surprised at a century. It is kind of basic knowledge that one tends to think is generally well known, like red blood cells carrying oxygen, or most elements being made in stars and supernova.

>>... all we have to do is to set up perceptual
>>functions that use weighted combinations of formant frequencies
>>to produce signals that can be controlled by easily separable
>>manipulations of the parts of the vocal apparatus. The frequency
>>filters will have quite broad bandwidths, and locating them
>>properly for a given person's formants only has to be done once.

That's quite an ambitious "all." You are dealing with a fairly high dimensional space, in one way of looking at things. From another way of looking at it, "all" you are doing is creating a formant vocoder, and there are lots of those around. But it is not correct to say that one does this only once for a particular person's formants. If that were so, the problem of talker identification would be easy--just locate the person in the space of formant-vs.-intended vowel. It isn't so easy as that.

As Bruce keeps trying to get across, HOW one produces any particular word depends crucially on what other words one is trying to contrast it with. That is to say, in the current dialogue context, what words is the partner likely to be expecting, and in what ways do they differ from the particular word one wants to get across. If the partner is likely to be expecting the word you want, and no others with any great probability, then one can slur and elide (as with my "Pres'n') example. If the partner has (to the speaker's belief) no particular expectation, then the word must be pronounced with maximum clarity--each phoneme is spoken with maximum deviation from a neutral sound. But the next time the same word is spoken, the partner is likely to have a higher expectation for it, and the speaker uses more neutral sounds. You can see for yourself that this effect occurs, by taping a conversation in which someone introduces and then follows a new topic, and listening just to the word that defines the topic on its first introduction and on subsequent times it is spoken.

(I suspect the same thing is happening on a higher level of abstraction with the "ghetto" talk that Rick often indulges in. He feels that PCT is not understood to be distinguished from other approaches that invoke control theory, so has to emphasize the distinctions, mainly for those who he believes not to understand that the distinctions exist.)

>Eight to twelve parameters for production. That means 8 to 12 control loops
>in a simplified Little Baby Articulator, right, assuming that you allow each
>loop to both push and pull on its degree of freedom in the model for the
>articulators.

Yes, I think so. But not in the proof-of-concept syntax recognizer. That will have only 3 degrees of freedom, so we can visualize it as a tracking task in a 3-D space.

(To Bruce Nevin)
>We share Rick's uncertainty concerning the meaning of your phrase,
>"systematically mutual contrastiveness" and concerning how that concept
>would be implemented in a PCT model for speech.

I think this warrants a Durango discussion. There is a difference in the behaviour of control systems depending on whether the CEV is a function whose arguments are passive elements of the outer world or is an aspect of another control system. The use of language is to change one's perception of the state of another control system. It is not clear to me that this can be done without dealing with issues of internal models (of the other control system). Maybe it can, but I don't see how. In most discussions of PCT, the ECSs do not incorporate models (though I tend to imagine them as containing models as an aspect of the imagination loop).

In my Layered Protocol model, the critical aspect of how one handles a virtual message and the related feedback (the loop for which passes through the partner control system) depends on three levels of model for three different perceptions. The levels are (1) What I think the situation is, (2) What I think you think ..., and (3) What I think you think I think ... In theory, one could go down an infinte regress like that, but I have never found a need to go beyond three in any practical situation.

The three perceptions ("what the situation is", in the above) are (1) whether the recipient (which may be me or you) has interpreted the message, (2) whether, given that 1 is true, the recipient has got the message right, and (3) whether it is worth continuing. These are independent, although it looks as if 2 depends on 1. It doesn't. 3 is interesting, because there are two reasonse for 3 to be false. One is that the reference signal has shifted to match the current perception (i.e. that the message to be transmitted now matches the perceived state of the partner, even though that perception has not changed--this happens when a higher-level ECS has attained a zero error state in some other way). The other is that the loop seems to have gone out of control--nothing one does brings the perceived state of the partner toward the reference (i.e. there's no way to make the partner understand). That situation calls for reorganization to become highly probable--we try some other way to get the message across.

Now, back to contrast. Contrast is not between words spoken, but between words that the partner might expect to be spoken, insofar as the speaker models the partner's expectation. If the partner is a foreigner, for example, contrast levels will be maintained higher than if the partner is your life partner. You will tend to speak sloowwllyy and CLEARLY to the foreigner, but cryptically and perhaps in a slurred way to a very familiar conversational partner.

One shouldn't confuse partner modelling with preplanning. It has some elements of preplanning, but it can work only in a full-fledged control of perception mode. If you are talking into the open air, like a broadcaster, you have to talk more carefully, both in pronunciation and in choice of wording, than you do if you can see and hear your partner. Then, you are in a preplanning situation insofar as deciding what to say and how to say it (though of course not in the execution. As Bill pointed out, the auditory and other feedback is critical in speaking as you wish to speak; you must be able to perceive what you are doing in the "passive" world. In broadcasting it is the perception of what you are "doing" in other control systems that you lack.)

I'm not sure whether the above is any clearer about contrast than Bruce has been. It does not fit "contrast perception" in as a perceptual function in a classic ECS of any kind, so it fails to answer the request to implement it in a (standard) PCT model for speech. And I'm not sure it is what Bruce has been trying to say--probably not. But, like Bruce, I am convinced that contrast is the essential element of all linguistic structures, from phonotactics to rhetoric.

>Our remaining question, which is the major one we have for you, is whether
>this transition

[the formant transition between the stop release and the steady state of the vowel]

> has been identified in speech waveforms and, if so, whether
>linguists have taken it into account in their models? If the transition is
>an acoustic and perceptual invariant, then there is no longer any need for
>any variety of plan model and it should be possible to develop a PCT model
>for detection and production of such perceptions.

I know of no such invariants that have been discovered in the speech waveform. All there is is a set of correlations of magnitudes that are declared on CSG-L several times per week to demonstrate the total uselessness of the data. These correlations are sufficient to allow people to write newspaper columns, or M.Sc. Theses, using only voice recognition, so, useless though they may be to a pure PCT'er, they are useful to people with carpal tunnel syndrome. But there are no invariants of the "99+% correlation" kind in the speech business.

Martin

Date: Thu Jul 15, 1993 8:43 am PST
Subject: sound spectrography, o di du da day

[From: Bruce Nevin (Thu 930715 11:30:44 EDT)]

(Tom Bourbon and Andy Papanicolaou [930714.1404]) --

> >Bill Powers (930710.0930 MDT)]
>
> >What I'm doing here is busily re-inventing the wheel. Last night
> >I discovered formants. My sonograph resolves individual harmonics
> >of the fundamental voice frequency. When I do a glissando up and
> >down through an octave, maintaining a single vowel sound, all the
> >frequencies move slightly up and down the display (which covers
> >over 6 octaves) and the harmonics spread farther apart -- but the
> >clusters of bright patches on the sonogram remain in the same
> >place! I thought that raising and lowering the pitch would move
> >the positions of the patches, which is why I was thinking of
> >frequency-tracking filters to eliminate inflection effects.
>
> Would a "high-quality" sonograph do the same things? We don't have one, so
> we don't know.

First off (and reflecting agreement with Greg), it's expedient to use technical terms and names of equipment in ways normal for the workers you're addressing. The name for the acoustic phonetic widget you're reinventing is a sound spectrograph, and the image is a sound spectrogram. A sonograph (-gram) is used in maternity wards to obtain a visual image of a developing foetus in a noninvasive way. Carelessness in such matters erects unnecessary barriers between you and people whose expertise you would like to engage, who may be listening in on the net.

It would be helpful too if you would familiarize yourself with some of the basics of acoustic phonetics. The Lieberman & Blumstein book that I cited some months ago is a good one for this:

Lieberman, Philip & Sheila E. Blumstein. 1988. Speech physiology, speech perception, and acoustic phonetics. Cambridge: Cambridge Univ. Press.

This is not an expensive book. Bill, I think you said the library had ordered it through ILL for you? Lieberman is not to be confused with Alvin Liberman at Haskins, of course, and Blumstein has done extensive work with Kenneth Stevens at MIT.

This book would tell you that the vocal tract acts as a filter on the waveform produced by airflow through the vibrating vocal bands in the larynx. A reasonable approximation can be obtained by a physical model of two tubes, one of smaller cross-section entering one of larger cross-section. Moving the tongue occlusion forward, as for an [i] vowel, or back, as for an [u] vowel (while extending the oral chamber by protruding and rounding the lips) changes the ratios. The absolute frequencies of formants are determined essentially by the distance between larynx and, say, the velum: short for children, longer for women, longest for men. The most obvious differences between voices of these categories of people, fundamental frequency or pitch, is determined by characteristics of the larynx, however. It is an interesting fact that a deep voice is easier to understand than a high-pitched one, because (the fundamental being lower) there are more harmonics of the fundamental passed by each band of the filter: there are more data points defining the contour through time of each formant. You saw this in your sound spectrograph when the harmonics spread apart as the fundamental pitch rose, and came closer as you lowered the pitch again. Had you looked at some introductory material first, you would not have been surprised to see the formants stationary under these changes in the fundamental: you would have known that the vocal tract functions as an acoustic filter. Acoustic phonetics and articulatory phonetics are specializations of physics and physiology, not a bit of behaviorism or statistics in them to taint your work. It is only when people try to apply these fields to phonology, the sound-systems of languages, that explanatory entities like plans come in. It is not hard to recognize these intrusions and set them aside.

You don't have to worry about extracting pitch until you concern yourself intonation contours and stress, at least not for English.

> > Obviously what is being controlled is the auditory
> >signature of particular mouth configurations, excited by a
> >harmonic-rich noise source.

> That sounds promising, but is it *specifically* what the linguists had in
> mind when they began to talk about articulatory plans? Our impression is
> that Liberman and others were not addressing that particular possibility.
> Had they been, it seems that they would have spoken of the invariances you
> described -- did they ever notice such things? (These specific questions
> are directed more to the linguists on the net than to you, Bill.)

The auditory signal is often degenerate. There are cases that are clear, like /di/ vs. /du/ spoken carefully in isolation with pauses between them. One can project a common point of origin for /d/ and take it as a reconstructed invariant. The point of origin is reconstructed because it it is not present in either signal, but rather can be determined only by comparing a set of signals all known by independent means to contain the same sound /d/. Alternatively, and I think more likely, one compares the acoustic signature of different consonants with the *same* vowel, doing so for each of the different vowels. The vowel portion is easy to recognize as the same. The transition for /d/ can be seen to contrast with the transitions for /p/ and for /g/ in a parallel way before /i/ and before /u/. I suggest looking at pi-di-gi as compared with pu-du-gu, and seeing what you get. Out of this, at the pre-language, babbling stage, as the child learns to produce intended sounds, I surmise that there comes to be an association of acoustic perceptions of "same" sound with kinesthetic perceptions of "same" gesture. Whether the kinesthetic perceptions of gesture are controlled or not depends, it seems to me, on whether or not control of acoustic perceptions fails or becomes impossible under some circumstances. I surmise that it does, and that under those circumstances of acoustic degeneracy kinesthetic perceptions of intended gestures are included among the inputs of word detectors. But perhaps you are right, and those who say the acoustic signal is inadequate by itself have not looked at their data in appropriate (PCT) ways. (Rick, this is not all I am going to say in support of this notion, so please don't just blurt out the first "yeahbut" that comes to your mind, wait a bit.)

I base my views on many hours experience transcribing speech in English and in other languages, in standard orthography, phonemically, and phonetically. I know that when this work is the least bit difficult it is only by reproducing for myself the sounds that I hear in the tape playback that I can determine what to write. I do not think this is because the elements of the graphical representation (letters, diacritics, etc.) are defined in articulatory terms; they are defined equally in acoustic terms. If it were merely a matter of the acoustic image, I would not have to do this. The kinesthetic component of the aggregate perception of speech is essential to the recognition process, it seems to me. Spend a few hundred hours at this before rejecting the notion on principle, please. (Wait, Rick. There's more to come.)

> >The result will naturally look a lot like controlling the
> >articulators, but that will be only because their positions are
> >closely connected with the auditory signatures actually under
> >control. It does mean that one can control kinesthetic sensations
> >from articulators and get fairly close by imagining the sound
> >that goes with them. But clearly, judging from the experience of
> >deaf people, kinesthetic control is not discriminating enough to
> >produce the clear sounds of speech. Deaf speech is like
> >deafferented motor control: it can be done with a lot of
> >learning, but the feedback makes all the difference for skilled
> >performance.

I have argued only that kinesthetic perception is co-input with acoustic perception for phoneme perception, not that kinesthetic perception supplants acoustic perception. I have suggested that kinesthetic perception is of secondary importance, given the social character of language. I have observed that when acoustic input is inadequate, people repeat words to themselves and try them for fit to the acoustic image, and I have suggested that abbreviated forms of this may be used to disambiguate acoustically obscure segments or syllables.

(Roman Jakobson argued that distinctive features were necessarily acoustic 40 or 50 years ago for this reason. Jakobson, Fant, & Halle pursued this strongly. Generativists have since abandoned a purely acoustic definition because there are no or not enough reliable acoustic invariants.)

> Maybe one thing that could come out of this would be
> >a way to detect the signatures and use some other sense for the
> >feedback channel.
>
> Has this possibility been explored before? In our ignorance, we ask
> the linguists. (Don't read that the wrong way.)

There has been some work applying signals from sound spectroscopy to areas of skin on the back. I don't know the status of this work. A man named Hugh Upton designed eyeglasses with an electret microphone, an analyzer, and a diode array in a projector casting a visual image on the backs of the eyeglass lenses. The analyzer comprised four band-pass filters: ~150-500 Hz for voicing detection, ~1100 Hz (narrow band) for the first formant of some back vowels, ~2500 Hz (narrow band) for F3 of front vowels and relatively low friction (sh, zh, ch), and 4500-85 Hz for the friction of /s/. This was to help him with sounds that it was hard for him to lipread. It is described in:

Gengel, Roy W. 1976. Upton's wearable eyeglass and speechreading aid: History and current status. In Hirsh, Eldredge, Hirsh, & Silverman (eds.) _Hearing and Davis: Essays Honoring Hallowell Davis_, St. Louis, Mo: Washington University Press.

This paper cites a 1968 paper by Upton.

I am confident there is other work that I don't know about.

> (Will you be at CSG in Durango?)

Unfortunately, no. Ironically, I will be going to a conference on the history of linguistics at Georgetown in August. This is because I was asked to give a talk there on Harris and when I declined because I can't afford the conference costs they waived all costs and provided room and board. It would be boorish indeed to turn them down, and here is an opportunity to make some contacts and to spend some time in a good research library. The irony is that most of what is going on at the conference holds very little interest for me, though I may find out differently when I am actually there. I would much rather spend the time in Durango with you folks. Maybe next year we'll be out of the fiscal woods.

I don't know anyone at Haskins and am unfortunately no expert on making contact with academicians or engaging them in discussions leading to their involvement with PCT. I've been working on it with a group (out of the Linguistic Association of Canada and the U.S. (LACUS)) who are concerned with establishing linguistics on a scientific rather than philosophical basis, so maybe I'll get better at it as I go along. I would like to talk with Lieberman at Brown, but haven't felt ready because there are just so many unresolved problems as to how PCT can begin to account for language. "Let me tell you what PCT can do for linguistics!" "OK, what?" "Well, actually, we haven't figured out how to account for the most elementary things about language yet. But look at how well PCT works with other forms of behavior that you're not so interested in." Uphill all the way.

> >I have argued (apparently unsuccessfully) that the gestures and their
> >acoustic correlates are defined by their systematically mutual
> >contrastiveness.

You and Rick Marken [939713.1300] ask what I mean by this.

Look again at pi-di-gi vs. pu-du-gu. The /i/ portion is the same. p-d-g contrast with each other ("mutually") before /i/. They are also mutually contrastive before /u/. All six acoustic signatures of consonant segments differ markedly from one another. However, the differences between them before /i/ are parallel to the differences between them for /u/. Their contrastiveness is systematic. If before /i/ the transition for one comes from a lower frequency than the transition for another, then in parallel fashion before /u/ the transition for the first one comes from a lower frequency than the transition for the other one.

(You see exactly the same preservation of difference under changes of absolute value in Harris's studies of acceptability-gradings of like-form sentences, which he used as criterion for transformation. If no one remembers that, ignore this, the parallel is not useful to you.)

In maximally careful and distinct speech, these differences are maximized. In less careful speech, the range over which the differences are placed may be reduced, but within that range the terms of contrast are still placed for maximum contrastiveness with one another. If there are three terms (actually, there are four, counting /j/, but that is complicated because of the fricative component) then the middle one does not become closer to one of the extremes while maintaining distance from the other extreme as the range is reduced in less carefully articulated speech. A reduction in gain affects the degree of contrastiveness of all contrasted terms, not just one: again, systemic. At the same time, the distance between /i/ and /I/ (peep vs pip) is reduced, and the distance between /I/ and /e/ and /E/ (paper, pep), and so on -- the whole "vowel triangle" is reduced in fact -- while within the reduced space afforded for contrast the terms of contrast are maintained.

This is why the Peterson & Barney display of vowels all over the acoustic map and overlapping with one another is wrong wrong wrongheaded. Statistical fallacy. Take one speaker on one occasion and the vowels are neatly contrastive. Take the same speaker on another occasion and the vowels are equally neatly contrastive. Combine several occasions or (as they did) numbers of speakers in one acoustic map and the value for /i/ on one occasion overlaps the value for /I/ on another occasion, and you cry in dismay that there are no reliable acoustic correlates of the phonemes. This is because of the preconception that phonemes are physical things in the acoustic image of speech or in the articulatory stream of speaking. They are not: they are contrasts. Contrasts associated with (in terms of) acoustic and kinesthetic perceptions, but fundamentally contrasts.

> >I have argued that this accounts for how we so readily accommodate
> >variability, though the point was probably unintelligible without the
> >first two being accepted, or perhaps merely unmotivated because others
> >here are not yet ready to consider the problems of variation and
> >diversity in language, which are perhaps analogous to recognizing the
> >"same" face as an invariant.

I am referring here to variation between one occasion and another for a given speaker, between one speaker and another, between speakers in one community and another, between speakers at one historical stage and another. There are invariances across these differences, such that one speaker understands the other and both would report that they are speaking and understanding the same language. One could repeat the words of the other verbatim and neither notice the differences, or, if the differences were called to their attention, they would agree that they did not matter. These invariances lie not in the acoustic images of sounds (concerning ourselves only with different pronunciations of the same words, syllables, phonemes) but in the parallels between the contrasts in one version and the corresponding contrasts in the other version. If my Philadelphia friend's repetition of "park" sounds like my "pork" it does not matter, because her "pork" contrasts with her "park" in a way nicely parallel to the contrast between my "pork" and "park". It is the contrasts that are primary and fundamental to language, not the absolute values or even ranges of values (statistical fallacy!) found in the acoustic images or in traces of articulatory gestures (myograms, etc.)

Can you not see how this is different from the rat pushing the bar with front foot, back foot, tail, nose, etc. to get the invariant food pellet? And that it is different from the driver maintaining invariant relation of car to road margins by variable means countering variable disturbances? Disturbances do not enter the picture, are not a factor in the language variation to which I refer. And while there is a distal aim of communicating meaning, analogous to the food pellet, that accounts for choosing different words and different ways of linearizing word/meaning dependencies into discourse and into sentences in discourse; dialect variation, degree of care in speaking, and so on are only marginally and indirectly connected with the meanings that one aims to convey. This is why I said (as Rick overlooked):

> >The problem of "context-conditioned variation" is not really the same problem.

> >In this case, the occasion for variation is always identifiable
> >within the same utterance, and can be understood as the speaker's control
> >of perception of one part of the utterance interfering with control of
> >the perception of another part of the same utterance.

What I mean is that control of the perception of producing a /d/ interferes with control of the perception of producing an immediately following (or preceding) /u/, and similarly for other phoneme sequences. The transition between them is an unintended byproduct of controlling the sequence /d/ /u/. The tongue cannot move instantaneously from the configuration that produces /d/ (whether followed by /i/, /u/, or any other vowel, or whether coming after a vowel) to the configuration that results in /i/. That is just a physical fact about the environment. We have been looking for ways to disturb speech, difficult because the oral cavity is remarkably insulated from environmental disturbance. But coarticulation effects ("context-conditioned variations") are examples of disturbances that have been under our noses all along. Try some tongue twisters.

> >What can be taken as invariant for the two occurrences of /d/ is a
> >projected point at which the transition to formant 2 of the vowel begins.
> >I have indicated this point by inserting the character _ into your
> >diagram. The transition reflects the movement of the tongue from its
> >position closing the oral cavity for /d/ to its position shaping the
> >vowel formants for either /i/ or /u/, respectively.
>
> Now this is an interesting possibility. Rick picked up on it and suggested
> that such an invariant in the speech waveform, associated with an invariant
> in perceived speech, could be used as the reference signal in a simple PCT
> model which would then drive the articulators to produce the intended
> perceptual signal. If such a model were successful, there would be no need
> to assume an additional step, in whch the model (or a person) imagined the
> articulatory geatures needed to produce the perceptions. Rick already
> covered that topic.
>
> Our remaining question, which is the major one we have for you, is whether

> this transition has been identified in speech waveforms and, if so, whether
> linguists have taken it into account in their models? If the transition is
> an acoustic and perceptual invariant, then there is no longer any need for
> any variety of plan model and it should be possible to develop a PCT model
> for detection and production of such perceptions.

The projected starting point for /d/ is not actually in any of the waveforms. It can be determined only by comparing waveforms that have previously and independently been identified somehow as containing the "same" /d/. Given this (which I outlined earlier in somewhat more detail), your comments here are puzzling to me. Do you still intend to say this? If so, please clarify.

If the acoustic signal is all that is controlled, then masking the sounds of one's speaking with white noise in headphones should interfere drastically with speaking. When I tried this, it did not. I don't know how it would be after a long period. Reports of people who became deaf as adults show there are problems speaking. My subjective impression is that I was able to control most consonants easily, that fricatives presented more of a problem. Voiceless, tense fricatives f s sh felt about as secure as the stops, though sh was less sure; voiced fricatives v z zh were more wobbly, and the vowels and r most wobbly of all. This seemed to me to be due to the progressive reduction in kinesthetic perceptions in the above series of English phonemes. Where there is actual contact of an active articulator with a passive articulator, I felt secure; where there is reduced contact or no contact my speech felt increasingly unanchored and liable to error. I didn't notice any actual error in a recording of my speech under these conditions, but this was presumably because the trial period was too short. It was from this that I concluded some months ago that control of consonants is predominantly kinesthetic and control of vowels (plus the fricatives and lrwy) is predominantly acoustic. There was no question about voiced-voiceless, and my observation is that deaf people have no difficulty with the voicing contrast, though voice quality is out of control for them.

Hope this helps to clear things up. Now back to work.

Bruce bn@bbn.com

Date: Thu Jul 15, 1993 9:35 am PST
Subject: Uh Huh

From Rick Marken (930715.0800)] Greg Williams (930715) --

[Sorry, just couldn't resist reading it, Scarlett]

>I myself have brought a few individuals
>"remotely 'on the right track'" to the net; few of them remain even
>semi-active now. I believe (Rick, you've got a delete key) that an
>aura of "we're right and you're wrong" on the net is at least partly
>to blame. If others dispute this, well, we see things differently. At
>least (Rick, if you're still there!) I am trying to understand and
>care about other views on this.

This reminds me of my now rare attempts to have a discussion with my dear ol' mom. I keep forgetting that my mother's idea of a discussion is for her to say something and for me to agree -- period. I still occasionally make the mistake of responding to one or another of her more .. well... "inventive" ideas with something other than "uh-huh". If I respond with "well, I don't know; what about the case in which ... " then she immediately replies with " well, I have a right to MY opinion, don't I?" I keep forgetting that all I'm supposed to do in "discussions" with my mom is agree with everything she says. Apparently, this is the way discussions with nonPCTers are supposed to go, too. Then they'll like us PCTers just like my mom likes me -- when I remember that she's always right, that is. So, in the interests of solidarity with other control theorists who are studying the behavior of living systems, here is my reply to everything they say:

Uh Huh.

You can clip this out and send it to all those nonPCTers who are turned off by our obstinacy. See, you CAN learn things from your mother.

>But a few have participated in discussions on the net. And I have seen
>more than one leave (in my humble, not-cared-about-by-Rhett-Marken
>opinion) because of the incredible flak they got from some PCTers
>(especially Bill) not on their models per se, but on their not
>"getting it" (meta-issues).

Bill made specific suggestions regarding the Beer bug model, for example, and was even starting to make helpful prototypes of walking bugs based on controlled perceptions. Beer left because he didn't want to even HEAR suggestions about how to develop his model based on PCT. I doubt if it will bring Beer back, but you can clip out the above reply ("Uh huh")and sent it back to him as my revised comments on his model. If you get anyone new recruited to CSGNet you could send them my reply above before they see all the hostile, meta issue obfuscation from Bill, Tom and the rest of us meta - issue obfuscators.

>The fact is that some once-active participants on the
>net have dropped out after arguments about meta-issues.

Like who? And what meta issue? Izack Bar Kana (sp?) dropped out because he just never bought the idea that control systems control perceptual signals. That's a meta issue???? If people don't want to accept the basic proposition of PCT (that living control systems control perceptual signals) then they have no reason to get involved with PCT. We didn't cause Izack to drop out -- he dropped out because PCT made no sense (and was not useful) to him. Seems like he made a reasonable choice. Suppose that this were "Natural Selection Net"; would you expect creationist biologists to get real interested in the model? Maybe for religious reasons they would -- just to say "no, no it can't be so" but I think they might get turned off by one of the little meta issues on that net.

Oh, and remember, if, from now on, you want to know if any nonPCT application of control theory (or any other theory, for that matter) is consistent with PCT, you know my answer:

Uh huh.

Agreeably Rick (Momma's Boy) Marken

Date: Thu Jul 15, 1993 9:42 am PST
Subject: Martin on contrast

[From: Bruce Nevin (Thu 930715 12:57:18 EDT)] Martin Taylor 930715 09:30

We're talking about intersecting aspects of the same thing in your discussion of contrast and mine. Lowering the gain on care of pronunciation of a given word after its first mention is an excellent example of why pronunciations vary, fitting into the range of possibilities I tried to delineate. This example is especially interesting because the relation to meaning is unusually close, in the form of redundancy after first mention. It is the simplest case of the reduction system in operator grammar, which also covers all of morphophonemics and the "extended morphophonemics" of syntactic paraphrase.

Boy, do I wish I were going to make it to Durango this year!

Bruce

Date: Thu Jul 15, 1993 1:14 pm PST
Subject: Hancock on Stats

[from Gary Cziko 930715. 0250 UTC] Tom Hancock (930713.0735)

Tom, sorry I didn't get back to you sooner. I missed your message the first time around. I see Tom Bourbon has already responded, but I'll add my 2 cents and hope it makes a significant (p < .05) difference.

>I appreciate that an R square of .57 is partially useless. (For
>precision sake, isn't it a .43 factor of uselessness rather than
>.66 as you said. I believe that the coeffecient of alienation does
>not demand a square root as you indicated-what would be the point?)

The coefficient of alienation (k) is indeed the square root of 1 minus r squared. This provides an index which is the ratio of the standard error of prediction to the standard deviation of Y (the variable being predicted). In this context you can think of the SD of Y as being the standard error of prediction when the mean of Y is used as the prediction of Y for each subject, that is when X (the predictor) is ignored completely.

So a k of .66 means that your error of prediction is .66 of what it would be if you didn't use the X variable at all. This is less than 1.00, to be sure and in some contexts might be quite useful (it sure is a lot better than most k's I've seen in educational research which are routinely in the .80s and .90s), but may not be very useful for predicting any individual's behavior. Indeed, one could argue that a correlation in itself provides no logical basis at all for prediction unless it is indicative of what is going on at a deeper level. I can show you a demo of this at Durango if you'll be there--my (soon to be famous) sliding rubber-band-on-a-string demo.

>>But I'm not sure what you mean by a trend here.

>The trend for the individual subjects was as follows: Every
>subject (16) in the third of the subjects who had the highest
>achievement showed significant effects (alpha = .0001) for the
>PCT generated prediction.

I'm still not sure I follow this, since you are suggesting that you did a statistical test for each subject in your study, which I'm not sure you did. But I agree that if an interesting finding was found for ALL the subjects in a certain group, then you may have discovered an interestng, reliable phenomenon.--Gary

Date: Thu Jul 15, 1993 1:48 pm PST
Subject: Greg - right-on!

Ken Hacker [July 15]

Greg, I must comment on your recent points about PCT correctness. I have never seen such an interesting blend of good intellectual discussion peppered with self-acclaim and concomitant wonderment about why new folks get a little exasperated with the group, as on this conference. But it's part of the territory, I guess. Your point about new people, I think, is a good one. Why put up with BS if you don't have to; it's that simple. On the other hand, I keep tuning in and coming back because I think that PCT offers something I am not getting from others sources. It reminds me of a friend who explained to me how he as stayed married for so many years. He said marriage is 90% good and 10% bullshit, but the 90% makes the 10% tolerable. I think that CSG may be analagous to his idea of marriage.

More seriously, I think that there are two ways that PCT theorists tend to communicate with outsiders. First, there is the way that Powers generally operates: it is based on the claim that others can apply PCT to the work they do and come up with new kinds of data and theory. The others is the assualt approach which attemts to transform any doubter into a whirling dervish by overwhelming him/her with engineering math, control models, or articles and books that must be read. I find the second approach to be offensive and non-productive. I think the first one works best and allows someone new to PCT to test their ideas gradually. Experts in human reorganization should know this.

I am presently collecting thoughts, writing two papers, and reorganizing certain internal standards in related to PCT. I have read the American Behavioral Scientist articles (enjoyed many of them) and am now finishing up Runkel's book (good reading also). In closing, let me say, you will never be accepted by someone you attack. It is better to help someone see the logic you are using and to appreciate it without assualt or conflict. So far, I think that PCT has made substantial progress in gaining supporters and researchers in many disciplines. But the goal of gaining acceptance should never be anything but a natural follow-up to better formulation.

Hasta, KEN

Date: Thu Jul 15, 1993 2:43 pm PST
Subject: Re: flak

[From Bill Powers (930715.0900 MDT)] Greg Williams (930715) --

One person's "incredible flak" is another's attempt to avoid "compromising scientific integrity." And one person's "meta-issue" can be, to another, the main point.

I see my responsibilities to science differently from the way you appear to see them. My objective is not to merge into a general scientific effort, but to nurture an idea and make sure that in the process of being adopted by scientists, it is not lost through assimilation into older and contrary ideas. That, to me, is far more important than going along to get along. The incredible flak I have given some people has been aimed at preventing this assimilation, which would have occurred at the cost of giving up what seem to me the fundamental concepts of PCT. I will not compromise with anyone on those concepts; to do so would negate the most important thing I have been trying to do during my whole adult life. I'm surprised that you would expect me to do so simply to earn the approval of established science.

Your skepticism and your playing the role of Devil's Advocate have been very useful; they have forced me to confront difficult questions and think through things that had only been sketched in. But these roles have sometimes shaded into something less useful: demanding that PCT provide answers in a way that implies that the established approaches have already achieved satisfactory answers. This steers the arguments into conjectural regions where PCT is operating without facts just as much as other scientific conjectures are. While discussions in such areas are useful and suggestive, they necessarily fall back on principles, extrapolating from the known to the unknown. And this, of course, brings other principles into direct conflict with the principles of PCT. If, at this point, one argues that the weight of scientific thought is against PCT, this gives undue credibility to accepted scientific principles, and makes PCT appear to be just another fringe idea on the sole basis that not many people understand it yet.

You seem to be claiming that there is nothing in PCT that other users of control theory don't already know. If that's true, why don't we hear the same sorts of ideas about behavior coming from them? If the reason were that they had considered PCT concepts and found them wanting or incorrect, we would be hearing arguments against these concepts, and rejections of the PCT approach would be cast in quantitative terms showing where the errors are. This is not what has happened. The truth is that most of these ideas have simply never occurred to other control theorists; even though they have an excellent understanding of control theory, they have not seen what it implies about organisms, about the very principles of living systems. To come to that sort of realization, one must organize the control system diagram differently from the customary ways, and be far more careful about matching the components in the abstract diagram to those of the real system. And there must be certain basic realizations, such as what it is that a control system actually controls.

You might think, as I once naively thought, that if the problem is simply that certain concepts had never occurred to a person, that person, on hearing about them, would experience a revelation and explore the new ideas with enthusiasm. But that expectation is based on the tacit assumption that this person had never had any ideas at all about the more global aspects of living systems, and was simply waiting for suggestions. That is far from true.

Everyone engaged in science has been bombarded with speculations about all the same ideas, speculations that have grown up over 350 years of modern science, mostly without benefit of control theory. Thus it is that Milsum can offer a perfectly good (if overelaborate) model of neuromotor tracking behavior, and then offer the idea that the person doing the tracking is being influenced by "subconscious standards," appealing not to a broader control-theoretic model, but to Freud. Where the concepts of PCT are missing, older ideas have filled in the gap, and have effectively prevented further development.

So introducing the ideas of PCT even to researchers using control theory is more than a matter of demonstrating a more comprehensive model. It requires dislodging other explanations of behavior that have become fixed in place, accounting for aspects of human behavior that are not explained by piecemeal control models of low-level physiological processes. PCT challenges those older explanations, and that is where it meets resistance from other control theorists just as much as from personality theorists. The resistance does not come out of control theory, but out of a general background of beliefs and world-views that considerably antedates control theory. Even now in 1993, that background of beliefs and world-views persists, even in those who are applying control theory. And that background developed without any awareness of the process of control.

I have never had any difficulties with control engineers at the level of specific applications of control theory. All the difficulties have arisen when I show how the basic diagram can be rearranged to show the basic PCT principles; the relationship of output to disturbance, the stabilization of input and the control of input relative to internal reference signals. This rearrangement apparently comes as a vast shock, and my way of speaking about control systems on the basis of the rearranged diagram is just too unfamiliar to be tolerated. Look at the reaction of our Israeli participant on the net, before he signed off. He eventually saw what "control of input" was supposed to mean (I think), but ended by saying that he just couldn't get used to giving up the idea that output is controlled, and said goodbye.

If you don't understand control of input, you won't understand control of perception. If you don't understand control of perception, you won't see the relationship of control theory to personal experience. And in modeling the brain, you won't see how one level of controlled perceptions can form the basis for another level of perception and control, so you won't see how coordinations can happen, how people can interact with abstract aspects of the environment, how people can accomplish one thing as a means of accomplishing something else. In short, you just won't get the picture.

To understand PCT you need more than control theory. You need to grasp the meaning of control theory at all levels of human organization, and you need to connect it with your private world. The process of learning PCT is a continual struggle against conventional ideas that you yourself have accepted; one by one you have to see how they are displaced by PCT, and refuted. This requires an active effort, a deliberate re-examination of the commonplace. The people who have done this, or are close to doing it, stay with us. Those for whom the effort is too costly turn away again.

Best, Bill P.

Date: Thu Jul 15, 1993 2:50 pm PST
Subject: Vocoder, contrast

[From Rick Marken (930715.1400)]

Martin Taylor (930715 09:30)

Bill Powers said:

>>... all we have to do is to set up perceptual
>>functions that use weighted combinations of formant frequencies
>>to produce signals that can be controlled by easily separable
>>manipulations of the parts of the vocal apparatus. The frequency
>>filters will have quite broad bandwidths, and locating them
>>properly for a given person's formants only has to be done once.

Martin replies:

> From another way
>of looking at it, "all" you are doing is creating a formant vocoder,
>and there are lots of those around.

Is a formant vocoder a perceptual control system (which Bill described above) or did I just hear a "nothing but"?

>As Bruce keeps trying to get across, HOW one produces any particular word
>depends crucially on what other words one is trying to contrast it with.

How do you know with what other words one is trying to contrast the present word?

>There is a difference in the
>behaviour of control systems depending on whether the CEV is a function
>whose arguments are passive elements of the outer world or is an aspect
>of another control system.

Let x1, x2 ... xn be the arguments of the CEV function -- which I presume means that p = f(x1,x2...xn) with f being what you call the CEV function. It sounds like you are saying that there is a difference in the behavior of the control system with CEV function f() depending on whether the arguments are passive or an aspect of another control system. Let's say that the arguments are the 3-D coordinates of the note sitting next to me. Are you saying that the control system that controls f(x1,x2,x3) is different depending on whether the values of the arguments are determined "passively" -- meaning, I presume, by the laws of physics or by another control system (like my secretary who just plopped it down here)? The overt behavior of the control system might be different (because it might come into conflict with the other control system if that control system is also controlling f(x1,x2,x3)) but the control system itself is still the same, right?

>Now, back to contrast. Contrast is not between words spoken, but between
>words that the partner might expect to be spoken, insofar as the speaker
>models the partner's expectation.

This sounds to me like an attempt to control a higher level variable (the extent to which you perceive yourself as being understood) by varying lower level perceptions of acoustical and articulatory variables. People certainly do vary their speech in this way, though whether they are ALWAYS controlling their perception of a contrast between what they are saying and an imagined perception of what they don't want to be heard as saying, seems doubtful -- I bet people do it just to be annoying sometimes.

Bruce Nevin (Thu 930715 11:30:44 EDT)--

> >I have argued (apparently unsuccessfully) that the gestures and their
> >acoustic correlates are defined by their systematically mutual
> >contrastiveness.

>You and Rick Marken [939713.1300] ask what I mean by this.

>Look again at pi-di-gi vs. pu-du-gu. The /i/ portion is the same. p-d-g
>contrast with each other ("mutually") before /i/. They are also mutually
>contrastive before /u/. All six acoustic signatures of consonant
>segments differ markedly from one another. However, the differences
>between them before /i/ are parallel to the differences between them for
>/u/. Their contrastiveness is systematic. If before /i/ the transition
>for one comes from a lower frequency than the transition for another,
>then in parallel fashion before /u/ the transition for the first one
>comes from a lower frequency than the transition for the other one.

So "systematic mutual contrativeness" is a description of an acoustical invariant. The absolute "acoustic signatures" of p-d-g are different but relational signatures (parallel formant movement) are the same (within consonants) so that the "contrast" between p-d-g exists across vowels. This is a very handy acoutical feature, no?

>In maximally careful and distinct speech, these differences are
>maximized. In less careful speech, the range over which the differences
>are placed may be reduced, but within that range the terms of contrast
>are still placed for maximum contrastiveness with one another.

Ah, this is where the problem arose. You are talking as though the physical "contrast" is the aspect of the acoustical (or articulatory) perception that is being controlled. This is unnecessary; all you need are perceptual functions that respond maximally to each of these features. A higher level system could control the contrast between the outputs of these functions -- but "contrast control" is not required in order to produce what turn out to be contrasting perceptions. For example, I could have someone make trapezoids and rectangles of various sizes on the computer screen. The absolute lengths of the opposite sides of these figures are always different (just as the absolute formant transitions of p-b-d are always different) but there is a systematic contrast between the ratios of their sides (just as there is systemtic contrast in the formant trans- itions); the rectangle ratio is always close to 1; the trapezoid ratio is always <>1. So you could say that the subject is making the distinction between trapezoids and rectangles by controlling the contrast between them; in fact, all the subject has to control is the ratio of sides -- when this ratio is kept near 1 you get a rectangle; when its <> 1, you get a trapezoid. Same with b-p-d with the different following vowels being like the different sizes of rectangle and trapezoid.

Best Rick

Date: Thu Jul 15, 1993 5:07 pm PST
Subject: explaining a training task

[From Tom Hancock (930715.1700)

Rick Marken (930709.0800) and whoever cares to peruse some attempts at higher level investigation,

RM
>Maybe we can brainstorm on the net about possible ways to do
>this kind of research.

>I don't know if I can help you but there may be someone else on
>the net who can--which would help us both.

The following is my attempt to apply PCT to modeling some human functioning in a training environment. If possible I would like to go through iterations along this line with a few subjects. How's it seem? Where are the holes, etc.?

THE TASK
The subject is tasked to learn 27 complex associates via a computer-based drill. There are 27 name labels each associated with one combination of 3 configurations, having three possible levels each (a simplification of a real training task). For example, the subject would be perceiving that the spiked waveform, at the top position on the screen, with the medium tone is called dog house. The task is to identify the name label upon display of the waveform and tone. After each trial the subject rates how certain he is that his response was correct (0% to 100%). Then each subject is informed whether he was correct or not and is simultaneously displayed a list of the characteristics of the correct choice (formerly called instructional feedback), e.g Form = spiked, Position = top, Tone = medium. ---------

I see two controlled variables of concern: 1. the associations with the name label, 2. the higher level goal in doing this task. I will discuss each in order following the format in Runkel.

THE CONTROLLED VARIABLE
1. I am guessing that one of the controlled variables in the task is the perception of an association between the 3 configurations (form, position, tone) with their name label. Subjects should vary in how much their associations are distinct or vivid (vivid- BCP); and this degree will be partially determined by a second controlled variable (discussed below).

IF SUBJECT IS NOT CONTROLLING THAT VARIABLE
2. If association is not controlled then the responses times should not vary systematically with each presentation of an item; the screen display will be largely ignored; and the subject will not spend much time accessing previous memories for that item. As a validation of this logic we should see few correct responses where association is not controlled.

DISTURBANCE
3. The disturbance should be the presentation of a new item, three simultaneous configurations in relationship, for which is associated one correct name label. The magnitude of the disturbance should vary as the subject begins to bring various name associations under various degrees of control.

MEASURE THE EFFECTS OF THE DISTURBANCE
4. I take it that one measure of the effects of the disturbance is the response time to choose a name label. Time taken to select a correct response should vary according to the magnitude of the disturbance. A longer time should indicate a greater error signal has been opposed. Also, I assume that the certainty rating is generated from a perception of the error signal--the difference between the present time perceptions and the retrieved or imagined reference for the correct. Thus, there should be a negative correlation between response time and certainty rating- with the mean time at100% certainty being the minimum for every subject.

THE COUNTER EVIDENCE
5. The only time that the subjects response times should show no systematic pattern would be when the subject was intent on getting done with the drill as fast as possible (confirmed by interview). These subjects would not be controlling for association of configurations with name labels. And these subjects would have the lowest correct rate.

THE OPPOSITION TO THE DISTURBANCE
6. Post experiment interviews and some think-aloud protocols have indicated that the subjects are trying to remember the last time they saw the correct with its associations or they are trying to reconstruct their previous mnemonics or they are scanning the screen for any related clues. In other words: they are imagining (BCP); they are trying to perceive previous memories of the correct without being short-circuited. Thus, in the comparison between the present time perception and the reference for the correct memory there should be an error signal produced which will be output

into more memory location or imagination until the error is reduced as much as presently efficient. In the case where there is no error, at certainty of 100%, the response time should generally be at a minimum since there is no opposition to disturbance.

HIGHER LEVEL CONTROLLED VARIABLE
1. The reference for the degree of association, the distinctness or vividness of the association, would be determined by a higher level control system: the goal or concern in doing this task. For example,
 a. I want to understand or be certain of each association , or
 b. I want to perceive correct messages on the screen, or
 c. I want something else such as getting done with this task.

IF SUBJECT IS NOT CONTROLLING THAT VARIABLE
One means of testing that variable is with the post-response information, which informs of correctness and displays a descriptive list of the associates of the correct and incorrect response.
2a. If the subject is not trying to maintain a sense of certainty (or distinct associations, I assume) for responses, then that subject's times spent with post-response information should show no consistent relationship to certainty ratings.
 b. If the subject is not controlling for perceiving correct messages then that subject's times spent with post-response information should show no difference following correct or incorrect messages.
 c. (When a subject just wants to be done with the drill, and some do, then he will treat all post-response information about the same. This very real case will be left out of the discussion following, but has not been left out of my data gathering.)

DISTURBANCE
3a. The disturbance should be the unreduced error signal (still firing) from the previous task--rating certainty following choosing a name label thought to be correct. This signal should still be firing. Again this disturbance has been indicated by the certitude rating.
 b. The disturbance in the second case (where the subject is controlling the perception of correct responses) should be the perception of a message saying that the response was incorrect.

MEASURE THE EFFECTS OF THE DISTURBANCE
4a. Even when the subjects are correct, they should spend more time with post-instructional information when they were less certain of being correct. There should be an inverse relationship between certainty and time. In addition as a way of validation these subject should have a higher rate of correctness-- specifically, even low certainty corrects will tend to perseverate.
 b. When the subjects have received a message that they are incorrect they should spend more time with the message. As a way of validation, these subjects should tend to have a higher error correction rate.

THE COUNTER EVIDENCE

5. Again, only subjects who want to get done quickly and who do not get many responses correct should demonstrate no relationship between certainty and post-response information time and between correctness and time.

THE OPPOSITION TO THE DISTURBANCE
In the general case, opposition to the disturbance could come as follows: The subject would rehearse the trigram (three configurations) with the name label and he might produce a natural language mediator between the three associates and name (a mnemonic device).
6a. In the case where the subject is controlling for certainty and when he has been uncertain, then the post-response information would be used in order to stimulate distinct associations (through repetition and/or mnemonics), match the standard for higher certainty corrects and thus increase the perception of certainty.
 b. In the case where the subject is controlling for seeing correct responses and sees an incorrect message then he would use the post-response information to oppose the disturbance. But if this subject were not controlling as well for certainty, then the time following correct messages would be the same, regardless of certainty.

This type of work seems ideal for me: I can gather data in somewhat controlled environments, but the applications are very direct into the applied domain. Certainty ratings seem to be a way to tap error signals and post-instructional information time seems to be a way to test control. Using certainty ratings (or similar metacognitive devices) is a way to help students be aware of their error signals and the way the subjects treat post-instructional information is a way to help the teacher and students see what their reference standards are.

Thomas E. Hancock

Date: Thu Jul 15, 1993 5:30 pm PST
Subject: Re: Stats and higher level investigation

[Tom Hancock (930715.0250)]

Gary Cziko (930713.0735) and Tom Bourbon (930714.0735),

My understanding of the coefficient of alienation comes from such as the Encyclopedia of Statistical Sciences (1982). No matter, either way you have a point, the variance is not all accounted for.

My point, again, is that something (.66 or .34) seems to be better than nothing. I'm approaching this partially as a casting of nets, since no one has done any testing of specimens like this.

TB
>Did you do a mathematical trend analysis on your data, or are you ..
>saying there is a trend when .. some people had scores like the
>ones you expected.

For my findings with higher level control, I used some typical exploratory data analysis techniques: just exploring data (with PCT glasses) and seeing what is there. So I just said there was a trend.

Tom, what type of mathematical "trend analysis" are you refering to? With time series designs? No, I did not do that.

GC
>I'm still not sure I follow this, since you are suggesting that you
>did a statistical test for each subject in your study.

Yes, I usually run all analyses on a subject by subject basis. In the case cited, every subject (16 of 16) in the top third (mean correctness) was distinguished with a significant effect (alpha = .0001) for studying post-instructional information longer after incorrects. This trend was not at all the case in the lower third (2 of 16). It is assumed that the top group was controlling for seeing correct responses while the bottom third was not.

Another interesting finding is that the top third seems to control for being certain, but the other thirds do not. In the top third 11 of 16 had significant effects for certainty (studying post- instructional informati on longer if certainty is low, even if the response is correct.) It seems that the other subjects who are just concerned about seeing corrects or getting done with the drill, don't bother with the post-instructional information if they are correct, even if they have been wild guessing. This distinction is like in our classes: some want to understand or learn, some want to just be correct or just tell me what to do, and some want something else, like getting out of class.

Re the CSG meeting. I was thinking about coming, but I have not been sure yet that I would be freed up to come. Even if I could the deadline for registration is past!

Tom Hancock

Date: Thu Jul 15, 1993 6:02 pm PST
Subject: Re: Vocoder, contrast

[Martin Taylor 930715 21:30]   (Rick Marken 930715.1400)

>Bill Powers said:
>
>>>... all we have to do is to set up perceptual
>>>functions that use weighted combinations of formant frequencies
>>>to produce signals that can be controlled by easily separable
>>>manipulations of the parts of the vocal apparatus. The frequency
>>>filters will have quite broad bandwidths, and locating them
>>>properly for a given person's formants only has to be done once.
>
> Martin replies:
>
>> From another way

>>of looking at it, "all" you are doing is creating a formant vocoder,
>>and there are lots of those around.
>
>Is a formant vocoder a perceptual control system (which Bill described
>above) or did I just hear a "nothing but"?

Just where above did Bill describe any kind of perceptual control? Bill described an output apparatus, which is what a formant vocoder is.

>>As Bruce keeps trying to get across, HOW one produces any particular word
>>depends crucially on what other words one is trying to contrast it with.
>
>How do you know with what other words one is trying to contrast
>the present word?

I take it you wrote this before reading any further in the message you are commenting on. And you, of all people, should "know" better than to use "know" in that context.

All I can say is: read it again, Sam.

>>There is a difference in the
>>behaviour of control systems depending on whether the CEV is a function
>>whose arguments are passive elements of the outer world or is an aspect
>>of another control system.
>
>Let x1, x2 ... xn be the arguments of the CEV function -- which I presume
>means that p = f(x1,x2...xn) with f being what you call the CEV function.
>It sounds like you are saying that there is a difference in the behavior
>of the control system with CEV function f() depending on whether the
>arguments are passive or an aspect of another control system.

I didn't say or try to imply that the structure of the control system was different depending on whether the CEV was a passive aspect of the world or a perceived aspect of something internal to another control system. I said, and meant, that the *behaviour* of the control system would be different. And so it will. Feedback loops that contain only dissipative components give a control system quite different dynamic possibilities from those that contain components with complex gain, even in linear systems. And we are dealing with systems that are far from linear, and by no means dissipative. So yes, the behaviour will be different.

>>Now, back to contrast. Contrast is not between words spoken, but between
>>words that the partner might expect to be spoken, insofar as the speaker
>>models the partner's expectation.
>
>This sounds to me like an attempt to control a higher level variable
>(the extent to which you perceive yourself as being understood) by
>varying lower level perceptions of acoustical and articulatory variables.

That's taken for granted. It is WITHIN that kind of structure that we are dealing with the control of contrast phenomenon.

>People certainly do vary their speech in this way, though whether they
>are ALWAYS controlling their perception of a contrast between what
>they are saying and an imagined perception of what they don't want to
>be heard as saying, seems doubtful -- I bet people do it just to be
>annoying sometimes.

People make puns, too, and play other games with the contrasts. All part of the same deal--the fundamental factor in using language is to perceive that the partner perceives something in a way you want to perceive the partner perceiving. You have, as usual, a reference for a perception (your own). That perception is of the partner's perception, evaluated by whatever PIF you can use in the ECS that does the control. Your output through the physical world may be vocal, gestural, or anything else. Your PIF may be taking input from the partner's voice, movements, or whatever. But you are trying to bring to a reference value a perception for which the CEV is within the partner. To do that, it is USUALLY most effective to do what normally gets the partner to the desired state, i.e. to use distinct language that contrasts with other things the partner might (as you perceive it) expect you to say.

As to that "ALWAYS"--is behaviour ALWAYS the control of perception?

Martin

PS. I am off the air from now until we meet in Durango. I'd be happy to continue this discussion there.

Date: Thu Jul 15, 1993 6:44 pm PST
Subject: Vocoder, Contrasts

[From Rick Marken (930715.1930)] Martin Taylor (930715 21:30) --

>Bill Powers said:
>
>>>... all we have to do is to set up perceptual
>>>functions that use weighted combinations of formant frequencies
>>>to produce signals that can be controlled by easily separable
>>>manipulations of the parts of the vocal apparatus. The frequency
>>>filters will have quite broad bandwidths, and locating them
>>>properly for a given person's formants only has to be done once.

>Just where above did Bill describe any kind of perceptual control? Bill
>described an output apparatus, which is what a formant vocoder is.

Well, maybe it wasn't that obvious except to a DITW (dyed in the wool) PCTer. Notice that phrase "produce signals that can be controlled". Bill was describing a phoneme generator based on control of the outputs of the frequency filters. The hook ups to the vocal apparatus are needed, not to generate "outputs" for the observer (as is the case with a vocoder) but to generate these perceptual signals (the outputs of the filters) for the phoneme generator itself. As a side effect, this device (when it's built) will produce phonemes for the observer's enjoyment.

>As to that "ALWAYS"--is behaviour ALWAYS the control of perception?

Purposeful behavior is always the control of perception but the other things we call behavior are alwasy irrelevant side effects of the control of perception.

See ya at the meeting. Then we can discuss our contrasting views on contrasts.

Until then, control well.

Best Rick

Date: Thu Jul 15, 1993 8:33 pm PST
Subject: Re: dudidu

[From Bill Powers (930715.2100 MDT)] Bruce Nevin et al.--

At the moment I'm working on what I believe you, Bruce, may mean by "contrasts." A contrast between two sets of inputs exists, I assume, if a perceptual function responds differently to the two sets. I'm assuming that to build up perceptions of different phonemes that will be stable under varying conditions, we must first set up a number of such perceptual functions, each perceiving a different weighted sum of the inputs. The inputs I'm using right now are the outputs of broadly-tuned filters dividing the (my) total bandwidth into about equal segments, overlapping enough to cover the whole range. This is simply a low-resolution sound spectrograph (thanks for the correction, Bruce). The output of each filter is rectified and smoothed to give a quasi-dc signal representing the equivalent amount of signal at the center frequency of each filter. As formants change, the relative magnitudes of the filter outputs also change. I use about a 0.1 sec time constant after the rectifiers.

Before weightings can be worked out, unwanted variations in the signals have to be removed. There's already one control system that converts the raw audio input into a constant-amplitude wave. The output signal of the control system contains information about amplitude that can be used in dealing with consonants and syllable separations, but that's not the current track; I'm just trying to distinguish vowels. I know that's been said to be easy, but I can't do it yet.

The next control system will be one that perceives the sum of all filter outputs and adjusts the gains (all by the same amount) to give a signal having a constant mean output amplitude averaged over all the filters. This is equivalent to normalizing all the signals to a common basis. I'm also going to try normalizing to the magnitude of the lowest-frequency filter output, and the middle one. This will result in all the filter output magnitudes being referenced to the average, and after subtracting the average each filter response then becomes relative to all the other filter responses as positive and negative deviations -- in other words, contrasts. Again, the output of the control system that maintains this constant average output will contain useful information.

The next step is to look at the filter output variations that are left, for a lot of different vowels, to get an idea of how the patterns differ. I've already tried this on an unstabilized signal (one reason I'm going to stabilize it), and have found one weighting combination that creates a nice sequence ih, ee, eh, aah, ah; a single weighted sum that varies smoothly from negative to positive as the formants shift. This should work even better once the frequency responses are normalized.

The "control systems" I'm talking about, by the way, are the kind that Tom Bourbon asked about -- they're entirely within the input function. They're just being used as ways of computing things, as in an analog computer.

I hope to end up with a set of input functions using different weightings that will produce output signals that vary along relatively independent dimensions. This may or may not end up resembling the "sound triangle" you speak of. I'm aiming for dimensions that are sensitive to different kinds of articulator configurations.

I've done a lot of moaning and hooting while figuring out how my mouth works, and have discovered what seem to be the main effects:

1. Elevation of center of tongue in back or middle or both
2. Raising and lowering tip of tongue
3. Widening tongue to shut off air flow past the margins
4. Opening-closing the glottis (or something) to control nasalization
5. Building up pressure against a closed configuration of the mouth.
6. Opening or closing the center of the lips.
7. Widening or narrowing the rounded lip opening
8. Allowing air to escape with almost-closed configurations.
9. Turning vocalization on and off while expelling air (seems to be done far back in the throat -- vocal cords? When vocalization is on with a closed (air-tight) configuration, the result is a vocalized constant. Needless to say, that can't go on for long.

These are obviously controllable perceptions, as my only way of knowing about them directly has been by feel. There are combinations I never use (such as margins of tongue spread and raised, center depressed, producing a very strange "L" or "R" sound). I suppose these come into play in other languages. I'm just modeling myself -- variations later.

The trick is going to be to teach myself to vary just one articulator dimension at a time, while recording the sound. If I had a realistic software vocal tract model I could probably do all this better, but I suspect that such a model would use horrendous amounts of computing time. If I had gobs of money I'd have one built. It shouldn't be too hard to get all the right motions with little servos pushing things around. Anybody got a few hundred K lying about getting moldy?

Bruce, I appreciate the fact that I could just look most of this up in the literature. But I like being in direct contact with the data instead of getting it through someone else's interpretations. I may end up doing nothing more than following well-worn grooves, but perhaps I might accidentally wander off in some productive direction if I don't get too familiar with the grooves first. I don't like to read too much of what other people are doing until I've gone as far as I can alone. I'm just too suggestible. I've had more than one train of thought derailed by too much reading. I've obviously read some of the literature, but the trick is to stop when you can feel the other guy pulling you in some arbitrary direction. It's too easy to get caught in a local minimum. The other guy may have a good solution, but maybe one can accidentally fall into a better one.

Point of possible interest. One of the ground-rules I'm working under is that all computations in this model have to work in present time only: no looking ahead, no looking back; just deal with the current value of the current signal. Of course with integrators and tuned circuits you get cumulative effects, but its only the present-time value that counts. What you can't do with this modeling approach is collect a complete run and then analyze it. You can't, for example, compute the average value of the signal by adding up all the numbers and dividing by N. It has to be computed as the signal passes by. Working models have to be designed this way. of course I have to collect a complete run to sample fast enough, but then I just pass the signal one point at a time through the rest of the model -- more slowly.

Best, Bill P.

Date: Fri Jul 16, 1993 3:28 am PST
Subject: Assimilation vs. even being heard

From Greg Williams (930716) Bill Powers (930715.0900 MDT)

>I see my responsibilities to science differently from the way you
>appear to see them. My objective is not to merge into a general
>scientific effort, but to nurture an idea and make sure that in
>the process of being adopted by scientists, it is not lost
>through assimilation into older and contrary ideas. That, to me,
>is far more important than going along to get along. The
>incredible flak I have given some people has been aimed at
>preventing this assimilation, which would have occurred at the
>cost of giving up what seem to me the fundamental concepts of
>PCT. I will not compromise with anyone on those concepts; to do
>so would negate the most important thing I have been trying to do
>during my whole adult life. I'm surprised that you would expect
>me to do so simply to earn the approval of established science.

I am not trying to denigrate your objective; rather, I am trying to get you to reconsider the means by which you are attempting to achieve that objective. In my opinion (and in the opinions of some netters who have posted private agreement with this), the flak and hyperbole and what might be summed up as "proprietary" stance of some PCTers (as perceived by some nonPCTers) appears not to be helping achieve your objective very well. Those sorts of approaches are leading, in at least some cases, to an IGNORING of PCT ideas by nonPCTers, rather than EITHER an assimilation OR an adopting of the ideas by the nonPCTers involved.

What I'm calling attention to is not a constant phenomenon. Many times on the net, you and other PCTers have had quite reasonable (again, in my opinion) dialogs and (apparently) reached mutually valued understandings (not always AGREEMENTS). But many other times the dialogs have degenerated because of claims by PCTers (apparently) perceived by nonPCTers as grandiose or obfuscating or not very important or unsupported by sufficient evidence or even insulting. Note that I am talking about the perceptions of the nonPCTers involved -- I can fully appreciate that the PCTers involved have quite different ideas about what is going on. The "objective" (mutually shared subjective) facts are that some individuals with (my opinion) the potential for teaching PCTers in important areas (as well as learning from PCTers) dropped out prematurely. If I were in a lighter mood, I would phrase my advice as "lighten up!" A certain amount of concern about having PCT ideas assimilated and not kept properly distinct from other notions is legitimate, it seems to me. But overworry resulting in hyperbolic-sounding statements and chip-on-shoulder-sounding arguments don't appear very productive for getting PCT ideas even CONSIDERED, seriously, by many nonPCTers.

>You seem to be claiming that there is nothing in PCT that other
>users of control theory don't already know.

I tried to be explicit in earlier posts about this. Your detailed HPCT hypotheses certainly go far beyond anyone else's models I've seen. The problem as I see it is more on the other side: you sometimes act as if there is nothing in nonPCT control models that PCTers don't already know. Some nonPCTers DO seem to think that there is nothing in PCT more than what they already are using in their models. I think that they are missing an opportunity to learn a lot by making this mistake. But this kind of mistake doesn't seem to be always unilateral.

>To understand PCT you need more than control theory. You need to
>grasp the meaning of control theory at all levels of human
>organization, and you need to connect it with your private world.
>The process of learning PCT is a continual struggle against
>conventional ideas that you yourself have accepted; one by one
>you have to see how they are displaced by PCT, and refuted. This
>requires an active effort, a deliberate re-examination of the
>commonplace. The people who have done this, or are close to doing
>it, stay with us. Those for whom the effort is too costly turn
>away again.

I think that, given the difficulties you see in learning PCT, "flak" (as perceived by nonPCTers attempting to learn PCT) would make the process even harder. (At least that is what some private posters have told me.) I hope that it does not continue to be so hard that few of the best and the brightest don't give up.

I simply can't resist this, for Rick ("uh-huh") Marken:

In the words of another screen immortal, I'll be back, Greg

Date: Fri Jul 16, 1993 5:04 am PST
Subject: for Bill P.

[From: Bruce Nevin (Fri 930716 08:47:26 EDT)] Bill Powers (930715.2100 MDT)

Bill, I will have to study your post this weekend. Please look at my response to Rick just now, and see if we are indeed intending the same kinds of perceptions when we use the word "contrast".

I am in sympathy with your concerns about being led to a spurious "local minimum" through your vulnerability to suggestion. (BTW, how would you model suggestion and vulnerability to suggestion?) This is a characteristic of adolescent striking forth at whatever age and in whatever realm, no? The cost is a conflict with one's desire to be heard and understood, to have others agree. I know the dilemma. It becomes necessary, I think, to find ways of being less fragile. I could not have resumed my PhD work (set aside in 1970, or maybe 1974) and I could not be now working on my dissertation and various contributions to the field without having decided to trust my crap detectors. Though I have often had to hold my nose and plunge ahead to satisfy requirements! I suppose you have done the same many times, and that your "I wanna do it myself!" note only states an extreme within a larger space of possibilities across which you try to knit together a meaningful fabric.

There is another side. Even when reading it is essential to experience what people are describing in one's own terms and not substitute imagined perceptions based upon the words and pictures. I think failure to do this is much more of a pitfall. Actually doing the kind of self- experimentation that you are doing would be adequate safeguard while learning the relevant physics and physiology--and that is all that I am suggesting that you do. I think that it is easy to distinguish the physics and physiology from the interpretation and theory in a book like Lieberman and Blumstein's, to find out how the vocal tract functions as a filter while ignoring stuff about automatisms, for example. Lieberman has a very strong skepticism about the orthodox views of phonologists, by the way, and I think that is helpful.

Bruce bn@bbn.com

Date: Fri Jul 16, 1993 5:35 am PST
Subject: Rick on contrast

[From: Bruce Nevin (Fri 930716 07:47:50 EDT)] Rick Marken (930715.1400)

> So "systematic mutual contrativeness" is a description of an
> acoustical invariant. The absolute "acoustic signatures" of p-d-g
> are different but relational signatures (parallel formant movement)
> are the same (within consonants) so that the "contrast" between
> p-d-g exists across vowels. This is a very handy acoutical feature, no?

Rick, I thought the ordinary input function (PIF) had access only to current real-time input. The comparison of p d g adjacent to a given vowel, and the perception that their differences there are parallel to their differences adjacent to various other vowels, is not part of reall-time perceptual input when you hear or produce "du". Nor is this parallel present in the acoustical signal "du". When you say this parallel between differences between consonants is a "very handy acoutical feature," you clearly cannot mean that it is a feature of the acoustical image "du". What do you mean? It seems to me that this feature is a perception of a relationship between relationships between an acoustical feature in the environment and other remembered or imagined perceptions of acoustical features that are not presently in the environment. How is this "very handy acoustical feature" perceived? How would you model control of this perception?

> >In maximally careful and distinct speech, these differences are
> >maximized. In less careful speech, the range over which the differences
> >are placed may be reduced, but within that range the terms of contrast
> >are still placed for maximum contrastiveness with one another.
>
> Ah, this is where the problem arose. You are talking as though the
> physical "contrast" is the aspect of the acoustical (or articulatory)
> perception that is being controlled. This is unnecessary; all
> you need are perceptual functions that respond maximally to each of
> these features. A higher level system could control the contrast
> between the outputs of these functions -- but "contrast control"
> is not required in order to produce what turn out to be contrasting
> perceptions. For example, I could have someone make trapezoids
> and rectangles of various sizes on the computer screen. The absolute
> lengths of the opposite sides of these figures are always different
> (just as the absolute formant transitions of p-b-d are always
> different) but there is a systematic contrast between the ratios of
> their sides (just as there is systemtic contrast in the formant trans-
> itions); the rectangle ratio is always close to 1; the trapezoid ratio is
> always <>1. So you could say that the subject is making the
> distinction between trapezoids and rectangles by controlling
> the contrast between them; in fact, all the subject has to control
> is the ratio of sides -- when this ratio is kept near 1 you get a
> rectangle; when its <> 1, you get a trapezoid. Same with b-p-d with
> the different following vowels being like the different sizes of
> rectangle and trapezoid.

Rick, you have a marvellous ability to select one passage out of extended argumentation and respond to that passage in isolation from the remainder. I often have the feeling-response "well, he's ignoring most of what I'm saying, I'll just return the compliment!"

Try this: Control four degrees of obliqueness of quadrilaterals. In the extreme cases, A is a rectangle, B is a trapezoid with sides at 60 degrees, C has sides at 30 degrees, and D has the sides collapsed to 0 degrees (overlapping in a horizontal line, perceptible as a "trapezoid" only by the fact that the length of the line is greater than the length horizontally of the controlled quadrilateral.) (Ignore left-right symmetry--either pick one, or collapse the symmetrical cases as equivalent.) Stipulate an environmental condition such that it requires greater effort to approximate the extremes (rectangle, line). A higher-level elementary control system (ECS) is controlling a perception that another similarly organized control system correctly identifies which of A, B, C, D is intended. If the other has misidentified one, or if for any reason you (the control system) expect that the other may misidentify one, gain goes up on an ECS controlling perception X, and greater effort is expended, with a result that A, B, C, and D are farther apart in the "space" between 0 and 90 degrees obliquity. Without such expectation, gain goes down in control of X, and A, B, C, and D are closer together in that space. The set of them may be produced one side of center (closer to 90) by some group of control systems, so as to avoid the 0 case which may be confused with a horizontal line (let's say angular orientation of a line happens to be controlled by another ECS) and with a similarly collapsed scalene triangle (also independently controlled). But another group of control systems has a convention of producing them closer to 0, which is risky, but they do this because they are controlling a perception of not being members of the first group. This is something like the situation with language.

What is perception X? I have been calling it contrast. The analogy to language is to the relation b d j g in English (or p t ch k, a series that contrasts with the first, or m n ng ...).

Bruce bn@bbn.com

Date: Fri Jul 16, 1993 7:34 am PST
Subject: -graphs and -grams

From Tom Bourbon [930716.0947]

Just a brief note on the lighter side of our discussions about representations of speech sounds.

Bill Powers posted about his work designing, building and playing with what he called his "sonograph," which produces "sonograms." In a reply to him, I repeated his terms and said something about what a sight (and sound) those sessions of his must be.

Bruce Nevin corrected us:

>First off (and reflecting agreement with Greg), it's expedient to use
>technical terms and names of equipment in ways normal for the workers
>you're addressing. The name for the acoustic phonetic widget you're
>reinventing is a sound spectrograph, and the image is a sound
>spectrogram. A sonograph (-gram) is used in maternity wards to obtain a
>visual image of a developing foetus in a noninvasive way. Carelessness
>in such matters erects unnecessary barriers between you and people whose
>expertise you would like to engage, who may be listening in on the net.

In a recent post, Bill even thanked Bruce for the correction.

Hold on there, young fellow! Bill and I can cry "ageism" and accuse you of erecting unnecessary barriers, an actionable offense in post-modern American society!

Spectrographs (devices that produce representations of spectra) and spectrograms (the representations) have been around a long time, like Bill and me. When applied to sound, I believe the early generic names were audiospectrograph (-gram). When applied to the sounds (songs) produced by birds, the terms were, and still are, sonographs (-grams), a usage that significantly predates the bouncing of sonar pings off of foetuses. During the early 1960s, while I was a student working with an ornithologist-evolutionary theorist, we used sonographs and sonograms as tools to establish evidence for and against various taxonomic classifications of birds. (Names like Stein, Marler and Borror come back to me as the authorities at that time.) Naturally, when Bill described his homemade device, and what to me seemed to be his musical exercises for producing waveforms, his use of sonograph (-gram) seemed to be appropriate -- his post conjured up for me a vision of the Greater Southwest Colorado Bearded Ridge Warbler, a generally reclusive species that forms flocks only during the last few days of July and the beginning of August each year. During that period, its warblings begin before sunrise and continue into the early hours of the morning.

Let this be a lesson to all who are tempted to declare agreement with that Devil's Advocate, Greg Wiliams.

Until later, Tom Bourbon

Date: Fri Jul 16, 1993 8:22 am PST
Subject: Contrast; attitudes

[From Bill Powers (930716.0930 MDT)] Bruce Nevin (930716.0847) --

I'll get to the reading, don't get me wrong. I want to wait until I get stuck. I'm already pretty much aware of the physics involved; the physiology is a bit more hazy, particular with respect to nomenclature, but that's just a lookup problem when the time comes.

I understand what you're trying to say about contrast, but I would REALLY like to find a way of achieving the same thing without having to compare each word or sound being heard against all the words or sounds in memory. At the phoneme level, some sounds are genuinely ambiguous; they could in fact be taken as either one sound or another. When that's true, only a higher level system can select one of the possibilities as most likely, in context. There's no point in trying to find differences between sounds when they are not in fact different. So I'm not going to try to solve that problem. All I want for now is a set of perceptual functions that will serve up signals indicating the sounds that are in the auditory signal, with as much discrimination as possible. More discrimination may be possible than seems immediately reasonable; even if one sound component is the same, others around it may differ, and by normalizing to the mean we in effect make the "same" sound be perceived differently. I think that will produce at least some of the effect you're talking about, and without having to search memory. I know the brain is a big fast parallel machine, but I also don't think it does anything the hard way when there's an easier way available. So it's worth looking for the easier way.

Greg Williams (930716) --

I'm grateful that you recognize some episodes of reasonableness in my constant output of flak. I will try to increase the number of them so they become more noticeable. I do hope that you will turn an equally critical eye on people who purport to want to learn about PCT. Perhaps your comments might help improve some of their attitudes, too.

Best to all, Bill P.

Date: Fri Jul 16, 1993 8:41 am PST
Subject: Re: Rick on contrast

From Tom Bourbon [930716.1028]

>[From: Bruce Nevin (Fri 930716 07:47:50 EDT)] Rick Marken (930715.1400)

[Bruce quotes Rick who quotes from: Bruce Nevin [930715.1130]
>> >In maximally careful and distinct speech, these differences are
>> >maximized. In less careful speech, the range over which the differences
>> >are placed may be reduced, but within that range the terms of contrast
>> >are still placed for maximum contrastiveness with one another.

[Rick]
>> Ah, this is where the problem arose. You are talking as though the
>> physical "contrast" is the aspect of the acoustical (or articulatory)
>> perception that is being controlled. This is unnecessary; all
>> you need are perceptual functions that respond maximally to each of
>> these features. A higher level system could control the contrast
>> between the outputs of these functions -- but "contrast control"
>> is not required in order to produce what turn out to be contrasting
>> perceptions. For example, I could have someone make trapezoids
>> and rectangles of various sizes on the computer screen. The absolute
>> lengths of the opposite sides of these figures are always different
>> (just as the absolute formant transitions of p-b-d are always
>> different) but there is a systematic contrast between the ratios of
>> their sides (just as there is systemtic contrast in the formant trans-
>> itions); the rectangle ratio is always close to 1; the trapezoid ratio is
>> always <>1. So you could say that the subject is making the
>> distinction between trapezoids and rectangles by controlling
>> the contrast between them; in fact, all the subject has to control
>> is the ratio of sides -- when this ratio is kept near 1 you get a
>> rectangle; when its <> 1, you get a trapezoid. Same with b-p-d with
>> the different following vowels being like the different sizes of
>> rectangle and trapezoid.

[Bruce]>
..

>Try this: Control four degrees of obliqueness of quadrilaterals. In the
>extreme cases, A is a rectangle, B is a trapezoid with sides at 60
>degrees, C has sides at 30 degrees, and D has the sides collapsed to 0

>degrees (overlapping in a horizontal line, perceptible as a "trapezoid"
>only by the fact that the length of the line is greater than the length
>horizontally of the controlled quadrilateral.) (Ignore left-right
>symmetry--either pick one, or collapse the symmetrical cases as
>equivalent.) Stipulate an environmental condition such that it requires
>greater effort to approximate the extremes (rectangle, line). A
>higher-level elementary control system (ECS) is controlling a perception
>that another similarly organized control system correctly identifies
>which of A, B, C, D is intended. If the other has misidentified one, or
>if for any reason you (the control system) expect that the other may
>misidentify one, gain goes up on an ECS controlling perception X, and
>greater effort is expended, with a result that A, B, C, and D are farther
>apart in the "space" between 0 and 90 degrees obliquity. Without such
>expectation, gain goes down in control of X, and A, B, C, and D are
>closer together in that space. The set of them may be produced one side
>of center (closer to 90) by some group of control systems, so as to
>avoid the 0 case which may be confused with a horizontal line (let's say
>angular orientation of a line happens to be controlled by another ECS)
>and with a similarly collapsed scalene triangle (also independently controlled)
.
>But another group of control systems has a convention of producing them
>closer to 0, which is risky, but they do this because they are
>controlling a perception of not being members of the first group. This
>is something like the situation with language.
>
>What is perception X? I have been calling it contrast. The analogy to
>language is to the relation b d j g in English (or p t ch k, a series
>that contrasts with the first, or m n ng ...).

I have the feeling we are all still talking past one another, but that we are getting closer than we were back in May or June when there were other discussions about mutual contrastiveness. Back then, I emphasized the idea that when people speak, that is not "what they are doing." They are not, as their highest-level intentions, specifying and defending perceptions of speech sounds, or of contrasts between them. Rather, the person varies the actions that produce speech any way necessary to control other, higher- level, perceptions.

I think Rick implies the same idea in his example of a person controlling patterns on a computer screen; the person can produce patterns that are different without specifying or controlling a feature that we might label "contrast." Bruce offers a counter-example in which a person manipulates shapes with the intention of seeing another person correctly identify them. When the first person's intended results do not occur, the person manipulates the gain on a different perception, X (contrast), until the second person is seen making the correct identification. Bruce's example is closer to the situation confronting a person who intends to use language to affect the actions of another (a degree of closeness I am certain Bruce intended).

In a hierarchical system of ECSs, perceptual categories and the degree of distinctiveness between them can be controlled and the ECS that controls them can be equipped so as to vary its gain. Those are all features Bruce desires for his model. But there is another way to achieve varying degrees of distinction (contrast) between the patterns A, B, C and D. Instead of varying the gain on a special-purpose ECS that controls distinctiveness, the gains on the ECSs that produce A, B, C and D can be varied.

In the first person, the higher-level ECS that is controlling for a perception of some degree of correct identifications by the second person specifies only that perception. When the desired degree of correctness is not perceived, the error signals need not necessarily (but might) lead to the first person setting a reference to control a different perception, X or contrast. Instead, the error can serve as a reference for an ECS that varies the gain of the loops that produce A, B, C and D. (This gain-varying ECS might be thought of as part of a reorganization system.) The result would be changes in the gains of the lower loops, which would result in changes in the "precision" or "tolerance" of the products (A, B, C and D), which *might* result in changes in the degree of correctness of identifications by the second person. A system of this kind does not require (a) that the first person have a "model" of the other person's perceptions, (b) a strategy for varying the gain of the lower-level loops (the random E. coli method would do just fine), or (c) references for perceptions of contrast. Notice that the alternative I suggest does not rule out the perception and control of contrast; I merely want to emphasize that they need not always be features of a system for which part of the loop by which it achieves control of its own perceptions passes through and depends on the actions of other systems like itself.

Until later,
Tom Bourbon

Date: Fri Jul 16, 1993 9:34 am PST
Subject: Hancock, Contrasts, Nice Net

[From Rick Marken (930716.1000)] Tom Hancock (930715.1700)--

>Rick Marken (930709.0800) and whoever cares to peruse
>some attempts at higher level investigation,

I will get to it this weekend (I hope) if someone else doesn't get there first (hint hint to other PCTers). I think that a lot can be learned from going over your proposal with PCT glasses on.

Bruce Nevin (Fri 930716 07:47:50 EDT) --

>When you say this
>parallel between differences between consonants is a "very handy
>acoutical feature," you clearly cannot mean that it is a feature of the
>acoustical image "du". What do you mean?

I mean that you could build a /d/detector that would produce maximal output when the /d/ transition occurs; this detector would ignore the absolute level of the transition. Since the relative transition is the same in all /d/'s and distinct from the transitions for b and g this detector will respond to the d feature while the b and g detectors are quiet.

>It seems to me that this
>feature is a perception of a relationship between relationships between
>an acoustical feature in the environment and other remembered or imagined
>perceptions of acoustical features that are not presently in the
>environment.

All I'm saying is that it is not necessary to determine these relationships in order to recognize phomemes. That is, you don't have to compare the acoustical input to stored values of every possible acoustical input that it might be (it would have to be all possible inputs -- how would you know that the comparison set was only b-d-g, for example, until you have determined that the input WAS either b d or g?).

>How is this "very handy acoustical feature" perceived?

Bill P. will probably have that one figured out in a week. You could find out today by calling the people who built the voice recognition system that Martin posted about. That system does a fairly good job of perceiving the difference between d-b-p. I would bet it looks at a slice of the input, divides it into transition and constant postions, identifies the constant portion (the vowel) and the then applies a logical rule to determine the consonant, like "IF the vowel is /i/ AND the second formant in the transition is changing in such and such a way THEN the transition corresponds to /d/".

> How would you model control of this perception?

That's the easy part (except for the problem of building a model of the articulators that create the sounds that are percevied and controlled).

>Rick, you have a marvellous ability to select one passage out of extended
>argumentation and respond to that passage in isolation from the
>remainder. I often have the feeling-response "well, he's ignoring most
>of what I'm saying, I'll just return the compliment!"

Believe me. I'm just trying to understand and get the essential points.

>Try this: Control four degrees of obliqueness of quadrilaterals.

Good, I'm glad you buy my analogy of control of shapes on a computer screen to control of phonemes by a speaker.

> Stipulate an environmental condition such that it requires
>greater effort to approximate the extremes (rectangle, line).

I presume this means that there is a non-linear feedback function; if the shape is being affected by handle position, for example, then more "throw" or possibly more force is required in the regions of handle movement that get the display toward "retangle" or "line" -- right? So it's harder to move the handle in the extreme handle postion regions which produce the retangle or line shape (assuming no disturbances -- a BIG assumption) -- right?

> If the other has misidentified one, or

>if for any reason you (the control system) expect that the other may
>misidentify one, gain goes up on an ECS controlling perception X, and
>greater effort is expended, with a result that A, B, C, and D are farther
>apart in the "space" between 0 and 90 degrees obliquity.

A change in gain will not affect the displayed shapes (at least, not consistently -- it will affect them if the gain goes so low that you lose control). You are saying something here that makes no sense to me. Time to break out the diagrams and the math.

>Without such expectation, gain goes down in control of X, and A, B, C, and D
>are closer together in that space.

I am just not understanding the set up here (the diagrams and math would help). As it sits, the example you describe makes no sense to me.

>This is something like the situation with language.

GREAT! So let's try to clarify this situation -- with diagrams and math, and I could turn it into a computer demo and we could see exactly what you think is happening with phoneme generation (it's a lot easier to create a demo that controls the shapes of quadrilaterals than one that controls the "shape" of phonemes).

So let's go with this example. Could you draw me a diagram of your "quadralateral controller"? We can work from there.

Please try this Bruce. I think it would really help us converge to a common understanding.

Greg Williams (930716) --

>But many other times the
>dialogs have degenerated because of claims by PCTers (apparently)
>perceived by nonPCTers as grandiose or obfuscating or not very
>important or unsupported by sufficient evidence or even insulting.

Could you give some examples?

>I think that, given the difficulties you see in learning PCT, "flak" (as
>perceived by nonPCTers attempting to learn PCT) would make the process even harder.

How about this -- a CSGNet "role play". Why don't you (Greg) pick one of the particularly egregious cases of a PCTer giving flak to a non-PCTer and post the segments of the interchange that you think are particularly bad -- that show the least sensativity to the non-PCTer and that, if done differently, would result in the non-PCTer becoming an ardent PCTer and giving tons of grant money to CSG, some of which could be used to purchase Bill Powers' vocal tract board. Then show us how this interaction SHOULD have been done; what the flak giving PCTer should have said -- how, why, whatever.

You've been very good about pointing out what we PCTers do wrong -- at least in general terms (now I know, for example, that I post like an insensative, narrow-minded boor); how about helping us out by explaining how we can do it right. Please be very specific (as you know, I'm very literal minded); show us the right way to have conducted some of the discussions that you might be talking about: the Beer Bug debate, the control of perception discussion that lost us the Israeli engineer, the information in perception debate, the social control debate, etc. Pick one and tell us the right way to do it, OK? Maybe this could be a topic at the CSG meeting?

I think most of the problems we have on the net are a result of the fact that this written medium is so damn HOT for some reason. I think you'll agree that our interactions in person (with the non-PCTers who sometimes show up at the meetings -- or when we go to Cybernetics meetings) are USUALLY more convivial, no?

Best Rick

Date: Fri Jul 16, 1993 9:47 am PST
Subject: Re: -graphs and -grams

(Tom Bourbon [930716.0947]) --

Humor well taken, Tom. I hope I haven't offended. Maybe we should revive the terminology developed by Alexander Graham Bell for his Visible Speech apparatus.

I have an old 286 and no A/D hardware, so I suspect I won't be warbling with visible side effects for a while. (Eastern loon?) But Bell and others heard and identified formants before there were any devices for displaying them, and Helmholtz (I think) had an excellent understanding of the principles and issues.

I have made some preliminary overtures to people in the R&D division of BBN about using their equipment for the experiment Rick and I were talking about a while back, and for which he sent me some BASIC code for generating noise. Still keeping my nose clean, however, so as to stay employed.

(Bill Powers (930716.0930 MDT)) --

Some of the nomenclature: larynx = glottis, it's the velum that is lowered for nasalization by letting air out through the nose, the broad flap with the uvula hanging from it that is visible at the top of the throat in the back of the mouth.

> I would REALLY like to find a way of achieving the same thing
> without having to compare each word or sound being heard against
> all the words or sounds in memory.

A given vowel sound is UNAMBIGUOUSLY /i/ because of its distance in acoustic space from other phonemes that might be confused with it. That's for speaker A. The SAME SOUND is UNAMBIGUOUSLY /e/ for speaker B (or for speaker A on a different occasion) for the SAME REASON. It can be for the because the whole system of vowels is shifted in acoustic space, preserving the relation of contrast between /i/, /e/, and other phonemes.

Call this systematic shift, if you like.

Consider the analogy to slant of quadrilaterals. A 50 degree slant constitutes trapezoid B in group x of control systems and the same 50 degree slant constitutes trapezoid C in group y. Yet members of group x recognize as group y's C what would be a B among themselves, because the whole system is shifted over in the geometric space being used to express the distinctions.

(Perhaps this is only a calibration step, subordinate to the recognition of words -- which requires only recognition of highly predictable words at the start of interaction.)

Remembering of words that might contrast is important in the early stages of learning language. It is that, evidently, that limits vocabulary to a small set of short words. Then the child learns to control word contrasts as fungible elements constituting words--the phonemes or the phonemic components or the distinctive features. Suddenly vocabulary learning takes off. Contrast with particular other expected words is important after that only when there is ambiguity or misunderstanding with respect to some similar word -- "I said atTRACT, not atTACK!"

The number of phonemic distinctions in a language is fewer than 100, maybe fewer than 25 (depending on how you think they are represented and organized), not so taxing on memory. However, the Land normalization scheme may indeed turn out to be the easy way. I'm not attached to the notion of controlling contrast, and I don't know how words are distinguished and recognized. That's why I haven't spelled out how all of this should be controlled when Rick challenges me to. I just know that contrast is more fundamental than the sounds that are contrasted, paradoxical as that may seem. It's not just a matter of recognizing sounds as one recognizes tables and chairs. They stay put; the correlation of sounds with phonemes doesn't (systematic shift).

(Tom Bourbon [930716.1028])

Some X varies the gain on the ECSs for phonemic contrast. You suggest X might be part of the reorganization system. Whatever works, say I. I'm just a control system, pragmatic to the core.

Bruce bn@bbn.com

Date: Fri Jul 16, 1993 10:27 am PST
Subject: daring young man on the flying trapezoid

[From: Bruce Nevin (Fri 930716 14:10:47 EDT)] Rick Marken (930716.1000)

> Bruce Nevin (Fri 930716 07:47:50 EDT) --

> >When you say this

> >parallel between differences between consonants is a "very handy
> >acoutical feature," you clearly cannot mean that it is a feature of the
> >acoustical image "du". What do you mean?

> I mean that you could build a /d/detector that would produce maximal
> output when the /d/ transition occurs; this detector would ignore the
> absolute level of the transition. Since the relative transition is the
> same in all /d/'s and distinct from the transitions for b and g this
> detector will respond to the d feature while the b and g detectors
> are quiet.

Rick, you are now not talking about "this parallel between differences between consonants," which is determined only by comparing present input with remembered or imagined records of other acoustic images. You are now talking about a feature of the present acoustic image, the present input from the environment. You are apparently unaware of having changed the subject.

Read it again, please.

> >Try this: Control four degrees of obliqueness of quadrilaterals.

> Good, I'm glad you buy my analogy of control of shapes on a computer
> screen to control of phonemes by a speaker.

> > Stipulate an environmental condition such that it requires
> >greater effort to approximate the extremes (rectangle, line).

> I presume this means that there is a non-linear feedback function; if
> the shape is being affected by handle position, for example, then
> more "throw" or possibly more force is required in the regions of
> handle movement that get the display toward "retangle" or "line"
> -- right? So it's harder to move the handle in the extreme handle
> postion regions which produce the retangle or line shape (assuming no
> disturbances -- a BIG assumption) -- right?

Yes, there is some kind of cost to the control system, such that it won't spread the four shapes to the extremes of the space between 0 and 90 degrees unless it is motivated to put out extra effort and incur that cost.

The vocal tract is pretty well protected from disturbances. About all you can do is put objects in the mouth, diddle with inertial/gravitational effects on the body (mouth included), or interfere with drugs. I think it's OK to talk about the normal case, talking without significant disturbance to the articulators in the oral cavity. But if you can see some sources of disturbance that I'm overlooking, please do tell me. They might present some ways of applying the Test.

> > If the other has misidentified one, or
> >if for any reason you (the control system) expect that the other may
> >misidentify one, gain goes up on an ECS controlling perception X, and
> >greater effort is expended, with a result that A, B, C, and D are farther
> >apart in the "space" between 0 and 90 degrees obliquity.

> A change in gain will not affect the displayed shapes (at least, not
> consistently -- it will affect them if the gain goes so low that you
> lose control). You are saying something here that makes no sense
> to me. Time to break out the diagrams and the math.

> >Without such
> >expectation, gain goes down in control of X, and A, B, C, and D are
> >closer together in that space.

> I am just not understanding the set up here (the diagrams and math
> would help). As it sits, the example you describe makes no sense
> to me.

X may be Tom's reorganization widget rather than a controlled perception, I don't care. There is some motivation for the control system to push the envelope within which A, B, C, and D are distinguished, pushing against the cost it incurs by approaching the 90 and 0 limits.

Let's say the obliquity starts as follows: A=30 B=40 C=50 D=60. The other party doesn't recognize a C for what it is, confuses it with B or D. So our control system presents a sequence of shapes including quadrilaterals with obliquity range increased from 30 degrees (the range above) to 60 degrees: A=20 B=40 C=60 D=80. (Corresponding ranges for other controlled shapes, such as scalene triangles, is also increased.) OK, so now the other guy can tell that when you produce a C shape you mean a C shape, and so on. It happens that the absolute value of obliquity that you are now using for C (60 degrees) is the value that you previously were using for D, but that's OK because D is now at 80 degrees.

A change in gain affects the range over which the degrees of obliquity are spread. As the trapezoids are distributed over the available space in a maximally separated way, yes, there are changes in trapezoid shape, i.e. in the absolute values of obliquity.

Don't try to identify a trapezoid here with /d/ in the phoneme case. You could able to make an analogy between the trapezoid shape (obliquity) and a single feature of a phoneme, such as voice onset time, which distinguishes /d/ from /t/.

I'll see what I can do with diagrams, but not immediately. Meantime, please read what I have said a bit more carefully. Thanks.

Bruce bn@bbn.com

Date: Fri Jul 16, 1993 11:58 am PST
From: Control Systems Group Network
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767

TO: Henry James Bicycles Inc / MCI ID: 509-6370
TO: Gary Mobley / MCI ID: 538-7445
TO: Edward E. Ford / MCI ID: 591-3466
TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu
Subject: Re: -graphs and -grams & daring young men

From Tom Bourbon [930716.1333]

Bruce Nevin: 16 Jul 1993 13:28:10
..

>(Tom Bourbon [930716.1028])
>
>Some X varies the gain on the ECSs for phonemic contrast. You suggest
>X might be part of the reorganization system. Whatever works, say I.
>I'm just a control system, pragmatic to the core.

My allusion to the reorganizing system was not intended as a suggestion that X is in the reorganization system, but as a reminder that loops that know nothing about "what is supposed to be going on" can alter the functioning of the system in such a way that an observer might *think* the system is controlling perceptions that might not even exist in the system.

It seems to me that the different positions taken in the discussion on contrast are organized around different interpretations of the first sentence in your brief reply. When you say "ECSs for phonemic contrast," I think of two possible meanings that you might intend: (1) X varies the gains on the ECSs -- the ECSs that are designed to control phonemic contrast; or (2) X varies the gains on the ECSs -- X does so *for the purpose of* controlling phonemic contrast. And I can think of the meaning I intended: X does not know about or care about phonemic contrast; it merely senses error signals from another ECS and, when they exceed a reference magnitude, or perhaps merely as they vary, the output of X is a change in the gain of other loops -- the ones that directly affect A, B, C and D. The loop on top (which is not X and which does not act on A, B, C and D) has a reference only to see the other person correctly recognizing or identifying A, B, C and D.

Nowhere in the system I suggest is there an ECS with a refernce for, or perceptions of, contrast per se. Again, I realize that there *could* be such an ECS, but I maintain that one is not *necessary* and that without one, the system can still act in a way that convinces an observer there is an ECS for contrast. With a system like the one I suggest, "whatever works" might turn out to be more simple than even a pragmatic control system had envisioned.

=================================
[From: Bruce Nevin (Fri 930716 14:10:47 EDT)]

Subject: daring young man on the flying trapezoid

Alright, already! We give up. You have made your point about puns.
Eastern Loon it is!

(Rick Marken (930716.1000)) --

> Bruce Nevin (Fri 930716 07:47:50 EDT) --
..

[Bruce]
>> > Stipulate an environmental condition such that it requires
>> >greater effort to approximate the extremes (rectangle, line).

[Rick]
>> I presume this means that there is a non-linear feedback function; if
>> the shape is being affected by handle position, for example, then
>> more "throw" or possibly more force is required in the regions of
>> handle movement that get the display toward "retangle" or "line"
>> -- right? So it's harder to move the handle in the extreme handle
>> postion regions which produce the retangle or line shape (assuming no
>> disturbances -- a BIG assumption) -- right?

[Bruce]
>Yes, there is some kind of cost to the control system, such that it won't
>spread the four shapes to the extremes of the space between 0 and 90
>degrees unless it is motivated to put out extra effort and incur that
>cost.

[Tom -- now]
There it is again, and it is still slipping past. I mean the implication
that the system has (must have) references for contrast, and now for level
of effort and for cost incurred. Perhaps so. In some cases, undoubtedly so.
But (break out your sonograms), it ain't necessarily soooo. (An allusion to
a song some of us, and probably Alexander Graham Bell, would recognize.)

>The vocal tract is pretty well protected from disturbances. About all
>you can do is put objects in the mouth, diddle with
>inertial/gravitational effects on the body (mouth included), or interfere
>with drugs. I think it's OK to talk about the normal case, talking
>without significant disturbance to the articulators in the oral cavity.
>But if you can see some sources of disturbance that I'm overlooking,
>please do tell me. They might present some ways of applying the Test.

[Tom - now]
The vocal tract is indeed pretty well protected from disturbances, but the
loop through another person is not, and that is the loop that matters, isn't
it? Disturbances there affect the match between a speaker's intended
and actual perceptions of what other people are doing and the speaker varies
the actions of articulators, which vary the resulting acoustic events, and
the person continues varying things until the desired perceptions occur.
All of this *can* occur with no control of either *varying*, or *contrast*.
It is easy to create, or to identify, disturbances in the "other person"
part of the loop and to see how the speaker varies outputs to eliminate the
effects of those disturbances. But if, in the presence of disturbances,

the contrast between outputs changes, that *need not*, but might, mean the
speaker acted to control contrast.

..

Bruce]
>> > If the other has misidentified one, or
>> >if for any reason you (the control system) expect that the other may
>> >misidentify one, gain goes up on an ECS controlling perception X, and
>> >greater effort is expended, with a result that A, B, C, and D are farther
>> >apart in the "space" between 0 and 90 degrees obliquity.

[Rick]
>> A change in gain will not affect the displayed shapes (at least, not
>> consistently -- it will affect them if the gain goes so low that you
>> lose control). You are saying something here that makes no sense
>> to me. Time to break out the diagrams and the math.

[Bruce]
>> >Without such
>> >expectation, gain goes down in control of X, and A, B, C, and D are
>> >closer together in that space.

[Rick]
>> I am just not understanding the set up here (the diagrams and math
>> would help). As it sits, the example you describe makes no sense
>> to me.

[Bruce]
>X may be Tom's reorganization widget rather than a controlled perception,
>I don't care. There is some motivation for the control system to push
>the envelope within which A, B, C, and D are distinguished, pushing
>against the cost it incurs by approaching the 90 and 0 limits.

[Tom - now]
At the conclusion of that exchange, the various positions on contrast seem
closer together, but (to sing my song again) the control system producing
A, B, C and D need not know about an "envelope within which" they are
distinguished. If the system is ignorant on that count, then it probably is
not motivated to push against that envelope, but rather acts to control
something else -- its intended changes in the actions of other systems. For
instances where my alternative system is at work, an observer might still
easily notice, and be misled by features and events such as contrasts,
envelopes and costs.

(It is disappointing to know the Eastern Loon will not extend its range to
Colorado this month.)

Until later,
Tom Bourbon
Department of Neurosurgry
University of Texas Medical School-Houston   Phone: 713-792-5760
6431 Fannin, Suite 7.138   Fax: 713-794-5084

Houston, TX 77030 USA tbourbon@heart.med.uth.tmc.edu

Date: Fri Jul 16, 1993 12:06 pm PST
From: Control Systems Group Network
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767
TO: Henry James Bicycles Inc / MCI ID: 509-6370
TO: Gary Mobley / MCI ID: 538-7445
TO: Edward E. Ford / MCI ID: 591-3466
TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu
Subject: Contrasting subjects

[From Rick Marken (930716. 1230)]

Bruce Nevin (Fri 930716 14:10:47 EDT) --

>(Rick Marken (930716.1000)) --

>> Bruce Nevin (Fri 930716 07:47:50 EDT) --

>> >When you say this
>> >parallel between differences between consonants is a "very handy
>> >acoutical feature," you clearly cannot mean that it is a feature of the
>> >acoustical image "du". What do you mean?

>> I mean that you could build a /d/detector that would produce maximal
>> output when the /d/ transition occurs;

>Rick, you are now not talking about "this parallel between differences
>between consonants," which is determined only by comparing present input
>with remembered or imagined records of other acoustic images. You are
>now talking about a feature of the present acoustic image, the present
>input from the environment. You are apparently unaware of having changed
>the subject.

Boy, am I.

>Read it again, please.

I didn't save the old post and I'm completely confused about what you
mean by the above (as well as by the rest of your post, which I will try
to get to later -- but you didn't give me anything close to the kind
of clarification I need to start building the quadralateral control demo).
I have to go to a meeting now but if you get a chance, could you please

explain what your problem is with the /d/ detector notion? What subject did I change?

Oh, one little point. In a previous post you said:

>I just know that contrast is more
>fundamental than the sounds that are contrasted, paradoxical as that may
>seem. It's not just a matter of recognizing sounds as one recognizes
>tables and chairs. They stay put; the correlation of sounds with
>phonemes doesn't (systematic shift).

Have you ever seen what happens to the retinal image of a table or chair as you move about? Do you seriously believe that the problem of speech perception is any different than the problem of perception in general. If you do, then that's pretty much a show stopper for me. I've studied perception for over 20 years and I've never seen anything anywhere to suggest that speech is any different than any other kind of perception.

Best

Rick

Date: Fri Jul 16, 1993 2:15 pm PST
Subject: uh-uh

[From: Bruce Nevin (Fri 930716 16:09:50 EDT)]

(Rick Marken (930716. 1230)) --

Rick, I have just resent today's posts to you. I hope you will be able to find the time to read them carefully before you throw them away.

> What subject did I change?

The projected point of origin common to Formant 2 for occurrences of /d/ before various vowels can be perceived only by comparing the different occurrences of /d/ before vowels. It is NECESSARILY not a perception based solely upon input of a single acoustic image, say /di/. It is a relationship between several acoustic images, one of which may be input in real time but they could all be in memory or imagination. That was the subject under discussion.

You began talking about perception of some feature of a single acoustic image in real time, say /di/. That is the changed subject. It is not clear to me what you think this feature is. Or is the projected point of origin of F2 apparent from the acoustic image of /di/ all by itself, in some way that I don't understand?

The projected points of origin for formants F1 and F2 associated with /b/ are different from those for /d/ (before the same several vowels). Likewise those for /g/.

These features (or metafeatures, features at one remove from any single actual acoustic image) have a relation to one another that we as investigators can see. They are separated from one another. It appears to be this separation that enables people to distinguish the consonants from one another. We as investigators can refer to this relation as contrast. I will leave unresolved whether "contrast" is a perception of language users, and whether it is a controlled perception that is part of language. That notion is clearly a disturbance.

> >seem. It's not just a matter of recognizing sounds as one recognizes
> >tables and chairs. They stay put; the correlation of sounds with
> >phonemes doesn't (systematic shift).
>
> Have you ever seen what happens to the retinal image of a table or chair
> as you move about? Do you seriously believe that the problem of
> speech perception is any different than the problem of perception in
> general. If you do, then that's pretty much a show stopper for me. I've

Is it the case that the retinal image that you see as a brown chair next to a green sofa in one person's house appears to you in another person's house (same orientation, same lighting conditions, same background) to be a brown footstool next to a green chair? If so, then there is indeed a parallel.

(Tom Bourbon [930716.1333]) --

I just grepped through my posts today searching for "ECS" and I find no occurrence of the phrase you quote me as saying, "ECSs for phonemic contrast." As above (and as I said in at least one earlier post), I am willing to leave up in the air whether contrast is a controlled perception or a byproduct of other perceptions.

However, any "feature" that depends upon comparison of multiple tokens (/d/ before different vowels, or /i/ after various consonants, for example) might as well be termed contrast. I don't see what you gain by merely avoiding the term.

Control of perception of being understood is at a higher level, as shown by the fact that it is at least as often a different choice of words rather than clearer articulation of phonemes that opposes disturbance of that perception.

No ECS perceives the "envelope" or the cost of pushing it as such. It is simply a fact that speaking with maximum distinctness costs an effort that people are generally too lazy to expend, why bother. That's all. So the maximum acoustic distance between /i/ and /a/ and /u/ is greater than the actually exploited distance in a typical pronunciation of words containing these vowels. If a person is speaking with maximum distinctness, then the vowel /I/ of "bin" overlaps with the vowel /i/ of "bean" as heard on a more relaxed occasion. The vowels are differentiated as much as possible within the range given on each occasion; the "relaxed" range of possible differentiation is smaller than the maximum. I don't care how this is accounted for in a PCT model. However, it is a fact that must be accounted for. And it scotches the notion that some feature of the acoustic image [I] is always going to be recognized as the phoneme /I/, since on the "relaxed" occasions it is actually the phoneme /i/. It is not the case that [I] is ambiguous (as Bill suggests): on each occasion, the range of possibilities determines a unique and unambiguous interpretation of the acoustic image.

I'm getting loony, and now I'm at risk of missing my train. I'll be back Monday. Have a good one.

Bruce (16:56)
bn@bbn.com

Date: Fri Jul 16, 1993 3:02 pm PST
From: Control Systems Group Network
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767
TO: Henry James Bicycles Inc / MCI ID: 509-6370
TO: Gary Mobley / MCI ID: 538-7445
TO: Edward E. Ford / MCI ID: 591-3466
TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu
Subject: Re: Stats and higher level investigation

From Tom Bourbon [930716.1650]

>[Tom Hancock (930715.0250)]
..

>TB
>>Did you do a mathematical trend analysis on your data, or are you ..
>>saying there is a trend when .. some people had scores like the
>>ones you expected.
>
>For my findings with higher level control, I used some typical
>exploratory data analysis techniques: just exploring data (with PCT
>glasses) and seeing what is there. So I just said there was a trend.

That is what I thought. There is nothing inherently wrong with the time-honored eyeballing of data, but back in the dark ages, "trend" had a more precise (although not necessarily more legitimate) meaning. (See below for just a few more details.)

>Tom, what type of mathematical "trend analysis" are you refering

>to? With time series designs? No, I did not do that.

In stat texts around the 1960s and 1970s, trend analysis was almost always a topic, often it earned a whole chapter. But fads and fashions come and go in statistics, and certainly in textbooks, so the subject is not always there today.

Trend analysis was used in "causal" experiments of the Independent Variable-Dependent Variable variety. It was a way of identifying whether there was a "trend" (a statistically significant functional relationship) for the DV across levels of the IV. Trends were usually identified as "linear" or "curvilinear," with the latter category including quadratic, cubic, and in a few brave texts, quartic. Almost all texts said no one could make sense out of anything above a quadratic relationship.

Sometimes the subject was labeled as "orthogonal comparisons" or "orthogonal polynomials." That is the form that survives today as "contrast analysis" or "focused tests of significance." Ralph Rosnow and Robert Rosenthal are two names widely associated with that topic.

>GC
>>I'm still not sure I follow this, since you are suggesting that you
>>did a statistical test for each subject in your study.
>
>Yes, I usually run all analyses on a subject by subject basis. In the
>case cited, every subject (16 of 16) in the top third (mean
>correctness) was distinguished with a significant effect (alpha =
>..0001) for studying post-instructional information longer after
>incorrects. This trend was not at all the case in the lower third (2
>of 16). It is assumed that the top group was controlling for seeing
>correct responses while the bottom third was not.

In light of the things I mentioned above, what you have just described is not a "trend," but a result. In the top 1/3 of your people you found one thing -- result. In the bottom 1/3, you did not find that thing -- result. But *that* fact might reveal a trend, in the old-fashioned sense: the greater the amount of whatever defined your upper and lower 1/3rds, the greater the proportion of people who had significant results.

Not that any of this will help you, but you asked.

Until later,

Tom Bourbon

Date: Fri Jul 16, 1993 3:04 pm PST
From: Control Systems Group Network
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: Hortideas Publishing / MCI ID: 497-2767
TO: Henry James Bicycles Inc / MCI ID: 509-6370
TO: Gary Mobley / MCI ID: 538-7445
TO: Edward E. Ford / MCI ID: 591-3466
TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L%UIUCVMD.bitnet@vm42.cso.uiuc.edu
Subject: Contrast' synchronous detector

[From Bill Powers (930716.1530 MDT)]

Bruce Nevin (930716) replying to Rick:

>Rick, you are now not talking about "this parallel between
>differences between consonants," which is determined only by
>comparing present input with remembered or imagined records of
>other acoustic images. You are now talking about a feature of
>the present acoustic image, the present input from the
>environment. You are apparently unaware of having changed the
>subject.

Rick voiced my objection much more clearly than I did. If you're comparing present-time auditory input to remembered images, you don't initially know what the current auditory input is, so you must check every input for contrast with every memory image. Not only that, but you must treat every auditory input as if it is completely unknown: that is, you may be looking for the contrast between an ee and an oo, but you may also be looking for the contrast between aah and ih -- all with the same input. In order to check contrast you must have two already-discriminated signals, one from the current audio input and any one of those in memory. Without that pre-processing, you have nothing to check for contrast.

Now what about the memory images? Here, presumably, we have ideal examples of each possible -- let's call it "phoneme." Perhaps "canonical" phonemes. How, then, is "contrast" detected? It must be through some mechanism that receives the current signal and the remembered "canonical" signal, and produces a signal indicating the degree of difference or contrast. The greater the signal, the greater the contrast of the current input with that canonical phoneme in memory.

But is that what we want? If some totally new sound were to appear, it would contrast with ALL stored phonemes -- but that wouldn't tell us what it is. This "contrast detector" responds maximally when the input is NOT like the stored image -- at least as I interpret the term contrast. This means that for any one input, you would get a large number of different contrast signals. Then you have to postulate a mechanism for selecting one of those signals as meaningful.

It seems to me that what we really want here is just an ordinary input function that responds to the degree to which a particular input pattern is detected. Given a set of such input functions, each tuned to a different input pattern, one of them would respond more than the others for a given input, and that would be taken as the phoneme being heard. Contrast would be indicated by a lack of signal from the other detectors: the less the signal, the greater the contrast between the input and the particular pattern that input function is designed to detect. This is more like what we need: we don't need a whole lot of contrast detectors reporting loudly that this input is NOT a particular sound.

Elsewhere you point out that the very same sound can mean /e/ when said by one person and /i/ when said by another. But this sameness is in terms of the objective sound spectrogram. If the perceptual signals from the various phoneme detectors are averaged and the average is subtracted from each signal, what began as an /e/ could very well end up as an /i/ as it should. This is because the normalization removes any differences from one person to another that affect ALL frequencies (such as differences in the size of the mouth cavity). With the use of weighted sums, even more sophisticated ways of removing common pattern changes can be implemented. Once past this stage, the following detection stage receives a more or less standard input from which the major individual differences have been removed.

General note for modelers:

I have now tried a "synchronous detection" filter and it works like a dream. Basically you connect the audio signal to the input of an ordinary tuned filter. Then you multiply the output of the filter by the input. The output and input will be exactly in phase at the center frequency of the filter, so their product will always be positive. For input frequencies away from the center frequency, there will be a positive or negative phase shift which tends to 90 degrees far from the center frequency (leading or lagging). Thus the product for those frequencies tends to zero. Smoothe the product and you have an envelope-following signal. Sounds complicated, but here's the code:

```
   x += input + ff*y - x*damping;
   y -= ff*x;
   out += sf*(input* y - out);
```

By picking x to represent the output, you generate an "out" signal that is zero for frequencies at the center frequency, negative for those below and positive for those above -- the input function and comparator of a frequency-follower.

Best,

Bill P.

Date: Fri Jul 16, 1993 3:24 pm PST
Subject: CSGnet role play

[From Dag Forssell (930716 1310)] Rick Marken (930716.1000)

>How about this -- a CSGNet "role play"....
>
>Please be very specific....

A splendid idea, but why shove all the burden onto Greg? Rick, you have recently demonstrated an intense interest in applications of PCT. Last fall, you indicated an interest in _Freedom From Stress_ by Ed Ford. Ed or I sent you a courtesy copy, as I recall. You promised us a review at the time.

Perhaps a time has come when you can kill two or even three birds with one stone.

A) Read _Freedom_ in full and give the net a review of what you observe and how the book can be improved or extended.

B) Building on the many role plays in the book and the rationale presented for how and why they are conducted the way they are, offer the net your own suggestions about a CSGNet "role play."

C) You will get a lot of ideas for your own book (you are writing one, right?) on how to apply PCT - things to do as well as things not to do. You can't lose.

I for one look forward to your review and to the PCT canon on CSGNet etiquette.

Best, Dag

Date: Fri Jul 16, 1993 7:54 pm PST
Subject: Projected points of origin

[From Bill Powers (930716.2130 MDT)] Bruce Nevin (930716.1609) --

As I understand you, you're saying that the perception of /d/ is hypothesized to be connected with a "projected point of origin" of F2 (projected in time, backward, from the waveform?). However, it also seems that this projected point does NOT predict perception of /d/ except through comparison with projected points of origin found when other vowels follow the /d/, for a given speaker. Is this correct?

Assuming I understand: This seems to have led to the conclusion that in order to identify a /d/, it is necessary to compare a given point of origin with points of origin occurring with different vowels for the same speaker. This has evidently been done successfully with sound spectrograms; I presume that an experimenter, using this method of projected points of origin, can successfully predict when a subject will say that a /d/ has occurred.

The question I want to ask is why we should assume that the brain also uses this method in identifying a /d/. But I had better wait to see if I have understood correctly what you're saying.

Best, Bill P.

Date: Fri Jul 16, 1993 8:18 pm PST
Subject: Vocal Tract Normalization

[From Gary Cziko 920717.0400 UTC]

Bruce Nevin (Fri 930716 16:09:50 EDT) said to Rick Marken (930716. 1230):

> And it scotches the
>notion that some feature of the acoustic image [I] is always going to be
>recognized as the phoneme /I/, since on the "relaxed" occasions it is
>actually the phoneme /i/. It is not the case that [I] is ambiguous (as
>Bill suggests): on each occasion, the range of possibilities determines
>a unique and unambiguous interpretation of the acoustic image.

Here's another example of the same type of phenomenon I recently ran into in Philip Lieberman's book _Uniquely Human_ (p. 47).

". . .the word _bit_ spoken by a large adult male speaker can have the same formant frequency pattern as the word _bet_produce by a smaller male. Yet we "hear" the large person's _bit_ as _bit_ rather than as _bet_.

"Human listeners always normalize speech signals in terms of probable length of a speaker's vocal tract. Experiments using artificial speech to confuse listeners show that listeners will interpret the same acoustic signal as a different vowel, depending on whether they believe the speech is being produced by a shorter or longer vocal tract (Ladegoged and Broadbent, 1975; Neary, 1978)."

Thanks goodness we all SPELL the words more or less the same way or big guys like Bill Powers would have even more difficulty communicating with smaller guys like Greg Williams (and vice versa!).--Gary (medium-sized-and-better-figure-than-Marken) Cziko

P.S. On the subject of phonemes, Mary Powers and I make a distinction between the three words "Mary," "merry," and "marry." Does anyone know why? And are there other dialects of English which make this three-way distinction (everybody here in Heatland USA says all three words identically).

Gary Cziko

Date: Fri Jul 16, 1993 8:38 pm PST
Subject: Modeling Perception

[From Rick Marken (930716.2100)] Bruce Nevin (Fri 930716 16:09:50 EDT) --

>You began talking about perception of some feature of a single acoustic
>image in real time, say /di/. That is the changed subject.

I'll let Bill Powers' (930716.1530 MDT) comments on contrasts stand as my reply to this profoundly confusing (to me) discussion. Bill sums up with my constantly iterated point:

>It seems to me that what we really want here is just an ordinary
>input function that responds to the degree to which a particular
>input pattern is detected. Given a set of such input functions,
>each tuned to a different input pattern, one of them would
>respond more than the others for a given input, and that would be
>taken as the phoneme being

There is a perceptual side to the PCT model that is EXTREMELY important to understand; we don't spend nearly enough time discussing and explaining it (I think). It is VERY important to understand that, ultimately, a control system controls only unidimensional variables; all a control system "sees" is the size of a signal. It is the perceptual functions that determine what the control system experiences and controls. The concept of a perceptual function is apparently not an easy one to understand. Since my modem is flakey I can't do it this weekend but I think it would be good to have a few classes on perceptual modelling in PCT. This aspect of PCT, by the way, is quite compatible with much of conventional psychological work on perception. Perhaps that's why there seems to be a disproportionate number of perceptual psychologists in CSG.

The fact that we have a big disconnect on the perceptual side of the model is suggested by this statement by Bruce Nevin:

>Is it the case that the retinal image that you see as a brown chair next
>to a green sofa in one person's house appears to you in another person's
>house (same orientation, same lighting conditions, same background) to be
>a brown footstool next to a green chair? If so, then there is indeed a
>parallel.

I would really like to know how this relates to speech? Are you saying that the exact same acoustic image (time frequency pattern on the basilar membrane) sounds like /di/ when spoken by one person and /bi/ when spoken by another (the differnet people being the differnet houses and /d/ and /b/ being the chair and footstool)?

Curiouser and curiouser Rick

Date: Fri Jul 16, 1993 10:23 pm PST
Subject: Neural signals

[From Dag Forssell (930716 2310)] Rick Marken (930716.2100)

>There is a perceptual side to the PCT model that is EXTREMELY
>important to understand; we don't spend nearly enough time
>discussing and explaining it (I think). It is VERY important to
>understand that, ultimately, a control system controls only
>unidimensional variables; all a control system "sees" is the size
>of a signal.

I am working on the rubber band demo script Bill provided in early April. It was included in the closed loop on demos.

The point Rick is making is well spelled out here. Bill points out that the position of a rubber band knot against an easel or blackboard is TWO variables. Up/down and right/left. Each has its own reference signal and perceptual signal. The error signals tell you:

A) too high North
B) too low South
C) too much right East
D) too much left West

The frequency of the neural current can represent the momentary distance.

Since neural currents transmit frequency only, they can't transmit negative values. Therefore, there are actually four signals required. So, to control a two dimensional location in a plane with an ordinary coordinate systems requires not two control systems as one would think, but four.

Perhaps we should say that the position of a knot in relation to a target is FOUR variables, each with its own control system.

I expect to make these distinctions in the future when I play with rubber bands. I anticipate that making some of these points will show something about how thoroughly thought out PCT is. Thanks, Rick.

Best, Dag

Date: Sat Jul 17, 1993 10:27 am PST
Subject: Contrast model?

[From Bill Powers (930717.0530 MDT)] Bruce Nevin (930716.1609) --

I'm still worrying away at this business of how contrast is defined. The problem seems to be somewhere in the following:

>The projected points of origin for formants F1 and F2
>associated with /b/ are different from those for /d/ (before
>the same several vowels). Likewise those for /g/.

>These features (or metafeatures, features at one remove from
>any single actual acoustic image) have a relation to one

>another that we as investigators can see. They are separated
>from one another. It appears to be this separation that
>enables people to distinguish the consonants from one another.
>We as investigators can refer to this relation as contrast. I
>will leave unresolved whether "contrast" is a perception of
>language users, and whether it is a controlled perception that
>is part of language. That notion is clearly a disturbance.

A sound spectrograph is an artificial perceptual function that receives a sound waveform and from it creates a visual display. The y axis represents frequency, and the x axis represents time. In measuring a projected point of origin, I imagine that the investigator extrapolates backward by some means to a hypothetical point in time and frequency at which the formants F1 and F2 appear to have begun. This projection is geometric, not auditory, and it uses information that spans a period of time, not as it occurs.

The investigator makes several such determinations for spectrograms of consonants at the onset of various vowel sounds. In comparing these spectrograms, it is found that the projected point of origin varies as a function of the vowel. Only by comparing such sets of spectrograms is it possible to distinguish, say, the /t/ in /ti/ produced by one speaker from the /d/ in /du/ spoken by another person. The distinction is not in the projected point of origin, but in the relation to the projected points of origin with other vowels following the consonant, on other occasions, for a given speaker.

The subsequent reasoning appears to be this:

1. If the investigator can successfully predict what a given consonant will be called by a human listener, a predictor of the consonant identification has been found.

2. That predictor predicts a human perception.

3. Therefore the brain of the listener identifies the consonant by the same means that the investigator used.

4. Because the investigator uses a sound spectrogram for the analysis, the brain must also use a sound spectrogram.

5. Because the investigator makes the identification from spectrograms by projecting backward to points of origin, the brain of the listener must also have available sound spectrograms, and make similar projections of the points of origin.

6. The comparison that the investigator makes is therefore also made by the brain of the listener, with the current measures of points of origin being compared with previous measures held in memory.

7. Hence the perception of a consonant results from perceived contrasts between measures of points of origin determined from spectrograms taken on different occasions.

Schematically,

```
                       LISTENER              memory
                                            /     |
                                          /       |
            ear ----spectrogram -- features --contrast -- i.d.
          /
        /
      /
Raw sound
      \
        \
          \
           device --spectrogram -- features -- contrast --i.d.
                                         \        |
                                           \      |
                    INVESTIGATOR             records
```

Is this what is being proposed?

Best,   Bill P.

Date: Sat Jul 17, 1993 1:46 pm PST
Subject: Modeling Perception

[From Gary Cziko 930717.2130 UTC]

Rick Marken (930716.2100) said to Bruce Nevin (Fri 930716 16:09:50 EDT):

>There is a perceptual side to the PCT model that is EXTREMELY
>important to understand; we don't spend nearly enough time discussing
>and explaining it (I think). It is VERY important to understand that,
>ultimately, a control system controls only unidimensional variables;
>all a control system "sees" is the size of a signal. It is the
>perceptual functions that determine what the control system experiences
>and controls. The concept of a perceptual function is apparently not
>an easy one to understand. Since my modem is flakey I can't do it this
>weekend but I think it would be good to have a few classes on
>perceptual modelling in PCT. This aspect of PCT, by the way, is quite
>compatible with much of conventional psychological work on perception.
>Perhaps that's why there seems to be a disproportionate number of
>perceptual psychologists in CSG.

I'd certainly appreciate such a "class" since I have a very hard time seeing how the "top-down" aspect of perception fits into PCT (in which the perception functions always appeared to be diagrammed "bottom-up"). I have a suspicion that this might be Bruce Nevin's problem as well.

How do you fit expectation into perception, for example, the fact that we tend to hear what we expect to hear and don't seem to pay much attention to the details of the sounds unless there is a higher-level error of involving comprehension.

The first time I might Bill Powers in 1990 I asked him about how the bottom-up vs. top-down views of perception fit in to PCT (actually, it was still just CT then), and I remember his exact words: "It's a closed loop, man!" I hoped then that my understanding would arrive when I better understood how closed loops operated, but I still have the question.

Martin Taylor and his wife Insup have developed a model of perception called the Bilateral Cooperative (BLC) model which includes both "bottom-up" and "top-down" processes, so perhaps Martin can help me see how the BLC model fits in a PCT model of perception.--Gary

Date: Sat Jul 17, 1993 6:34 pm PST
Subject: Top-Down as Perceptual Control

[from Gary Cziko 930718.0115 UTC] Gary Cziko 930717.2130 UTC said:

>...I have a very hard time
>seeing how the "top-down" aspect of perception fits into PCT (in which the
>perception functions always appeared to be diagrammed "bottom-up"). I have
>a suspicion that this might be Bruce Nevin's problem as well.

>The first time I might Bill Powers in 1990 I asked him about how the
>bottom-up vs. top-down views of perception fit in to PCT (actually, it was
>still just CT then), and I remember his exact words: "It's a closed loop,
>man!" I hoped then that my understanding would arrive when I better
>understood how closed loops operated, but I still have the question.

Why do I have to post a question to the net before I come up with an obvious answer that makes the question seem so dumb now?

Obviously, the top-down part of the loop has to do with higher-level control systems providing reference levels for lower-level systems. So the expectations are related to the controlled perceptions of higher-levels.

So then what would be an uncontrolled perceptual variable? Don't we ALWAYS have perceptual expectations? And if the higher-level can take very "sketchy" lower-level perceptual information satisfying the higher-level controlled variable, doens't this mean that the imagination is always working to some extent?--Gary

Date: Sun Jul 18, 1993 11:12 am PST
Subject: Re: PLANTS, ETC - RKC

FROM: Bob Clark (930718.03:00 pm EDT)

My Post of (930709.0930) concerned the phrase: "every living thing is
a control system," especially the word, "EVERY."

I certainly agree that SOME "living things" include control systems. My only objection is to the word, "EVERY." If ANY "living thing" is found without ANY control system, the use of "EVERY" is refuted.

Bill Powers (930710.1045 MDT) You quote me:

RC>It is my understanding that plants, in general, are "living
RC>things." They come in many sizes and shapes from single cells
RC>(bacteria) to giant Sequoias. For convenience, consider everyday
RC>trees and shrubs, etc. I am unaware of any (physically existing)
RC>variable that such plants control.

Your comments:

WP>There's some ambiguity here concerning how we identify control
WP>systems. At the level of organization at which we look at the
WP>nervous system, there are components which are not control
WP>systems (neurons, input functions, comparators, output functions)
WP>but which, when connected properly, constitute a control system.
WP>So whatever the level of analysis, there will be identifiable
WP>subsystems which are not control systems.

WP>Also, at any level of analysis, there will be side-effects of
WP>control actions which are not themselves controlled variables.
WP>For example, one reliable outcome of visual-motor tracking is to
WP>produce lactic acid in the muscles. Another is to heat the
WP>bearing holding the control handle.

WP>A human observer looking at an organism can't tell side-effects
WP>from controlled variables without investigating whatever control
WP>systems can be discovered. Science has so far looked at the
WP>behavior of organisms without considering the difference between
WP>intended and unintended effects of actions.

Well, yes, of course, Bill. I don't see why you repeat this material -- as you know, I have been familiar with these ideas for a long time. We discussed hormonal and similar systems "back when." I guess it is useful to repeat some important items from time to time.

PLANT SENSITIVITIES.
Gary Cziko (930710.0400 UTC) Bill Powers (930710.1045 MDT)

Since I raised the subject of plant physiology, I have been looking up what is known. My sources are not very current, but I find quite a bit has been known for some years.

For example, following up on the Posts noted, I find that Phototropism has been studied for many years. In some plants, at least, the growing tip is sensitive to illumination, resulting in the differential rate of growth of cells on opposite sides of the stem. This is shown by covering the tip and finding the phototropic effect disappears. It is also reported that fluids flow both upward and downward in many plants. This could transmit chemical signals in place of the nerve systems found in animals. Here the question is the nature of the controlled variable. The effect demonstrated maximizes illumination, rather than establishing a "desired condition" to be maintained. Perhaps the necessary experiments have not been performed, or perhaps either or both the reference signal and comparator are missing.

Many plants do display different versions of phototropism, but many do not. Thus, even if phototropism is somehow shown to involve one (or more) control systems, my objection to the word "EVERY" is not refuted.

Rick Marken (930704.1400) offers his experience with a Fly-Trap, and adds a suggested sequence of events. Yes, Rick, you end up with a "closed loop." Although the over-all long term operation returns the Trap to its original condition, or nearly so, each step in the sequence is complete (or very nearly) before the next begins. But nowhere is there any action in opposition to the original disturbance of the trigger hairs! Where, or what, is the "controlled variable?"

For a control system, presence of a closed loop is a "necessary condition," but not a "sufficient condition." Is this a "control system?"

I introduced the Fly-Trap as an example of a plant showing rather remarkable action -- remarkable because plants exhibiting any action are so rare. There are others such as the Mimosa, the Sensitive Plant. Such plants also suggest possible control systems.

OTHER TERMINOLOGY
Greg Williams has reported references from biological and physiological sources that seem to use control system concepts in their own terminology. These sources are not readily available to me, perhaps Greg can select and quote an example. I am particularly interested in finding the terms used for "comparator" and "reference signal." I suspect that seeming control systems are, in fact, "balance of force" systems (like the "ball in a bowl"), or "one way systems" tending to bring the "controlled variable" to an extreme value.

The TEST.
All of these situations require application of The Test. Frequently it is not hard to imagine suitable components that could combine to form one, or more, control systems. It is necessary to apply The Test to establish whether or not a controlled variable, and therefore a control system is present.

A few weeks ago, Bill presented a list of four questions to be answered favorably before concluding a variable is a "controlled variable." Unfortunately I failed to download those questions.

I hope you, Bill, or someone, will repeat those questions for those who, like me, failed to save them.

Regards, Bob Clark

Date: Sun Jul 18, 1993 5:18 pm PST
Subject: Control plants

[From Rick Marken (930718.1800)]   Bob Clark (930718.03:00 pm EDT)

>Rick Marken (930704.1400) offers his experience with a Fly-Trap, and
>adds a suggested sequence of events.

No sequence. Plant output influences plant input AT THE SAME TIME that plant input influences plant output.

> Yes, Rick, you end up with a
>"closed loop." Although the over-all long term operation returns the
>Trap to its original condition, or nearly so, each step in the
>sequence is complete (or very nearly) before the next begins.

If this were actually true then there would, indeed, be no control; just a sequential state, S-R device. But it is hard for me to believe that the effect of the fly on the sensory output of the hair cells is completed before the output of the fly trap (caused by the sensory output) has an effect on the fly.

> But nowhere is there any action in opposition to the original disturbance
>of the trigger hairs! Where, or what, is the "controlled variable?"

The disturbance to the hair cells is the fly. Doesn't the fly trap do something (like digest the fly) that might have an influence on that disturbance?

>For a control system, presence of a closed loop is a "necessary
>condition," but not a "sufficient condition." Is this a "control system?"

The closed loop is sufficient if all effects in the loop are occuring simutaneously and the sign of the effects around the loop is negative. Yes, this is a control system simply by virtue of the closed loop. The controlled variable in a negative feedback closed loop is always the sensory input variable.

>All of these situations require application of The Test.

You are implicitly doing the test when you identify the components of the closed loop. If Fly Trap experts already know that, say, bending of the hair cells is what causes the trap to spring, and if we know that the sprung trap has an effect on the disturbance to the hair cells, then we have a first approximation to ONE of the variables controlled by the fly trap -- "hair cell bend". This is not really a controlled perception -- it is an external correlate of a controlled perception -- what we call a controlled quantity.

Best   Rick

Date: Mon Jul 19, 1993 12:20 pm PST
Subject: Role Play,Perception, Hancock Experiment

[From Rick Marken (930719.1230)]   Dag Forssell (930716 1310)

>Read _Freedom_ in full and give the net a review of what you
>observe and how the book can be improved or extended.

After the meeting (or maybe at the meeting?). I promise.

As for the role play, it could really be done quite informally. I was just curious about how we can make the PCT point without seeming to be as "rejecting" of friendly perspectives as Greg seems to think we are. Just a sample reply or two is all I want to see. For example, what is the proper way to reply when a nonPCT control theorist says "the input to a control system guides the outputs" or something equivalent? Maybe the right thing to do is just ignore it and respond only to the things we can agree on with friendly nonPCTers-- things like "It's not baseball if it's played indoors".

Dag Forssell (930716 2310)--

>Since neural current transmits frequency only, they can't transmit
>negative values.

Good point. A neural perceptual signal is not only a single number, it is also always positive.

Gary Cziko (920717.0400 UTC) --

>Here's another example of the same type of phenomenon I recently ran into
>in Philip Lieberman's book _Uniquely Human_ (p. 47).

>"Human listeners always normalize speech signals in terms of probable
>length of a speaker's vocal tract. Experiments using artificial speech to
>confuse listeners show that listeners will interpret the same acoustic
>signal as a different vowel, depending on whether they believe the speech
>is being produced by a shorter or longer vocal tract (Ladegoged and
>Broadbent, 1975; Neary, 1978)."

My guess is that this is an articulatory - phonetician's - eye - view of perceptual contrast; the same signal (like a grey square) looks different (light or dark) depending on context (whether it is on a dark or light background, respectively). I don't think the speech phenomenon above depends on beliefs about vocal tract length (unless you are convinced that speech is a "special" perception, being perceived in terms of how it is produced instead of how it is transduced); it depnds on the context of sound in which the same acoustic signal is heard. This does not mean that assumptions about how a perception was produced never influence perception -- they do, and this is an example of what you call a "top down" process in perception, I believe. But this kind of "top down" phenomenon is common across perceptual "modalities" -- it is certainly not unique to speech. For example, assumptions about the distance of objects that produce visual perceptions have an influences your perception of their size. This is the basis of the wonderful Ames room illusion where people appear to shrink and grow because they are actually walking away and towards you; but because you assume that the edges of the room are equidistant from the observing point (because the room looks rectangular) you see the change in the person's retinal size as a change in the person's actual size -- a change that is even more impressive to me than that of the same acoustical signal going from /i/ to /e/.

Gary Cziko (930718.0115 UTC) --

> Don't we ALWAYS have perceptual expectations?

Your questions about perception are far more advanced and difficult than what I wanted to discuss about the perceptual model in PCT. All I want people to undertand is that ALL perceptions in the control model are the momentary values of unidimensional variables, the perceptual signals; a perception is just a number -- a single number -- in PCT. This is true of ALL perceptions -- whether it is the perception of a color or the percetion of a political position. This is what the model says, anyway. There is a great deal of work, mainly in the neurophysiology of perception, that is consistent with this basic idea. This is the work on receptive fields. It shows that the firing rates of individual afferent neurons depend on the states of fairly complex events occuring at the sensory surface. This is the aspect of conventional studies and models of perception that is completely consistent with PCT. I think it is VERY important to understand this work and to know its current status. Hubel and Weisel won a Nobel prize for their work on receptive fields so I think this is a very respectable (and reasonable) place for PCT to interface with "conventional" behavioral science.

Tom Hancock (930715.1700) --

Here are some comments on your proposed experiment:

>I am guessing that one of the controlled variables in the task
>is the perception of an association between the 3 configurations
>(form, position, tone) with their name label.

I don't see how this could be a controlled variable. The subject has no way to affect the assocation between configuration and name. The association is established by you. All the subject can control is the name they give when a configuration is presented; is this what you mean by an association? If so, then this association can be controlled because the subject can say whatever name he or she likes when shown a configuration (or say no name at all). Actually, it's not this association that can be controlled but something about it -- like whether the name said as the "associate" matches the name that was previously experienced as the associate.

>Subjects should vary in how much their associations are distinct or vivid
>(vivid- BCP);

How do you measure the distinctness or vividness of the association? How does the subject affect this aspect of the association?

>If association is not controlled then the responses times
>should not vary systematically with each presentation of an item;

I don't understand this. What does response time have to do with the aspect of the association that is controlled? Can't subjects take as much time as they want to say a name after a configuration is presented? Are they asked to control the time between presentation of configuration and saying the name?

> The disturbance should be the presentation of a new item,
>three simultaneous configurations in relationship, for which is
>associated one correct name label. The magnitude of the
>disturbance should vary as the subject begins to bring various
>name associations under various degrees of control.

I don't understand this either. If the controlled variable is the association between configuration shown and name said then the disturbance to this variable is the different configurations that are presented, right? If the subject is trying to control an association so that X is said when A is presented and Y is said when B is presented then A and B are disturbances that are corrected by saying X and Y respectively.

> I take it that one measure of the effects of the disturbance is
>the response time to choose a name label.

I can't see why this is the case. Why is response time a measure of the effects of a disturbance? Shouldn't the effects of the disturbance be measured in terms of the controlled variable?

I don't think it's really worth it to go on until I can get a better idea of what you think a subject might be controlling in this experiment. Based on what I know of your research, let me just suggest that what you might be interested in is the subject's ability to control the relationship between a a confidence rating and an association they make between a configuration and a name. Is this right? For example, the subject is shown a configuration and gives a name. This name is either right or wrong -- something that the subject might know if he or she has previously been told the name that goes with the configuration. Anyway, the subject's response to the configuraton can either be right (R) or wrong (W). The subject is then to give a rating of his or her certainty that the response is R or W. The con- trolled variable might be some measure of the association between a binary variable -- the subjects response (R or W) -- and a continuous variable -- the subject's rating of certainly (0.0 to 1.0?). There are various statistical measures of this variable -- which I think is called a point-biserial correlation or something like that. So this is one hypothesis about what the subject is controlling -- a correlation between ratings and R-W values. The problem with this controlled variable is that part of it is imagined -- the subject's perception of whether his/her answer is R or W. Subjects can't control the relationship between their rating and whether they were "really" R or W, can they?

Let's keep working on this.

Best Rick

Date: Mon Jul 19, 1993 1:19 pm PST
Subject: Trends and casting nets

[From Tom Hancock (930719)] Tom Bourbon [930716.1650]

TB
>>Did you do a mathematical trend analysis on your data, or are you
>>saying there is a trend when .. some people had scores like the
>>ones you expected.

TH
>For my findings with higher level control, I used some typical
>exploratory data analysis techniques: just exploring data (with
>PCT glasses) and seeing what is there. So I just said there was a
>trend.

TB
>That is what I thought.

>Trends were usually identified as "linear" or "curvilinear," with
>the latter category including quadratic, cubic, and in a few brave
>texts, quartic. Almost all texts said no one could make sense out
>of anything above a quadratic relationship.

TH now:
Ah, that's what you mean by trend analysis! Yes, all the trends were linear fits. My exploratory analysis of such a study usually includes polynomial regression analysis-with linear, quadratic and cubic components. I am curious about why you originally asked-- what difference does it make?

TB
>>Not that any of this will help you, but you asked.

TH now:
Yes, it helps--my vocabulary is better now. Feel free to lend any other help you might-I am hungry for it. I only have a few more weeks where I can indulge in keeping up on the net and making more sense of my enduring research concerns--back to teaching 14 credits a term! (On to somewhere else after next year!)

By the way, thanks for introducing me to CSG back in 1990. Do you ever go to those type of conventions any more (APS)? Also, I was wondering if anyone in CSG does any netcasting of higher level PCT predictions (a la Runkel)? Is the feeling that it is kind of fruitless until the lower levels get pinned down?

Tom Hancock

Date: Mon Jul 19, 1993 4:02 pm PST
Subject: Re: Trends and casting nets

From Tom Bourbon [930719.1705] Tom Hancock (930719)]

First a comment on trends, a few thoughts about your quest for a new
experimental design.
>
>Tom Bourbon [930716.1650]

>
>TB
>>>Did you do a mathematical trend analysis on your data, or are you
>>>saying there is a trend when .. some people had scores like the
>>>ones you expected.
>
>TH
>>For my findings with higher level control, I used some typical
>>exploratory data analysis techniques: just exploring data (with
>>PCT glasses) and seeing what is there. So I just said there was a
>>trend.
>
>TB
>>That is what I thought.
>
>>Trends were usually identified as "linear" or "curvilinear," with
>>the latter category including quadratic, cubic, and in a few brave
>>texts, quartic. Almost all texts said no one could make sense out
>>of anything above a quadratic relationship.
>
>TH now:
>Ah, that's what you mean by trend analysis! Yes, all the trends were
>linear fits. My exploratory analysis of such a study usually
>includes polynomial regression analysis-with linear, quadratic
>and cubic components. I am curious about why you originally asked--
>what difference does it make?

Only that the word "trend" has a specific meaning in statistical analyses and that most often whan people use the word they ignore that specific meaning. Instead, they mean nothing more than, "it seems to me." Not that there is anything wrong with things "seeming to" someone, but trends, like significance, have different meanings in statistical analyses. I may not like the way statistics are often applied in the behavioral sciences, but I like to play by the rules when I use them.

>TH now:
..

>By the way, thanks for introducing me to CSG back in 1990. Do you
>ever go to those type of conventions any more (APS)? Also, I was
>wondering if anyone in CSG does any netcasting of higher level PCT
>predictions (a la Runkel)? Is the feeling that it is kind of fruitless
>until the lower levels get pinned down?

My pleasure! I do still carry my PCT-equipped portable computer to conventional psychology meetings, but not very often. Most people walk right past the demonstrations and most of those who stop either openly blow it off, or say, "Oh, that's [just like/nothing more than] so-and-so's work." But once in a while a Tom Hancock comes along.

As for the higher-level predictions, some of us seem to have more than we can do trying to work our way through some of the lower levels, or maybe it's in our genes that we must focus on little things. (Linear causality, anyone?

It's in the news these days.) I think a lot of work in the social and behavioral sciences *could* be turned into Runkel-like net casting -- look at what the PCTish sociologists have done. But is seems a lot more fashionable and profitable for people to search for and find nice one-way causes. Does anyone else foresee the impending collision between advocates of genetic causes, on the one hand, and of social contagion and infection, on the other? Yet another round in the dance.

I have been thinking about your discussion, primarily with Rick, about your research design. If you want to modify your research so that it more clearly includes elements from PCT, you need a controlled variable -- Rick has been saying that, I believe. If you are studying associations between arbitrary geometric patterns (is that right?) and "real" words or objects, could you make the patterns continuously variable, affected by slow random disturbances, and look to see if there are any identifiable configurations your people control as "the right ones?" This would involve you looking for reference levels of environmental variables, with the participants setting their own references for perceptions of those variables. You might even let them develop their own sets of associations, starting with the easiest or most obvious ones from their point of view, then refining the set as they became more confident or experienced. Initially, they would only defend against disturbances to the few associations they had formed; later, when they negate your disturbances to previously undefended variables, you would know they had formed new associations.

At one level, this becomes a lot like the classic testing-as-game with young children in which we probe their knowledge of [control of] associations. ADULT: "What is that?" CHILD: "A duck." ADULT: "Yes! Ducks say meow!" CHILD: -- laughter, scorn, double takes, cries of "Silly!" or "You're dumb!", or any one of host of other spontaneous indications that you screwed up an obvious association. Come to think of it, this works even with discrete variables and it seems Runkle-like [also like Robertson, Goldstein, Mermel and Musgrave], doesn't it? And now that I think of it, Hugh Petrie talked about similar kinds of "testing" -- if not that specific one -- in his chapter, "The case against objective testing," in the AAAS book on control theory, years ago.

This is probably too simple minded for your needs, but perhaps you can imagine related changes in your design.

Sorry you won't be at SCG next week.

Date: Mon Jul 19, 1993 9:58 pm PST
Subject: PCT role play

[From Dag Forssell (930719 2245)] Rick Marken (930719.1230)]

>>Read _Freedom_ in full and give the net a review of what you
>>observe and how the book can be improved or extended.
>
>After the meeting (or maybe at the meeting?). I promise.
>
>As for the role play, it could really be done quite informally. I

>was just curious about how we can make the PCT point without
>seeming to be as "rejecting" of friendly perspectives as Greg
>seems to think we are. Just a sample reply or two is all I want
>to see. For example, what is the proper way to reply when a
>nonPCT control theorist says "the input to a control system guides
>the outputs" or something equivalent? Maybe the right thing to do
>is just ignore it ......

Rick, I really don't want to ruin your reading enjoyment by telling the secret before you get to it yourself. As Benjamin Franklin said: "Rick Marken is best convinced by that which he himself discovers".

If we are going to discuss it at the meeting, it seems to me that you would want to read it before the meeting, not after.

Best, Dag

Date: Tue Jul 20, 1993 3:10 am PST
Subject: Roll play

From Greg Williams (930720) Rick Marken (930716.1000)

>I think most of the problems we have on the net are a result of the
>fact that this written medium is so damn HOT for some reason.

You're going to have to learn not to be so metaphorical and loose in your talk if you're going to last on this net. Real PCTers don't talk about the "fact" of "hotness" of written (or elsewise) media! You need to realize that "hot" is a perception, and that only living organisms have perceptions. In this case, it is your perception (apparently). Problem two: I touch my computer screen and note that the letters glowing there are not "hot" to me. So we have another problem even on the intersubjective level. I wish you would get the concepts right before you go around acting like a representative of PCT. You could easily get others to start talking about, e.g., the "factual odor" of the net and such. Please -- perceptions aren't somewhere in the air between organisms! If you need an explicit hint on PCT correctness, note that the following is how you should have said the above, to remove all traces of possible accusations of PCT heresy: "I think most of the problems we have on the net are a result of my perception that this written medium seems so damn HOT to me for some reason." It really isn't that hard to get it right, so I hope that you'll straighten up soon.

As Bobby Dylan asked, "How does it feel?"

:->> [hearty laugh] I'll be back, Greg

Date: Tue Jul 20, 1993 7:24 am PST
Subject: role of contrast

[From: Bruce Nevin (Tue 930720 10:32:02 EDT)] Tom Bourbon [930716.1333]

(Me (Fri 930716 14:10:47 EDT))
> ... there is some kind of cost to the control system, such that it won't
> >spread the four shapes to the extremes of the space between 0 and 90
> >degrees unless it is motivated to put out extra effort and incur that
> >cost.

[Tom]
> There it is again, and it is still slipping past. I mean the implication
> that the system has (must have) references for contrast, and now for level
> of effort and for cost incurred. Perhaps so. In some cases, undoubtedly so.
> But (break out your sonograms), it ain't necessarily soooo. (An allusion to
> a song some of us, and probably Alexander Graham Bell, would recognize.)

No, I think this sort of constraint is a physical characteristic of the vocal tract, on a par with the acoustic filtering function of the vocal tract, something in the environment, hence, a disturbance. (But Alexander Graham died in 1922, Gershwin's Porgy & Bess is 1935. Too bad--it was a nice image of old AGB tapping his foot!)

(Bill Powers (930716.1530 MDT)) --

> It seems to me that what we really want here is just an ordinary
> input function that responds to the degree to which a particular
> input pattern is detected. Given a set of such input functions,
> each tuned to a different input pattern, one of them would
> respond more than the others for a given input, and that would be
> taken as the phoneme being heard.

Bill, how is the input pattern determined?

I am suggesting that these input functions are calibrated in part by a function that maximizes the acoustic distinctness of these patterns from one another. The tongue and lips have rest positions to which they return unless muscular effort (or some external disturbance) moves them from it. This amounts to a kind of inertia in the production of speech. Its effect is perhaps most evident when someone speaks in their sleep or under sedation. This inertia is a disturbance to a postulated contrast function controlling the acoustic distinctness of phonemic input patterns from one another. The gain, the degree to which the function resists this disturbance, may be varied. The appearance of there being such a function may be an illusion, a byproduct of something else. Anything that would account for the effect is perfectly fine with me. I don't require that there be a "contrast detector".

You don't need a "contrast detector" to recognize words; you need it to attune your inventory of phonemic elements to a particular speaker's particular way of effecting them in behavioral and acoustic outputs. (There is no a priori reason to assume that the phonemic elements are identical across all speakers of the "same" dialect, despite claims of Generative phonology. The pronunciations certainly vary.)

An image: spots moving apart or closer as a balloon expands or contracts. This accords with your notion of averaging and subtracting the average, but see farther below.

A graph with formant 1 frequency and formant 2 frequency on the two axes defines an acoustic space for vowels called the "vowel triangle." The limits are defined by how open the mouth can be for [a], and how close it can be for [i] and [u] without getting close enough to cause air turbulance (fricative sound). Those are the cardinal vowels. Other vowels are distributed between them and in the interior of the vowel triangle. (There are actually a front and back "a" defining the lower corners on the shortest side of a vowel trapezoid.) In "fortis" or more contrastive pronunciations, vowels approximate the limits of this acoustic space; in "lenis" or less contrastive pronunciations of the same words, the same vowels (phonemically same) are more centralized in this acoustic space, with the relative distances between them more or less proportionally preserved.

I say "that these input functions are calibrated in part" by the postulated contrast function because dialect differences involve other differences. I haven't time to get into that now in detail. But if you have fixed phoneme detectors, not subject to calibration, then your proposal would fail to accommodate the differences between speakers of different dialects, or switches from one dialect to another in a given speaker, in cases where formal speech is "speaking properly" i.e. emulating the model of a prestigious dialect, and not merely articulating the sounds with greater mutual contrast.

> If the perceptual signals from the various phoneme detectors are
> averaged and the average is subtracted from each signal, what
> began as an /e/ could very well end up as an /i/ as it should.

Surely you don't mean the signal whose relative strength in different phoneme detectors tells higher level systems which phoneme is present. That would not calibrate a phoneme detector to recognize the vowel in a Valley Girl pronunciation of "good" with no lip protrusion or lip rounding as /U/ rather than /I/. Nor can this be an averaging of signals input to the phoneme detector's perceptual input function (PIF), as those are the same for all. So it must be signals at some intermediate stage within the PIFs of phoneme detectors. What are you proposing? Perhaps you have something specific in mind in terms of weighted sums.

All I am claiming is that phonemic elements of a language are defined relative to one another, and not in absolute terms. This means that phoneme detectors are calibrated in terms of the requirement that phonemes contrast, and recalibrated if necessary so as to recognize the "same" phonemes in diverse behavioral and acoustic outputs. This calibration is possible only insofar as at least some words can be recognized without recognizing all the phonemes in them. Fortunately (rather, necessarily) most of the combinatorial potential of phonemes not realized in words. And if I hear on my daughter's Disney tape of Peter Pan a Valley Girl Wendy say "give me your foot and your ..." I can hear the contrast between the vowel of "give" and the vowel of "foot" and recognize the words, although the pronunciation of both is shifted relative to my own.

Your suggestions about averaging and weighted sums would accomplish this, if you can find a way to implement them. I don't say there is a need for a contrast detector, only that phonemes be defined in relative rather than absolute terms.

> General note for modelers

Are you on the verge of building a pitch extractor? A frequency- follower that can determine the fundamental frequency of a speech signal in real time?

(Bill Powers (930716.2130 MDT)) --

> The question I want to ask is why we should assume that the brain
> also uses this method in identifying a /d/. But I had better wait
> to see if I have understood correctly what you're saying.

You understand me correctly. (Your words appear to me to be a paraphrase of what I said.)

The claim is that there is no acoustic feature in the audio signal corresponding to /du/ that is also found in the audio signal corresponding to /di/.

The question then was, on what basis does the brain identify the same /d/ in both productions?

I pointed out the projected point of origin as something that they have in common with one another (and with all other /dV/ and /Vd/ productions, where V is any vowel). Rick siezed upon this as an acoustic feature that a phoneme detector could find in the audio signal for /du/ or for /di/. I argued that this projected point of origin is not actually present in the audio signal. I argued that it could be determined by comparing various audio signals, knowing in advance that they contained a /d/ at an identified point.

However, I also said more that you should consider together with this. I hypothesized that the controlled perceptions for consonants like /d/ are predominantly kinesthetic, with recognizable acoustic byproducts, and that the controlled perceptions for vowels are predominantly acoustic. I argued that the shift in F2 upward for /di/ and downward for /du/, plus the upward shift in F1 for both, is a side effect of the tongue moving from its point of closure for /d/ to the vowel. One may recognize the presence of /d/ from its any of differential side effects on different vowels, I suppose somewhat as one recognizes the presence of a dog from a tail or a paw or a bark. Phoneme perception is generally regardes as categorial.

I understand that there is good evidence that consonants are processed together in one part of the brain and vowels are processed together in another part of the brain, though I am not familiar with the literature; such a finding at least does not contradict my hypothesis.

Does this answer your (930717.0530 MDT) question? I do not require that the brain do something like sound spectrography. Nor do I claim that contrast is input to a perceptual input device whose signal (in sequence with others) is input to a word detector. (I assume that is the meaning of i.d. in your diagram.) Having a phoneme detector is fine. Contrast comes in to the question of how the input function of a phoneme detector is defined.

It is interesting that we can converse seamlessly with groups of people concurrently speaking different dialects. The "same" phoneme detectors can interpret a given phonetic feature as one phoneme in the speech of one person and as a different phoneme in the speech of the next. Only a linguist or phonetician is likely to notice that some pair of different words are phonetically the same in the speech of the two people (but are phonetically different within the speech of either person, and phonemically contrasted in the speech of either and in the conversation as a whole). An untrained observer is much more likely to observe that "pen" and "pin" are not phonemically contrasted in dialects spoken for example near the Wendy City -- I mean, the Windy City.

(Rick Marken (930716.2100))

> Are you saying
> that the exact same acoustic image (time frequency pattern on the
> basilar membrane) sounds like /di/ when spoken by one person and
> /bi/ when spoken by another (the differnet people being the differnet
> houses and /d/ and /b/ being the chair and footstool)?

No. /d/ and /b/ are too divergent. But take the /I/ of "bid" and the /E/ of /bed/. The pronunciation of "bid" by some folks where I grew up in Florida is about the same as my present pronunciation of "bed". Their pronunciation of "bed" is something like "bayuhd", so the distinction between the words is preserved. In Gary Cziko's neighborhood, I'll bet these words have identical vowels--in some dialects distinctions are lost. I say lost, because historical reconstruction shows that ancestors of those people did have a distinction between these vowels in their speech. Processes of change and reasons for them are beginning to be understood--ref. William Labov's work, etc.

So yes, the "exact same acoustic image (time frequency pattern on the basilar membrane)" MAY sounds like different phonemes when spoken by different people, if their pronunciations of other systematically contrasted phonemes shift around proportionally.

Or even when spoken by the same person at different times, if that person switches dialects (typically because "careful" or "correct" speech involves emulation of a socially esteemed dialect). These are true facts about language that PCT must somehow model.

(Gary Cziko 930718.0115 UTC) --

> Don't we ALWAYS
> have perceptual expectations? And if the higher-level can take very
> "sketchy" lower-level perceptual information satisfying the higher-level
> controlled variable, doens't this mean that the imagination is always
> working to some extent?

I think the answers are yes and yes much much more than we realize. But what do I know?

(Rick Marken (930719.1230)) --

I think experiments were done in which synthesized speech was coupled with visual images of people, and the interpretation of a nonsense syllable depended upon the image in a way completely consistent with the pronunciation expectable from a big or small person. Gary, does Lieberman give any details? But in general I agree with you, it is the context of other pronunciations that is most relevant, and the above is an artificial experimental setup.

Enough! Or too much. (End quote.)

Bruce bn@bbn.com

Date: Tue Jul 20, 1993 8:09 am PST
Subject: Stess, Rollin' Play

[From Rick Marken (930720.0800)] Dag Forssell (930719 2245) --

>If we are going to discuss it at the meeting, it seems to me that
>you would want to read it before the meeting, not after.

I will try to read "Freedom from stress" before the meeting -- but I'm pretty stressed out at the moment getting ready for the trip and dealing with Greg's roll play (see below).

Greg Williams (930720) --

[I presume that this is the start of my requested role play. I guess it's meant to be a representation of how I respond to people on the net. If so, then I am being perceived (at least by Greg) as a lot nastier than I am trying to be. So, as a model of how Greg sees my posts, it is very disturbing. I hope, however, that he will post some models of the "right" way to respond to posts that contain material that some of us PCTers disagree with (responses that communicate our disagreement -- remember the goal is to make PCT points without alienating nonPCTers; If the only way to do this is simply to agree with everything then the role play is unnecessary -- I already know how to agree with people I don't agreed with). Anyway, here is my reply to Greg's post -- with me in the role of the person who has been attacked by "Rick the Prick".

>You need to realize that "hot" is a perception, and that only living organisms
>have perceptions.

I agree. I used the word "hot" as a description of the state of a perceptual variable.

>I wish you would get the concepts right

>before you go around acting like a representative of PCT. You could
>easily get others to start talking about, e.g., the "factual odor" of
>the net and such.

As long as people understand that the words refer to perceptions I don't think there is a problem. "Factual odor" seems like a nice, poetic reference to a perceptual variable.

>If you need an explicit hint on PCT correctness,
>note that the following is how you should have said the above, to
>remove all traces of possible accusations of PCT heresy: "I think most
>of the problems we have on the net are a result of my perception that
>this written medium seems so damn HOT to me for some reason."

If these extra words would help keep it straight then I think that's a good idea. But I think your sentence here doesn't quite capture my intended meaning. I don't think the problems (that I perceive) are a result of just my own perception of the heat of the medium. I think (imagine) that each person on the net perceives (among other things) a certain level of "heatedness" in each communication. I am guessing that the perceived level of "heatedness" is higher for the same communication (same words) when this communication is done in writing compared to when it is done in person, in talking.

>It really isn't that hard to get it right, so I hope that you'll
>straighten up soon.

I try to get it right; I try to say what will make the communication clear. If by adding that "hot" referred to the state of a perceptual variable I would have made the point clearer then I agree thatI should have added it; I just took it for granted that we can now talk about ANYTHING and know we are talking about perceptions.

>As Bobby Dylan asked, "How does it feel?"

Like a rollin' stone.

>I'll be back, Hasta la vista, baby. Love Rick

Date: Tue Jul 20, 1993 8:26 am PST
Subject: Conference attenders

[from Mary Powers 930720]

Gary Cziko:

I was in Boulder over the weekend, and just got your message.

Here is the (not necessarily final) '93 conference list. (g) for guest.

Bourbon, Tom
Clark, R.K.
(g)Clark, Mary Ann

Cziko, Gary
Dennis, Brent
Duggins-Schwartz, Michelle
Ford, Ed
Forssell, Christine
Forssell, Dag
Forssell, Karin
Good, Fred
Good, Perry
Grandes, Toto
Hershberger, Wayne
Kelly-Wilson, Lisa
Kurtzer, Isaac
Larson, Larissa
Marken, Rick
(g)Marken, Linda
McClelland, Kent
McPhail, Clark
Miller, Dan
(g)Miller, Barbara
Moe, Jeffrey
Olson, Mark
Powers, Mary
Powers, William
Robertson, Richard
(g)Robertson, Vivian
Schweingruber, David
Schweingruber, Frank
Taylor, Martin
(g)Taylor, Ina
Tucker, Charles
Williams, Greg
Williams, William

plus 3 maybes

Date: Tue Jul 20, 1993 8:58 am PST
Subject: Plant control systems

[From Bill Powers (930720.0815 MDT)]

Bob Clark (930718.1500) --

>I certainly agree that SOME "living things" include control
>systems. My only objection is to the word, "EVERY." If ANY
>"living thing" is found without ANY control system, the use of
>"EVERY" is refuted.

Agreed.

What constitutes a control system? It is a system in which

a. Physical actions are regularly related to informational inputs from the environment. An "informational" input is one that affects signals inside the system in a unidirectional way. A "signal" is a low-energy variable that can alter the states of high-energy variables (neural or chemical signal, enzyme).

b. The informational inputs depend directly and continuously on effects of the same physical actions, as well as on independent influences.

c. The gain around this loop is substantially more negative than -1 at all frequencies below some finite limit.

In order for an organism NOT to be a control system in any regard, the only existing closed-loop relationships similar to those above would have to have very low or positive loop gains.

I think it would be impossible to find any living organism in which one or more closed loops with negative gain cannot be found. So the only real question remaining is the magnitude and sign of the loop gain involved. We can rule out large positive loop gains in most cases because we don't observe the oscillations or runaway conditions that would result. The only practical choice left is between zero (approximately) and negative loop gain. With zero loop gain, the physical actions of the organism would have no effect on any informational input from the environment. With negative loop gains large enough not to be lumped with the approximately zero loop gains, the related informational inputs would be controlled to some degree, as would the environmental conditions on which they immediately depend.

>Well, yes, of course, Bill. I don't see why you repeat this
>material -- as you know, I have been familiar with these ideas
>for a long time. We discussed hormonal and similar systems
>"back when." I guess it is useful to repeat some important items from time to time.

I repeat it because you occasionally forget it, or at least fail to apply it:

>For example, following up on the Posts noted, I find that
>Phototropism has been studied for many years. In some plants,
>at least, the growing tip is sensitive to illumination,
>resulting in the differential rate of growth of cells on
>opposite sides of the stem...

>Many plants do display different versions of phototropism, but
>many do not. Thus, even if phototropism is somehow shown to
>involve one (or more) control systems, my objection to the word
>"EVERY" is not refuted.

It is not logical to use the lack of one type of control system to "prove" that some plants contain no control systems at all.

All plants are sensitive to external conditions, including variously the direction of gravity, local moisture content, CO_2 concentration, and dozens of other variables. The traditional interpretation of these sensitivities is open-loop: the plant, it is said, simply "responds" to these conditions.

However, the response almost always alters the plant's relationship to the external causes of the variables, or alters the variables themselves (as plant respiration both depends upon and alters -- and controls -- local CO_2 and O_2 concentrations).

>For a control system, presence of a closed loop is a "necessary
>condition," but not a "sufficient condition."

The "sufficient condition" is completed by specifying that the loop gain is greater than 1 and negative.

In the case of the Venus Flytrap, I am not convinced that the trigger response is a closed loop in itself, because the immediate effect on the touch comes too late to affect the trigger response. I would view this response as part of an output function. The output variable in this case, I surmise, is the sensitivity of the trigger response itself. If this is part of a control loop aimed at controlling the nutritional state of the plant, we would expect the ingestion of flies to lower the sensitivity of the trigger response. If an unlimited number of flies were confined near the flytrap, I would expect the sensitivity to become very much less or even zero, preventing an excess accumulation of corpses being digested. Under traditional interpretations, we would call this "satiation" or "adaptation" (because the negative feedback effect has to be explained somehow).

Plants operate on a much slower time-scale than animals. We should think of the flytrap's output not in terms of individual triggering events, but, perhaps, in terms of triggerings per hour, or per day, and the input in corresponding terms: flies per hour, or per day, or simply concentration of nutrients derived from ingestions of flies at some rate. I'm not in a position to state the loop gain of this control system, but I'll wager that it is both negative and large.

>I suspect that seeming control systems are, in fact, "balance
>of force" systems (like the "ball in a bowl"), or "one way
>systems" tending to bring the "controlled variable" to an extreme value.

I couldn't disagree more. I don't believe that any such systems exist in organisms except as components of control systems: either balance-of-forces or extremum systems. You're offering the traditional concepts developed before anyone knew of control systems. All phenomena explained in such traditional ways need serious re-investigation.

>A few weeks ago, Bill presented a list of four questions to be
>answered favorably before concluding a variable is a
>"controlled variable." Unfortunately I failed to download those questions.

>I hope you, Bill, or someone, will repeat those questions for
>those who, like me, failed to save them.

Here they are:

RE: the controlled variable:

1. It must be affected by an output of the behaving system, and also by independent disturbing variables. A test for control requires the presence or application of known disturbances.

2. The effects of the output must remain nearly equal and opposite to the effects of disturbing variables, so that the putative controlled variable is maintained near some specific value.

3. If the output is prevented from affecting the controlled variable, that variable must change according to the effect of disturbances on it.

4. If the control system is prevented from sensing the state of the controlled variable, the effects of unpredictable disturbances on the variable must no longer be canceled by the variations in output of the behaving system.

Greg Williams (930720) --

>As Bobby Dylan asked, "How does it feel?"

Or as Eric Berne put it, "Now I've got you, you son of a bitch!" (p. 85)

Best, Bill P.

Date: Tue Jul 20, 1993 9:06 am PST
Subject: friendly comment

I am writing to make a comment about a thread of conversation on this net which has disturbed me. Although I am only a semiregular reader of the bulletin board, one aspect which I have noticed is the strongly worded rejections by some PCTers of some of the claims made by sociologists about human behavior, especially those involving "problem solving" as it is characterized by George Herbert Mead. I recall objections to the idea that people can "give themselves instructions" and that people can "change their reference signals." I believe that the version of PCT presented in BCP clearly allows for both of these, which are descriptions of creating reference signals by conscious thought.

BCP, I think, is such a remarkable book because it manages to "soak up" such a huge amount of data. Powers proposes a model to explain not just behavior, but memory, learning, cognition, consciousness and volition, among other things. He also makes clear that the model is only as good as its ability to explain the data. Some of the discussion on the net involves throwing out data which doesn't fit a particular interpretation of the model. (I suppose this is inevitable since limiting what is relevant data is one of the attributes of scientific paradigms or other system concepts.) I say a particular interpretation because clearly PCT a la BCP allows for things which are dismissed by some PCTers. Specifically, it allows for (1) thinking and (2) volition, one or both of which allows me to give myself an instruction and/or change a reference signal.

(1) In the imagination mode (p.223), I can try out a number of reference signals in my head before "flipping the memory switch." The example is the

chess game. Because the chess player has a reference signal of winning the game but has not won it, he is in a state of disturbance and therefore can generate a wide variety of reference signals, some of them quite creative, through thought. Similarly, since all of us are in a constant state of disturbance from birth till death, we can test various reference signals by imagination. Each time we flip the memory switch, we create new error which leads to new behavior. To say that people only behave because of disturbances and not because of new reference signals is inconsistent with BCP as I understand it. Again, because I am always in error, I can always create new and creative reference signals through thought.

(2) In designing his reorganizing system, Powers attempted not just to explain how new control systems can be created but to account for the phenomena--real data--of consciousness, awareness and volition. According to BCP (p.197-201), when I am conscious of something my reorganizing system is plugged into it and can reorganize the control systems which create and/or control the perception I am focusing on. Also, this system can stick reference signals anywhere in the hierarchy, i.e. volition.

"And we have given the system as a whole the ability to produce spontaneous acts apparently unrelated to external events or control considerations: truly arbitrary but still organized acts" (p.199).

Between the imagination mode and the reorganizing system a la BCP there is room for all the sorts of problem solving activities which were of interest to Mead and his followers. One of Mead's shortcomings, as I read him, is his concentration on problematic situations which require thought and consciousness. He doesn't provide a framework for behavior by already existing control systems to say nothing of all the many other phenomena which PCT can explain. The model presented in BCP is able to soak up so much data there is no reason why its adherents should want to dismiss any of it, let alone phenomena as important as thought and volition.

I think part of the theoretical narrowness may be related to methodological narrowness, specifically the idea that computer modeling is the best way to do research regardless of the problem being pursued. I have no objection at all to computer modeling or simulations. The paper which I hope to present in Colorado regards using the CROWD program to simulate Clark McPhail and Ron Wohlstein's collective locomotion experiment.

I have also spent the past year doing field research in a homeless shelter and am interested in the question: how do people with different reference signals in this or any other organization attempt to control their perceptions in an environment filled with other living control systems who create disturbances and provide resources. I consider this question important and I'm not going to stop my research until someone comes up with a way to model what's going on there. In fact, when someone is ready to model it, they ought to have the help of someone who has done the fieldwork. I hope my attempts to use PCT to understand what goes on in the homeless shelter will be evaluated on their own merits and not dismissed because some people who have applied or misapplied PCT to problems before the modelers got there got it wrong and upset people on the net.

I hope my comments will be viewed as the humble suggestions they are meant to be.

Dave Schweingruber dsg1072@uxa.cso.uiuc.edu

Date: Tue Jul 20, 1993 9:10 am PST
Subject: Re: role of contrast

From Tom Bourbon [930729.1027] Bruce Nevin (Tue 930720 10:32:02 EDT)
>
>(Tom Bourbon [930716.1333]) --
>
>(Me (Fri 930716 14:10:47 EDT))
>> ... there is some kind of cost to the control system, such that it won't
>> >spread the four shapes to the extremes of the space between 0 and 90
>> >degrees unless it is motivated to put out extra effort and incur that
>> >cost.
>
>[Tom]
>> There it is again, and it is still slipping past. I mean the implication
>> that the system has (must have) references for contrast, and now for level
>> of effort and for cost incurred. Perhaps so. In some cases, undoubtedly
so.
>> But (break out your sonograms), it ain't necessarily soooo. (An allusion
to
>> a song some of us, and probably Alexander Graham Bell, would recognize.)
>
>No, I think this sort of constraint is a physical characteristic of the
>vocal tract, on a par with the acoustic filtering function of the vocal
>tract, something in the environment, hence, a disturbance. (But Alexander
>Graham died in 1922, Gershwin's Porgy & Bess is 1935. Too bad--it was a
>nice image of old AGB tapping his foot!)

We seem to have gone full circle in discussing the origins of contrasts. I recall posting something back in late May or early June, just before leaving on a working vacation, in which I suggested that many of the distinct environmental consequences maintained by behavior say more about environmental constraints than about "control of contrasts" by the behaving system. At that time, I thought you were suggesting the opposite, but that was my interpretation of your posts and may not have been what you intended.

In your posts of the past few days, I think you have moved from what appeared to be a strong assertion that, if PCT is to provide a model for speech recognition and production, then it must as a prerequisite explain the control of contrasts, with contrasts as the objects of control. Now you seem to be saying (a) that many contrasts might result from environmental constraints and (b) that the degree of contrast in consequences of actions *sometmes* occurs as an unintended consequence (a "side effect") when a speaker controls for some other, probably higher-level, perception (such as seeing another person's actions change as a result of what the speaker says). Of course, it might be that you were saying that all along and my understanding has finally caught up with you. Either way, those more recent ideas (if I understand you correctly) seem fairly close to the "radical PCT view" some of us were expressing earlier in this round of discussion on speech.

As for AGB, I know he wasn't around when Gershwin produced Porgy and Bess. Neither was I. But, with you, I liked the image -- I even had him humming a few bars into the mouthpiece while experts predicted his device was only a toy and would never last.

>(Bill Powers (930716.1530 MDT)) --
>
>> It seems to me that what we really want here is just an ordinary
>> input function that responds to the degree to which a particular
>> input pattern is detected. Given a set of such input functions,
>> each tuned to a different input pattern, one of them would
>> respond more than the others for a given input, and that would be
>> taken as the phoneme being heard.
>
>Bill, how is the input pattern determined?
>
>I am suggesting that these input functions are calibrated in part by a function
>that maximizes the acoustic distinctness of these patterns from one another.
>The tongue and lips have rest positions to which they return unless
>muscular effort (or some external disturbance) moves them from it. This
>amounts to a kind of inertia in the production of speech. Its effect is
>perhaps most evident when someone speaks in their sleep or under sedation.
>This inertia is a disturbance to a postulated contrast function controlling
>the acoustic distinctness of phonemic input patterns from one another.
>The gain, the degree to which the function resists this disturbance, may
>be varied. The appearance of there being such a function may be an illusion,
>a byproduct of something else. Anything that would account for the
>effect is perfectly fine with me. I don't require that there be a
>"contrast detector".

Bruce, it seems that the "inertia" you describe is a disturbance to the production of speech sounds, period, not just to the production of contrast. It *certainly* is a disturbance to the higher-level perceptions the speaker controls by way of speech. Similar to the exaggerated effectiveness of inertial disturbance to speech during sedation is the disturbance to control using the hands, when the arm is "asleep." Control in general is disturbed, not just control of contrasts between various gestures or manipulations. I think you say, or imply, something like that later in your post, which I will not reproduce here.

Until later, Tom Bourbon

Date: Tue Jul 20, 1993 10:09 am PST
Subject: PCT applications

[From Dag Forssell (930720 1000) Rick Marken (930720.0800)

>I will try to read "Freedom from stress" before the meeting -- but

>I'm pretty stressed out at the moment getting ready for the trip
>and dealing with Greg's roll play (see below).

I enjoyed Greg's hearty laugh, too. I thought he was poking fun at you. I did not see it as the beginning of any role playing at all.

For my part, I sincerely believe that "Freedom" offers you well thought out suggestions on how to do *any* PCT role play *and* insight into how to deal with nonPCTers. My suggestion remains to table roll play with Greg and prepare for Durango by reading "Freedom" in sequence from cover to cover.

Best, Dag

Date: Tue Jul 20, 1993 10:51 am PST
Subject: Greg W., Higher Levels with Rick M.

[From Tom Hancock (910793.1100)

Greg Williams:
Hi! I believe your disturbances on the net may have a positive impact on the total scientific yield from the CSG--good work. However, in Rick's interactions with me not many of the behaviors you have refered to have been manifested. On the contrary!

Rick Marken (930719.1230) or anyone else interested:

Thanks again for your help.

TH
>>I am guessing that one of the controlled variables in the task
>>is the perception of an association between the 3 configurations
>>(form, position, tone) with their name label.

RM
>I don't see how this could be a controlled variable. The subject has
>no way to affect the assocation between configuration and name.
>The association is established by you. All the subject can control
>is the name they give when a configuration is presented; is this
>what you mean by an association?

Actually, I was thinking that the p is the association provided by me but that the p*, the subject's perception, is not the same. My experience is that some subjects are at least partially in a passive mode and they do not see a distinct, well-formed association (typically they short out on the tonal component) while other subjects are repeatedly scanning the elements of the configuration and are also executing programs for accessing previous category associates or mnemonics so that they will perceive a distinct association. So in this sense isn't this a controlled variable?

RM

>How do you measure the distinctness or vividness of the
>association? How does the subject affect this aspect of the association?

Re the distinctness or vividness of the association: I am assuming that recall is facilitated when the perceptions are distinct (this assumption is supported by cognitive psychology, Power's description of memory and vividness, the ancient Greek understanding, subjects' self-report, etc.). So when there is successful recall (which is rapid) I assume that there are distinct associations: perception of the associations with the name label which are both complete and also differentiated from any similar configurations. Yes, the subject has to identify the name label, but in doing this there must be the inhibition of any signals from other memories of associations which are similar to the one at hand.

Regarding measuring the distinctness: I have used self-reports and I have used response times as a means of getting at it.

RM
>Why is response time a measure of the effects of a disturbance?
>Shouldn't the effects of the disturbance be measured in terms of
>the controlled variable?

This gets to one of my central concerns. What does PCT say about differences in response times (within subjects). I have assumed that a faster response time is from a well-formed control system--good control--good understanding. In the case mentioned above the subject quickly inhibits any competing assocations or name labels, there is little persisting error signal for the subject's sense of the correct name label (indicated by a high certainty rating), and thus the response time is at a minimum. On the other hand my data indicate that when a subject is less certain and tends to make the incorrect choice then the response time is much longer. One tradition that has influenced my view is the ACT* model of J. R. Anderson--longer response times are indicative of multiple inputs (in my case from the memory and imagination connection) which fire without efficient inhibition of inaccurate inputs, while quicker response times are indicative of a subject who can clearly articulate how the inputs relate and differ.

RM
>I don't understand this. What does response time have to do with
>the aspect of the association that is controlled? Can't subjects
>take as much time as they want to say a name after a configuration
>is presented? Are they asked to control the time between
>presentation of configuration and saying the name?

Yes, the subjects can take as long as they want, but I assume that they generally do it as fast as they can, which speed varies depending on how well controlled the perception is, and also varies according to the higher level control of why the S thinks he is doing the task.

No, they are not asked to control time (except in an in-progress study where I tried to vary the time setting with instructions and then validated with a post-experiment questionnaire).

TH
>>The disturbance should be the presentation of a new item,
>>three simultaneous configurations in relationship, for which is
>>associated one correct name label. The magnitude of the
>>disturbance should vary as the subject begins to bring various
>>name associations under various degrees of control.

RM
>I don't understand this either. If the controlled variable is the
>association between configuration shown and name said then
>the disturbance to this variable is the different configurations
>that are presented, right? If the subject is trying to control
>an association so that X is said when A is presented and Y
>is said when B is presented then A and B are disturbances that
>are corrected by saying X and Y respectively.

As I tried to explain above, it is not simply a matter of saying the correct name label, the subject must create a distinct perception of the correct match, thus with each presentation of a new item on the screen there is a disturbance--the subject does not control what is presented but the presentation stimulates a change in the control for association-label match.

And related to the rest of the post, I have run out of time right now, but I intend to get to it as soon as I can--it deserves careful thought. In brief, it doesn't seem that the S controls confidence rating. That rating is simply indicative of his attempt to control association with the name label and match to the reference from previous perceptions (or imaginations). Perhaps, by sending this now I will get another post that will help with a more comprehensive understanding of your last paragraph. By the way, the part of my post (930715.1700) that you did not get to on the higher level control is the most interesting and clear-cut to me.

Tom Hancock

Date: Tue Jul 20, 1993 12:51 pm PST
Subject: Tom on contrast

[From: Bruce Nevin (Tue 930720 12:37:31 and 15:27:36 EDT)]

(Tom Bourbon [930729.1027]) --

> In your posts of the past few days, I think you have moved from what
> appeared to be a strong assertion that, if PCT is to provide a model for
> speech recognition and production, then it must as a prerequisite explain
> the control of contrasts, with contrasts as the objects of control. Now you
> seem to be saying (a) that many contrasts might result from environmental
> constraints and (b) that the degree of contrast in consequences of actions
> *sometmes* occurs as an unintended consequence (a "side effect") when a
> speaker controls for some other, probably higher-level, perception (such as
> seeing another person's actions change as a result of what the speaker says).

I think I started with something like the claim that what is "phonemic" about phonemes is not their sounds or phonetic features but their systemic relations of contrast to one another. I realize that my words may appear shifting and confused as I try to grapple with a problem to which I don't yet know the solution. There is also a problem of superposed points of view.

Your statement (a) above should be changed to refer like (b) to degree of contrastiveness rather than contrast itself. Your statement (b) overleaps too much. Surely, the distal perception of another's person's actions is relevant for choice words and syntax as correlated with meanings desired to be perceived by the other person. But phonemes in themselves have no correlation with meanings; they only serve to distinguish words from one another. The phonemic shape of words is in general arbitrary. Even onomatopeoia is conventional in culture-specific and language-specific ways, e.g. the representation of crowing or barking in words in different languages. So the control of phonemes has no direct connection to the perception of other people's actions, judged to be consequences of their understanding one's words. The consequences follow from their recognizing the words and the word-dependencies, and associating meanings with the words and word-dependencies. The phonemes have no role beyond the recognition of words. Frequently they are not sufficient by themselves for word-recognition, or in some cases they are not even necessary, being supplemented or supplanted by word dependencies (expectancies) and other modes of perception (meaning).

The reason this is so difficult for me to grapple with and talk about in a consistent way is the superposition of different perspectives on language. I have identified these as control of language in adults, acquisition or learning of that control in children, and the evolution of language in human prehistory. Discussions frequently shift unwittingly between the first two especially.

For an infant, there are no phonemes. Then there are intonation contours (assertion, exclamation, question) and a few short words, but still no phonemes. Then means of distinguishing words from one another begin to come clear to the child as something that can be controlled independently of the meanings of the words, and an early and rudimentary version of the phonemic system is there. Successive reorganizations refine this, reorganizations presumably driven by failure at the higher level of words and their meanings. As the phonemic system becomes more elaborate in its cross-connectedness among the phonemes, it takes more error (more persistent/recurrent and more consistent) for reorganization to happen. That's my picture of the process.

The child has to recognize the kinds of variants I have been discussing, where the same acoustic image constitutes a different phoneme, and where different acoustic images constitute the same phoneme. Until recently, my now 6-year-old would say "datty" when she was trying to speak especially clearly (my interpretation of her motivation, of course). This was her then current resolution of the fact that /d/ and /t/ are indistinguishable in "lax" speech when they occur between vowels--remember "it edited it" and "the latter ladder".

The problem facing an adult hearing a novel dialect is one with which she has had lifelong practice. The solution to the problem begins with recognizing at least some words, by any means that work, including but not limited to recognition of some phonemes. A word is recognized containing a pronunciation of phoneme X that would be perceived as phoneme Y, except that there is no similar word with Y in place of X, or if there is then that other word (with radically different meanings) could not possibly be what the speaker intended there. There must, it seems to me, be a process of recalibrating what constitutes not only an X phoneme (from this speaker), but also a Y phoneme. The range within acoustic and perceptual space that constitutes an X is extended. But that which constitutes a Y very probably is not merely reduced, but shifted, that is, extended in some other "direction" in that space, and so on with other sounds neighboring X, and those neighboring Y, and so on. If the person pronounces /t/ dentally, with the tongue tip behind the teeth rather than farther back on the alveolar ridge, then perhaps there are changes in the pronunciations that constitute a "ch" sound. And we expect /d/ and /j/ to be shifted in a parallel way. This is not just the linguist's analysis, a lay person mimicking such an accent does indeed generalize in these ways. (This supports an argument that "distinctive features" such as that common to /t d n/ are the appropriate representation of phonemic distinctions, rather than segments /t/, /d/, and /n/.) So it does seem to me, Tom, that there must be some control of perception of the relations between phonemes in acoustic and perceptual space.

One could argue that this sort of thing is a byproduct of controlling individual phonemes like /d/ and /j/, using averaging or weighted sums in some way such as Bill has hinted. I would like to see that argument fleshed out. But note that any function or process that averages or weights perceptual signals across phonemes so as to achieve a normalization defines the phonemes in relative terms, which is all that I want. And there must be sufficient difference between phonemes that they can be distinguished from one another. Together, these amount to contrast. Or so it seems to me.

> Bruce, it seems that the "inertia" you describe is a disturbance to the > production of speech sounds, period, not just to the production of contrast.

A disturbance to the distinctiveness of all the phonemes from one another does not (except in the limiting case of indistinctness) disturb control of any individual phoneme, so long as they can still be distinguished from one another. This is the difference between speech sounds and phonemes. Phonetics deals with discriminable sounds; phonemes concern just those differences that make a difference.

I have the feeling that this is not helpful to you, but I have to quit for now.

Bruce bn@bbn.com

Date: Tue Jul 20, 1993 12:58 pm PST
Subject: Perception, Talkin' PCT

[From Rick Marken (930720.1300)] Bruce Nevin (Tue 930720 10:32:02 EDT)

>The pronunciation of "bid" by some folks where I grew

>up in Florida is about the same as my present pronunciation of "bed".

>So yes, the "exact same acoustic image (time frequency pattern on the
>basilar membrane)" MAY sounds like different phonemes when spoken by
>different people, if their pronunciations of other systematically contrasted
>phonemes shift around proportionally.

All I'm saying is that this phenomenon is not restricted to speech. You are saying that the exact same acoustical signal is heard as a different phoneme when spoken IN THE CONTEXT of "pronunciations of other systematically contrasted phonemes". An easy to demonstrate visual analog of this same phenomenon is "brightness constancy" where the exact same amount of reflected light (the analog of an acoustic signal) can look different in different illumination conditions (the analog of the different speakers). For example, let's say that my desk is currently reflecting 10 units of light and the paper on my desk is reflecting 100 units. So the desk "signal" is 10, the paper signal is 100. The desk looks "dark" and the paper looks "light". Now I go and turn off the light (there is still a little light coming in the window). Now the signal from the desk is 1 and the signal from the paper is 10. But the paper still looks "light" and the desk still looks "dark". So the same visual signal (10 units of light) looks "dark" when the "speaker" (the ambient level of illumination) has a large "vocal tract" (high level of light) and "light" light when the "speaker" has a small vocal tract (low level of light). Doesn't this seem like a fair analogy? The point is that the visual system is able to produce a different perception (dark, light) from the same input signal (10 units) without any reference to "articulatory contrasts". I think the auditory system handles this in the same way -- the input functions, in both cases, produce perceptual signals that depend on relationships (spatial in vision, temporal in audition), not absolute values.

Dave Schweingruber says:

>I am writing to make a comment about a thread of conversation on this
>net which has disturbed me.

>I recall objections to the idea that people can "give themselves instructions" and
>that people can "change their reference signals."

Are you sure? I recall some objections to the idea that "clear instructions" are an important sine qua non for PCT experiments but not to the idea that people can (if they want) give themselves instructions. Nor do I recall objections to the idea that people can change reference signals (though I would say it differently -- in PCT, the outputs of higher level control systems within people vary the VALUE of the reference signal for the lower level control systems).

>Some of the discussion on the net involves throwing out data which
>doesn't fit a particular interpretation of the model.

If this were done I don't think it would be taken lightly by many of the PCTers I know. If data really doesn't fit the model then that would be a VERY important thing to know about. Could you give one example of data that was thrown out becuase it did not fit the PCT model?

Dag Forssell (930720 1000) --

>For my part, I sincerely believe that "Freedom" offers you well
>thought out suggestions on how to do *any* PCT role play *and*
>insight into how to deal with nonPCTers. My suggestion remains to
>table roll play with Greg and prepare for Durango by reading
>"Freedom" in sequence from cover to cover.

I promise to read the book (I've already read about 1/3). But I can't stand the suspense -- what insight can you give me about how to deal with non-PCTers?

Without having read the book, here is my unsolicited take on what is going on in the interaction between PCTers and nonPCTers. When a PCTer makes a correct PCT statement (with explanation) it is very likely to have a disturbing influence on some high level perception that is being controlled by the friendly non-PCTer. Since the non-PCTer can see the cause of the disturbance (Bill, Tom or myself) as well as experience the effect thereof, he or she is likely to conclude that we are being contrary, annoying and difficult -- after all, we are making it necessary for them to defend their perception, which ordinarily stays pretty much near it's reference. I don't see any way to present PCT honestly and not have people occasionally get upset at the messenger because one or another of the PCT messages is bound to be a distrubance, and a big one, to SOME cherished belief. No matter how nicely the messanger brings it, if it is a disturbing message the messanger gets in trouble. There might be ways to try to bring the message gently to the attention of the nonPCTer -- but once the essential disturbing point gets across then there will be resistence and, possibly, anger at the messanger. I really don't see any way out of this -- except by posting anonymosuly (there would still be the disturbance to the cherished perception but the messanger would be off the hook).

The only way people will eventually "come around" (if they ever do) is through reorganization -- which means abandoning some very precious and successful control systems. I'm not surprised that there are few people willing to do it. Most people want to incorporate PCT into what they already know, which is VERY understandable. Greg thinks this is the right way to go -- let nonPCTers incorporate PCT into what they already know. The problem is that, in order to do that, the nonPCTer has to keep some old control systems in effect -- which means that old misconceptions will be defended (these old misconceptions are probably what Greg calls "meta issues"). I guess I'd rather have someone get PCT fundementally right; but this means I'd like to see a person reorganize. But what I really want to see is the RESULTS of reorganization -- the reorganization process itself can be quite painful to watch.

Anyway, I can't see a way to NOT be a disturbance to nonPCTers other than by compromising the message of PCT. I imagine that you (Dag) would suggest first trying to determine the wants a person has that are currently not satisfied and then presenting PCT as a "good disturbance" that moves the wanted perception to the reference level. But I think this nice-sounding idea has problems. The reason I think this is that there have been numerous examples of people who, on there own, found things about PCT that satisfied their wants -- we didn't even have to TRY to make a good impression with Carver and Scheier, for example. The problem is that the wants satisfied by PCT usually make it possible to adopt only the language of PCT, but not the model. Carver and Scheier liked to talk about goals and self-regulation but they never bothered to learn how a control system works. So now we have these pro-PCTers, who got a message about PCT that satisfied their want, and who don't really want PCT.

The problem for the PCTer trying to present the PCT model honestly is not one of finding wants that are satisfied by PCT; the problem is the existence of all those other wants that make certain aspects of the PCT model anathema: the want to avoid mathematics, the want to use familiar statistical tools, the want to believe in the importance of the contributions of famous person X , the want to be respected by one's colleagues, the want to "explain it all, now", the want to use familiar methodological tools, the want to stay employed, the want to publish, the want to show how much you already know, the want to believe that one's last XX years of education in the behavioral sciences were not a waste, the want to believe that everyone who went before couldn't have been fundementally wrong, the want to fit in, the want to believe that theory Y must have SOME value, the want to be seen as a great genius, the want to be seen as an expert, the want to teach, the want to show that you already know that, the want to impress your girl (boy) friend, the want to seem clever, the want to be rich,

the want to "run before you walk", the want to have a good time, etc.

It's hard to think of anything any PCTer could say that would not be a MAJOR disturbance to one or another of the perceptions that correspond to one of those wants.

Best Rick

Date: Tue Jul 20, 1993 3:24 pm PST
Subject: PCT talk

[From Dag Forssell (930720 1500)] Rick Marken (930720.1300)

>I imagine that you (Dag) would suggest first trying to
a)
>determine the wants a person has that are currently not satisfied
b)
>and then presenting PCT as a "good disturbance" that moves the
>wanted perception to the reference level.

With a) you are talking PCT, and that is always a good idea. With b), I can't figure out where you are coming from or what your terminology means.

>I promise to read the book (I've already read about 1/3).

Great! But I am not going to post "Freedom" so you can read it on your office computer in order to reduce the suspense. Relax and enjoy it. I look forward to some fruitful discussions, where we have a common frame of reference.

Best, Dag

Date: Tue Jul 20, 1993 4:06 pm PST
Subject: Phoneme detectors; Instructions

[From Bill Powers (930720.1550 MDT)] Bruce Nevin (930720.1032) --

>Nor can this be an averaging of signals input to the phoneme
>detector's perceptual input function (PIF), as those are the
>same for all. So it must be signals at some intermediate stage
>within the PIFs of phoneme detectors. What are you proposing?
>Perhaps you have something specific in mind in terms of
>weighted sums.

At the moment, what I'm exploring is a set of input functions consisting of 6 to 12 tuned filters, broadly enough tuned so they overlap while covering the frequencies between about 50 and 4000 Hz. This provides a set of signals having magnitudes that vary with time as the audio signal goes by. For each sample, the outputs of all the filters are added together, the running average is compared with a reference magnitude, and the error signal varies the gain applied to the original input signal. So this control system maintains the SUM of the filter outputs at a more or less constant amplitude. If the output of one filter rises, all the others are depressed, so the sum is kept constant. This is the normalization step.

There will be multiple input functions at the next level. Each input function will receive all 6 to 12 filter output signals, weighting them with some pattern of weights and adding the results together to provide a perceptual signal. I don't know how many of these second-level input functions will be needed. Each input function will (ideally) define one axis in a multidimensional space.

Your description of the "vowel triangle" suggests a space with only two dimensions, employing two "formants". I am unsure about just what constitutes a formant in linguistic terms. As the mouth changes shape, various resonances appear, but they are not at constant frequencies nor are there always the same number of them. When diphthongs appear, curved lines appear on the spectrogram, sweeping though at least an octave. This path creates the visual appearance of some single thing changing -- is this what is meant by a formant? Or are formants fixed bands of frequencies, counted as such whether or not any sound energy appears in those bands?

I am dealing here with fixed bands of frequencies. Each second- level input function will assign a different pattern of positive and negative weights to all the filter outputs, so that as the mouth resonances change each second-order perceptual signal will vary in a different way. This will define a moving point in a space with as many dimensions as there are second-level input functions. The immediate problem will be to find a way of assigning weights that will make the distance between points in this hyperspace as large as possible for the different vowel sounds. I would like some way of doing this automatically, but so far it seems that something has to know what vowel is present at the input, so as with perceptrons, setting up this system will require a teacher. The nicest approach would be to use a reorganizing system to vary the weightings while some monitor system minimizes the proximity of each cluster of positions from all other clusters. I hope this becomes possible to do. This would probably correspond with your hypothetical "contrast function." My first attempts will depend on knowing the vowel that is present.

With this approach, the vowel triangle using only two formants would be a special case involving only two dimensions. Perhaps whoever came up with this triangle tried multiple frequencies and found that two were sufficient. But it's possible that by using all the frequencies, greater discrimination can be achieved.

>But if you have fixed phoneme detectors, not subject to
>calibration, then your proposal would fail to accommodate the
>differences between speakers of different dialects, or switches
>from one dialect to another in a given speaker, in cases where
>formal speech is "speaking properly" i.e. emulating the model
>of a prestigious dialect, and not merely articulating the
>sounds with greater mutual contrast.

An initial calibration will clearly be necessary. However, once this calibration exists, it may not be necessary to change it for different speakers of the same dialect. The patterns of weighting will establish "surfaces of indifference" in the n-dimensional space -- for each input function, such surfaces would define the different combinations of intensities that are heard as the same vowel. The output of each second-level perceptual function would indicate only movements normal to these surfaces.

A large dialect change might require recalibration. We'll just have to see how it comes out.

>> General note for modelers

>Are you on the verge of building a pitch extractor? A
>frequency- follower that can determine the fundamental
>frequency of a speech signal in real time?

Well, "real time" is a relative term when you're doing software simulations of hardware. But yes, this looks quite possible now, although I'm not pursuing that at the moment. A tracking filter is not hard to set up. If you initialize the filter so it starts at zero frequency, it will increase its frequency until it picks up the first harmonic with appreciable energy in it and locks onto it. I don't know how practical it would be with real speech.

Another approach I've briefly tried is to arrange the filter frequencies logarithmically. This means that differences between the outputs of various pairs of filters represent ratios of frequencies. A pattern of ratios can remain constant if all the frequencies change by a constant factor -- a simple algebraic translation up or down the spectrum will remove such differences. That could achieve the same result as a tracking filter, but more simply.

Also, it seems possible to take advantage of the fact that ALL frequencies are well-separated harmonics of the fundamental voice frequency. A tracking filter set to a high frequency would be excited by all lower harmonics. This would lead to a way of accurately perceiving the basic voice frequency, information that is NOT contained in the formants. Obviously we need a way to detect voice frequency per se in order to recognize and produce things like melodies. Also, there must be some explanation for why we can perceive octaves as having a special relationship. A filter that responds to multiple harmonics of a fundamental frequency would tend to restore itself to the same octave range once the frequency had changed by more than an octave. But of course we ALSO have to be able to perceive C2 from C6. The auditory system obviously has many kinds of input functions paying attention to different aspects of the raw audio input.

>> The question I want to ask is why we should assume that the
>>brain also uses this method in identifying a /d/. But I had
>>better wait to see if I have understood correctly what you're
>>saying.

>You understand me correctly. (Your words appear to me to be a
>paraphrase of what I said.)

OK. I guess what I'll be trying to show is that there is more than one way in which the invariance observed by the investigator could be created: as I said, more than one way to skin a cat, even though everyone agrees that it ends up skinned. My basic skepticism stems from the fact that the method of measuring points of origin is fundamentally visual and requires extensive 2-dimensional arrays of data to be available and to be continuously evaluated. There has to be a simpler way.

>I argued that the shift in F2 upward for /di/ and downward for
>/du/, plus the upward shift in F1 for both, is a side effect of
>the tongue moving from its point of closure for /d/ to the vowel.

I agree completely. I can produce /di/ and /du/ sounds with no such shifts, quite easily. All you have to do is shape your mouth to the vowel before saying it, or pronounce some other vowel prior to the /di/ or /du/. Eedie, or Urdu.

>Contrast comes in to the question of how the input function of
>a phoneme detector is defined.

Agreed.

>Only a linguist or phonetician is likely to notice that some
>pair of different words are phonetically the same in the speech
>of the two people (but are phonetically different within the
>speech of either person, and phonemically contrasted in the
>speech of either and in the conversation as a whole). An
>untrained observer is much more likely to observe that "pen"

>and "pin" are not phonemically contrasted in dialects spoken
>for example near the Wendy City -- I mean, the Windy City.

I think we have to be careful here. "Pin" and "pen" are two different words, although both can be used to mean a writing implement. The speaker who seems to be using the same word for the two meanings is probably not: there is probably a phonetic distinction. If so, my weighted sums should pick it up. Actually, I think you're saying that there IS a distinction, but that the untrained ear doesn't hear it.

Dave Schweingruber (930720.0957) --

>I recall objections to the idea that people can "give
>themselves instructions" and that people can "change their
>reference signals."

I did raise a semi-objection, but it wasn't to the idea of variable reference signals (obviously). The main point I was trying to make was that control systems don't ever set their OWN reference signals. If people give themselves instructions, it is one level of organization giving instructions to another. The giver of the instructions is one kind of system, the receiver a different kind, a subsystem.

My other cautionary objection is related. When we say "give instructions" it's almost certain that we mean giving the instructions in the form of words or symbolic acts. But a system that can receive symbolic instructions must receive them from a higher system that perceives in other ways, not symbolically. So the giver of symbolic instructions is not itself a symbol-manipulating system. Furthermore, speaking of reference signals as instructions is misleading for most of the lower systems, where reference signals are simply examples of the perception that the lower system is to match. A system that is controlling a relationship of distance between a hand and a target doesn't tell the lower system "move the hand to the left." It simply adjusts the position reference signal for the hand control system. The metaphor of giving instructions should really be limited to the level that is specifically concerned with symbol manipulations, the program level.

See you at the meeting!

Best to all, Bill P.

Date: Tue Jul 20, 1993 7:57 pm PST
Subject: Good disturbances -- FYI

[From Rick Marken (930720.2100)] Dag Forssell (930720 1500)

>>I imagine that you (Dag) would suggest first trying to
>a)
>determine the wants a person has that are currently not satisfied
>b)
>>and then presenting PCT as a "good disturbance" that moves the
>>wanted perception to the reference level.

>With a) you are talking PCT, and that is always a good idea. With
>b), I can't figure out where you are coming from or what your
>terminology means.

I just made up the term; I thought it would be clear to the guru of the rubber band demo. A "good disturbance" is a disturbing influence which moves the controlled variable toward the reference state. If you start off with the knot in the rubber bands about one inch to the right of the target (reference state) then a pull by the experimenter to the left is a "good disturbance". The analogy to the discussion of "dealing with nonPCTers" would be as follows: find a want that the nonPCTer has that might be helped by knowledge of PCT. Perhaps a desire to understand "goal oriented behavior" or "to have a cooperative shop". The perception of this variable is presumably not at the reference -- like the knot that is one inch to the right of the target. Telling the person about PCT(in the right way) is like the leftward pull on the rubber band -- as you describe PCT their perecption of the controlled variable moves toward their reference as a result (partly -- they might act as well) of your "disturbance". Rememeber, the word disturbance, in PCT, just refers to a physical variable that influences a perceptual variable. The effect of the disturbance at any instant (independent of the effects of the controller) can be to move the controlled variable toward or away from the current value of the reference state (determined by the current value of the reference signal)."Disturbance" may have been an unfortunate choice of names for this variable since it always suggests "disruption"; but some disturbances can defintely be "good" as anyone knows who has gotten an unexpected bit of help from the undetected bump in the green that ends up turning a disastrous putt into a birdie.

Best Rick

Date: Tue Jul 20, 1993 8:31 pm PST
Subject: Formants

[from Gary Cziko 930721.0400 UTC]

Bill Powers (930720.1550 MDT) to Bruce Nevin (930720.1032)

>Your description of the "vowel triangle" suggests a space with
>only two dimensions, employing two "formants". I am unsure about
>just what constitutes a formant in linguistic terms. As the mouth
>changes shape, various resonances appear, but they are not at
>constant frequencies nor are there always the same number of
>them. When diphthongs appear, curved lines appear on the
>spectrogram, sweeping though at least an octave. This path
>creates the visual appearance of some single thing changing -- is
>this what is meant by a formant? Or are formants fixed bands of
>frequencies, counted as such whether or not any sound energy
>appears in those bands?

Let me give Bruce a bit of a break and try to explain how I understand formats (primarily from Lieberman's books).

The formants are the local peaks of the filter function of the supralaryngeal tract when the mouth is shaped to produce a particular sound such as the [i] of beet. According to some graphs provided by Lieberman, for [i] pronounced at a fundamental frequency of 0.5 KHz, the first formant will be about .4 KHz, the second format will be at about 2.3 KHz, and the third at about 3.1 KHz.

However, if the fundamental frequency = 0.5 KHz there will be overtones at multiples only (e.g., 1.0, 1.5, 2.0, 2.5, 3.0 and 3.5 KHz) and so there will actually be no acoustic energy (in this case) present at the first, second and third formants! But they will nonetheless be heard by the listener (a type of auditory extrapolation?). If [i] is said with a different fundamental frequency, it may be the case that the second formant frequency (or third) will now have lots of energy.

It is intriguing to think of what happens as one sings [i] with rising pitch. Energy at the formant frequencies will rise and fall as the overtones approach and then past the formant frequencies. But the ratios of the "imagined" formants will remain the same and it all sounds like [i].

Lieberman says:

"Almost 200 years of research demonstrate that human beings are equipped with neural devices that, in effect, calculate formant frequencies from the speech signal. We do this even when very little acoustic information is present, as in the case on a telephone. We apear to have innate knowledge of the filtering characteristics of the human supralaryngeal vocal tract--a complex neural formatn frequency "detector" that calculates the formant frequencies on the basis of an internal representation of the physiology of speech production. Computer programs that go through a similar process are able to calculate the formant frequencies of unnasalized sounds with reasonable accuracy." [no further info or references given].

It seems to me that what we are able to do in picking up these formants even with they don't exist (in terms of audio energy) is perceive the speaker's mouth configuration.

Bruce, did I get this right?--Gary

Date: Wed Jul 21, 1993 3:11 am PST
Subject: Some "Player" YOU are!

From Greg Williams (930721) Rick Marken (930720.0800)

>[I presume that this is the start of my requested role play. I guess it's
>meant to be a representation of how I respond to people on the net.
>If so, then I am being perceived (at least by Greg) as a lot nastier than
>I am trying to be. So, as a model of how Greg sees my posts, it is very
>disturbing.

My nit-picking was intended as a parody -- semi-humorous and exaggerated -- of the sort of thing which has appeared on the net in reply to some nonPCTers, posted not just by you. Perhaps by seeing how others sometimes can perceive this sort of thing ("a lot nastier than I am trying to be"), some netters will learn to consider others' perspectives just a bit more.

>Anyway, here is my reply to Greg's post -- with me
>in the role of the person who has been attacked by "Rick the Prick".

I don't think you are a "Prick." You're a "Player" (as Isaac Kurtzer said at last year's CSG meeting). But I think Isaac would be disappointed in you as a ROLE Player -- you're supposed to act AS IF you are the person in the role you are playing, not yourself! In my (attempt at a) role play, you basically just rolled over and played dead -- not at all like a feisty nonPCTer. With that sort of behavior, I have no recourse except to give you a PCT-supporter certificate and banish you to the realm of CSGnet lurkers. Gee, everybody missed out on umpteen posts of ever-increasing length, full of escalating arguments.

BTW, there is something to be learned here. Note that nobody else jumped on the points I nit-picked. (I waited awhile before doing it.) Rick, after all, is PCT-all-the-way. I have to wonder whether the same sort of benefit-of-the-doubt would be shown to a nonPCTer -- it hasn't, on occasion, in the past.

>Bill Powers (930720.0815 MDT)

GW to Rick Marken>>As Bobby Dylan asked, "How does it feel?"

>Or as Eric Berne put it, "Now I've got you, you son of a bitch!" (p. 85)

This recasting of my comment is quite wide of the mark. I am not trying to "get" Rick, nor do I think he is an S.O.B. I do think that it is important sometimes to "walk in others' shoes" (or at least "try their shoes on") in order to understand that what PCTers sometimes believe they are accomplishing in certain dealings with nonPCTers might be quite different from those beliefs. Overall, my aim is to aid in the further recognition of PCT ideas (without their being compromised) by pointing out some self-defeating approaches. In my opinion, gratuitous-sounding (to me) comments like yours above aren't helpful in this regard.

>Rick Marken (930720.1300)

>Without having read the book, here is my unsolicited take on what is
>going on in the interaction between PCTers and nonPCTers. When a PCTer
>makes a correct PCT statement (with explanation) it is very likely to
>have a disturbing influence on some high level perception that is
>being controlled by the friendly non-PCTer.

Sometimes nonPCTers act as if PCT ideas contradict some of their own; but sometimes (some) nonPCTers (apparently) just don't see any big deal regarding PCT -- no contradictions to their own ideas, no new insights, just nit-picking about marginal issues. In the second case, it appears that some significant (to them) disturbances are felt by PCTers.

>Since the non-PCTer can
>see the cause of the disturbance (Bill, Tom or myself) as well as
>experience the effect thereof, he or she is likely to conclude that
>we are being contrary, annoying and difficult -- after all, we are making
>it necessary for them to defend their perception, which ordinarily stays
>pretty much near it's reference.

Ditto for my second case above, with PCTers and nonPCTers reversed.

>I don't see any way to present PCT
>honestly and not have people occasionally get upset at the messenger
>because one or another of the PCT messages is bound to be a distrubance,
>and a big one, to SOME cherished belief.

In my second case above, ditto with "nonPCT" replacing "PCT."

>The only way people will eventually "come around" (if they ever do) is
>through reorganization -- which means abandoning some very precious
>and successful control systems. I'm not surprised that there are few
>people willing to do it.

I'm saying: see it from the other side, too. You are among the "few," yourself.

>Most people want to incorporate PCT into what they already know, which
>is VERY understandable.

And, I think, you want to keep PCT isolated ("revolutionary") from all else, which isn't quite so understandable.

>Greg thinks this is the right way to go --
>let nonPCTers incorporate PCT into what they already know. The
>problem is that, in order to do that, the nonPCTer has to keep some
>old control systems in effect -- which means that old misconceptions
>will be defended (these old misconceptions are probably what Greg
>calls "meta issues").

This just isn't clear to me. "Misconceptions" often have a way of turning into "different viewpoints" (not necessarily contradictory) when examined closely. But close examination isn't possible with extreme polarization "built into" the arguments, taken for granted from the very first BEFORE HEARING OUT the other ("ALL nonPCT psychologists are wrong," etc.). The apparent guiding principle is that there is nothing from "Devils" which could contribute to PCTers' understanding. The Standard Operating Principle is too often Us vs. Them, rather than We're All in This Together.

>The problem is that the wants satisfied by PCT usually make it possible
>to adopt only the language of PCT, but not the model.

Or, I would say (in the case of some engineers), adopt only the model of PCT, but not the language.

>It's hard to think of anything any PCTer could say that would not be
>a MAJOR disturbance to one or another of the perceptions that
>correspond to one of those wants.

Here's one: tell nonPCTers about how much evidence PCTers have collected showing that the HPCT architecture corresponds to reality. Another: tell nonPCTers about how many remarkable experiments have been conducted by PCTers on high-level controlled variables. And another: how many CSG members there are, currently.

The bottom line: humility is a virtue, especially for somebody with something to be humble about. I believe that PCT could have a great future, but it is now in its infancy. It will never reach adulthood if it goes around trying to drive its parents' car (thumbing its nose at the cops!) at a tender age. There is plenty of time for street racing and showing off after adolescence.

At the close of the PA CSG meeting a few years back, Bill Powers, in reply to comments about developing a rapprochement with nonPCT psychologists, stated: "I'd rather win." He and other PCTers have won several battles with nonPCTers, but I think they are losing the "war." Perhaps it is time to reconsider tactics and to realize that sometimes losing a few battles can aid winning wars.

As ever, Greg

Date: Wed Jul 21, 1993 9:30 am PST
Subject: Good disturbance

Dag Forssell (930721 0840) Rick Marken (930720.2100)

>A "good disturbance" is a disturbing influence which moves the
>[other person's] controlled variable toward the [that person's]
>reference state.

Your explanation is clearer, thanks. I think the term "good disturbance" will prove stillborn in the PCT context, however.

We have often illustrated our contention that *all* disturbances are resisted with the idea of helping the little old lady carry her suitcase. The lady will *reduce* her effort.

Your case is not necessarily the same (though it may well be), since you specify that the other person is not controlling well - there is an error, despite the person's best effort. However, when it comes to things like "desire to understand 'goal oriented behavior,'" there may be a great number of variables that enter into the picture. Self-esteem, pride, (whatever that really is) etc., etc. The error you see (the engineers point of view) may be the result of internal conflict, not failure to control. You would need a large Marken spread sheet to sort it out.

When i read your sentence above, I was wondering: *Who's* controlled variable? You made it clear a few sentences later. I think this is a key question. It gives you a hint to keep in mind as you enjoy "Freedom."

Best, Dag

Date: Wed Jul 21, 1993 9:59 am PST
Subject: Formants, Players

[From Rick Marken (930721.0900)] Gary Cziko (930721.0400 UTC)

>It seems to me that what we are able to do in picking up these formants
>even with they don't exist (in terms of audio energy) is perceive the
>speaker's mouth configuration.

This must be a metaphor, right? All we usually know of a person's mouth configuration is its acoustical results. In-person we can see lip and mouth configuration to some extent but these perceptions are clearly not crucial for understanding speech, as evidence by the continued popularity of radio commnuications. Because of the typically tight relationship between certain articulatory configurations and their acoustical effects it is "as if" we are perceiving the speaker's mouth configuration when we are perceiving speech (based only on the acoutical input). I suppose that this may be all Liberman is saying -- that we perceive speech by comparing the acoustical input to the outputs of a model (in our head) of the articulatory system. Bill, Tom and I are trying to point out the problems with such an approach, but the proof will be in the pudding; it will be settled (maybe) only when someone invents a perfect speaker independent voice recognition system that turns sounds into words as accurately as a secretary taking dictation.

From Greg Williams (930721) --

>Sometimes nonPCTers act as if PCT ideas contradict some of their own;
>but sometimes (some) nonPCTers (apparently) just don't see any big deal
>regarding PCT -- no contradictions to their own ideas, no new
>insights, just nit-picking about marginal issues.

Of course they often don't see any big deal. That doesn't mean that there isn't any big deal, does it? Perhaps you could give an example of something that a nonPCTer thought of as "no big deal" that was, in fact, "no big deal" but that a PCTer made a big deal of.

>And, I think, you want to keep PCT isolated ("revolutionary") from all
>else, which isn't quite so understandable.

I don't see it that way at all. In the discussion of speech perception I have been earnestly trying to make points about perception that I learned in perception classes over 20 years ago. Much of the conventional work in perception is completly consistent with PCT, and vise versa. I have tried to suggest to behaviorists ways to do PCT type research in operant conditioning situations; I don't want PCT isolated; I'd love to see it working in the hallowed halls of rattus norviegicus. I have tried to show (in publications when I could) how PCT could be applied to "conventional" problems in motor control.

I do want to isolate PCT from being incorrectly applied by authoritative sounding people. But it seems to me that most of the isolating of PCT has been done by the "other guys". I just can't believe that you can't see that that is what's going on. It's a little weensey bit painful to be accused of being the cause of another person's bias.

Greg, dear sweet Greg. It seems to me that you are blaming the victim (PCT) instead of the victimizer. I might be a bit verbally rambunctious on CSGNet here but I have tried to be VERY careful and diplomatic in my communications in the domains of the nonPCTer (mainly in publications and conferences). I know it SEEMS like PCTers themselves must be somewhat responsible for the avoidance of PCT by most life scientists. But my experience over 12 plus years working with this model is that life scientists will listen for a moment and then reject PCT as "nothing but" such and such (which PCT is always NOT) or say that PCT can't be right because we already know (from statistical studies) that feedback is too slow, etc (pick your favorite misconception). Did WE make the life scientists respond in this way? This happens no matter how nice we are -- and it's hard to top Bill Powers and Tom Bourbon for nice. I see the net and CSG meetings as a place where rejected PCTers can let off a little steam -- sort of like a PCT ghetto. If the high falutin' conventional types ain't comfortable in our ghetto then tough -- they aren't very nice to us when we come, hat in hand, to theirs.

>This just isn't clear to me. "Misconceptions" often have a way of
>turning into "different viewpoints" (not necessarily contradictory)
>when examined closely.

Well, give me an example of a misconception that was really just a "different viewpoint"?

> But close examination isn't possible with
>extreme polarization "built into" the arguments, taken for granted
>from the very first BEFORE HEARING OUT the other ("ALL nonPCT
>psychologists are wrong," etc.). The apparent guiding principle is
>that there is nothing from "Devils" which could contribute to PCTers'
>understanding. The Standard Operating Principle is too often Us vs.
>Them, rather than We're All in This Together.

I agree that the claim that "ALL nonPCT psychologists are wrong" is mega-hyperbole and could be divisive and alienating. If I ever said that then it was just silly and inappropriate. I can't believe I said that but maybe in some context I did -- and I agree, if some reader is real hyper- sensative then it would be an annoying thing to read. But this is a net that's partly for fun, I think. If someone were driven away from PCT because they peeked in on the net and heard me say "All nonPCT psychologists are wrong" then they are really a bit over-sensative. The real debates on CSGNet have been over specifics -- information in perception, supra- individual control systems, speech perception, research methodology, feedforward models, environmental control, motor control, etc -- and I never detected the assumption that there was "nothing that the Devils could contribute". A lot of the Devilish data is ambiguous, many of the Devilish models are not models and if they are they don't work, but some Devils have done some pretty useful stuff and we refer to it -- then they're not Devils (I prefer using the term Devil's for people who have made flat out wrong statements about control in a highly authoritative manner).

>Here's one: tell nonPCTers about how much evidence PCTers have collected
>showing that the HPCT architecture corresponds to reality. Another:
>tell nonPCTers about how many remarkable experiments have been
>conducted by PCTers on high-level controlled variables. And another:
>how many CSG members there are, currently.

Happily -- the answers are "some, a few, very very few".

>The bottom line: humility is a virtue, especially for somebody with
>something to be humble about. I believe that PCT could have a great
>future, but it is now in its infancy. It will never reach adulthood if
>it goes around trying to drive its parents' car (thumbing its nose at
>the cops!) at a tender age. There is plenty of time for street racing
>and showing off after adolescence.

Bill answered this well some time ago. You imply that existing behavioral science is some great scientific edifice before which PCT must be humble and tactful. There are probably many valuable observations and models that exist in this edifice, it's true. We have never denied that. It's just that the basic assumption on which this edifice is built (the cause-effect model) is false, at least from a PCT point of view. I think that you don't believe that and you think that it is the responsibility of PCT to show it's "parents" exactly what is wrong with each and every experiment and model that they've built. Maybe this is where our problem lies. For me the basic facts of PCT are these:

1) Systems that are in a negative feedback relationship with respect to their environment are control systems.

2) Control systems control their perceptual inputs, maintaining them at secularly adjustable reference levels; they are perceptual control systems.

3) All organisms that have sensors and output systems (that can influence the effect of the environment on those sensors) are perceptual control systems.

That's it. That's all we know from PCT -- except, as you say, as few little applications of the PCT model to research. But once you accept these three facts about control then most of what has been published and modeled in the behavioral sciences can be safely ignored (read Powers' Psych Review paper if you want to know why). However, if you or anyone else thinks that there IS stuff in the "parental" ediface that we should take along now that we've moved out, please let us know -- we won't automatically reject it (as I said, I'm taking along the Hubel-Weisel, Lettvin, etc work on receptive fields).

Best Rick

Date: Wed Jul 21, 1993 11:12 am PST
Subject: Rick's analogy

[From: Bruce Nevin (Wed 930721 13:01:09 EDT)] Rick Marken (930720.1300)

> All I'm saying is that this phenomenon is not restricted to speech.

I'm not saying it is restricted to speech. I'm saying I don't know how to model this. If you know how to model the case for vision, and if your analogy is valid, then you can answer my question: how do you model this?

I have no attachment whatsoever to language being different or requiring extensions to PCT. My questions and proposals are due to my not being able to see how PCT can account for what is going on. The rest is due to difficulty I have getting you to recognize that what is going on is going on, or to difficulty you have getting me to recognize that what I perceive as going on is in fact an illusory byproduct of something that PCT can account for.

The analogy to "brightness constancy" seems to apply well to the differences in voices that are due to length of vocal tract. If so, good, we can peel that layer away and simplify the problem a bit. It is not at all clear that the analogy applies to the other kinds of differences-that-don't-make-a-difference within a phoneme vs. differences that do make the differences between phonemes.

Bruce bn@bbn.com

Date: Wed Jul 21, 1993 12:58 pm PST
Subject: Good Disturbance, Perception

[From Rick Marken (930721.1300)] I said:

>A "good disturbance" is a disturbing influence which moves the
> controlled variable toward the reference state

Dag Forssell (930721 0840) says:

>When i read your sentence above, I was wondering: *Who's*
>controlled variable?

It doesn't matter. The way I defined it, a "good disturbance" is an influence that moves any controlled variable toward its reference level. The controller that is responsible for the fact that the variable is controlled is irrelevant to this definition. Obviously, if two different controllers are controlling the same variable relative to different reference levels then it is very possible that the same disturbing influence will be "good" with respect to the reference level for one controller and "bad" with respect to the reference level for another.

>You made it clear a few sentences later. I think this is a key question.
>It gives you a hint to keep in mind as you enjoy "Freedom."

Oh, don't be so coy, Dag. I think I know PCT pretty well. "Freedom" is based on PCT. What am I going to find out that I couldn't figure out on my own from first principles?

This all started with you saying:

Dag Forssell (930720 1000) --

>For my part, I sincerely believe that "Freedom" offers you well
>thought out suggestions on how to do *any* PCT role play *and*
>insight into how to deal with nonPCTers.

So you were claiming that "Freedom" would tell me something about how to deal with nonPCTers that I couldn't figure out on my own from knowing PCT. My guess was that the insight from "Freedom" would be:

>trying to determine the wants a person has that are
>currently not satisfied and then presenting PCT as a "good disturbance" that
>moves the wanted perception to the reference level.

Now that you know what I meant by "good disturbance" could you tell me if this was a reasonable guess? Maybe there is more -- or maybe it is entirely different. Isn't the insight from "Freedom" something that could be described in a paragraph or two? If not, fine. I'll wait until the meeting. But until then, how about telling me if there was anything wrong with the following suggestion:

>The problem for the PCTer trying to present the PCT model honestly
>is not one of finding wants that are satisfied by PCT; the problem is the
>existence of all those other wants that make certain aspects of the PCT
>model anathema:

You see, I don't think there is a "right way" to deal with nonPCTers -- that is, a way that will accomplish what I think is Greg's goal --of getting the nonPCTer to see the merits of PCT. PCT is a disturbance to many perceptions that nonPCTers seem to be controlling quite successfully. What can make a person dismantle successful control systems? Why would one try to dismantle another person's successful control systems, anyway? The only "right" way to change a nonPCTer's mind is to present the scientific evidence -- the models, the data, etc. There isn't much of it (as Greg notes) but what there is of it is EXTREMELY strong (how much data did Galileo have with which to revolutionize physics?). If the nonPCTer is non swayed by this data (and most are not) then that's pretty much the end of the conversation as far as I'm concerned; when data and modelling don't work, the only sure way to change minds is by force; and I'm not in good enough shape to do that (but don't get cocky, Cziko -- I have two weeks!).

"Freedom" is a book about clinical applications of PCT. The changes in control strategies that Ed describes in that book are made by people who (when they came to Ed) were NOT successfully controlling the variables that they wanted to be controlling. These people WANTED to change - they might not have known how to change (Ed helped with that) but they knew something was wrong and they were interested in trying to fix it. This is not AT ALL like the situation of the typical nonPCTer, who is usually QUITE successful at controlling some subset of the list of wants that I catloged in a previous post. The nonPCTer is a happy, skillful individual who is curious about PCT for various reasons. But few nonPCTers are having problems that would be helped in any way by PCT knowledge -- these people are IN CONTROL, as evidenced by the extremely skillful way that they counter the disturbing effects of PCT ideas. All PCT would do for most of these people (once they understood it properly) is get them alienated from their colleagues, make them nearly unpublishable and limit them to friends like Bill, Tom and me. Yikes!

I don't believe (based on my understanding of PCT and my experience with nonPCTers) that there is ANY way to change the minds of the nonPCTer (any more than there is a way to change the mind of a happy, in control PCTer, as Greg also corretly noted). The only legitimate method for trying to change a nonPCTer into a PCTer is the usual scientific method -- testing and modelling. If "Freedom" suggests something other than that then I would really like to hear about it. Maybe I could just turn straight to the "converting nonPCTers into PCTers" chapter?

Bruce Nevin (Wed 930721 13:01:09 EDT)--

>If you know how to model the case for vision, and if your analogy is
>valid, then you can answer my question: how do you model this?

There are a many books on perceptual modelling that might give you an idea of how to model perceptual phenomena. An excellent old (and mathematically difficult) book is by Floyd Ratliff's "Mach Bands". Julesz's "Foundations of Cyclopean Perception" has some models of stereopsis (which create stereo images without monocular pattern data!!).

>The analogy to "brightness constancy" seems to apply well to the
>differences in voices that are due to length of vocal tract.

Good. Here's a simple model of a perceptual function that produces the phenomenon I describe (10 units of light being "light" in one context and "dark" in another). The input to this perceptual function is the light levels from the two adjacent points in the visual field; the paper and the desk -- call these points x1 and x2, respectively. The perceptual function computes the ratio of these two inputs (x1/x2) and reports (continuously) the result as a perceptual signal. This perceptual signal represents the perecived "lightness" of the paper. When the light in the room is on, the ratio is 100/10 = 10. When the light is off, the ratio is 10/1 = 10. So in both situations the perceptual signal is the same (10) so the paper is perceived as the same lightness even though its absolute illuminance is quite different (100 vs 10). This is a VERY simple model of a perceptual function. The model that produces this kind of constancy in phomeme perception will be a bit more complex -- more layers of perceptual functions for example -- but the principle (of a function converting inputs into a perceptual signal) will be the same.

>The rest is due to difficulty I have getting you to recognize that
>what is going on is going on, or to difficulty you have getting me to
>recognize that what I perceive as going on is in fact an illusory
>byproduct of something that PCT can account for.

I think I know what you are saying and I agree with your description of the problems involved in phomeme perception. These problems have been addressed (and often handled fairly successfully) by nonPCT models of speech perception. The problem of speech perception has not been solved by any means and PCT's hierarchical model of perception may suggest new approaches to solving it -- but it is not an inscrutably difficult problem; and I'm pretty sure that it doesn't require mental models of the articulatory process in order to solve it.

Best Rick

Date: Wed Jul 21, 1993 1:07 pm PST
Subject: resolving contrast

[From: Bruce Nevin (Wed 930721 16:37:41 EDT)]

Drat. I sent a reply to Gary (Thanks! Right on the mark!) and to Bill (pin and pen are both phonemically identical and phonetically indistinguishable--both vary over the same phonetic range--in the dialect I mentioned) but somehow did not include the file in my message. Just learned of the mistake from a message from the listserver that it wouldn't distribute an empty post, too late to rescue the file because I had already overwritten it with my note to Rick. Will look again at messages tomorrow if I get a break.

Bruce bn@bbn.com

Date: Wed Jul 21, 1993 1:57 pm PST
Subject: Higher level control and Education

[Tom Hancock (930721)] Tom Bourbon (930719.1705)

>I have been thinking about your discussion, primarily with Rick,
>about your research design. If you want to modify your research so
>that it more clearly includes elements from PCT, you need a
>controlled variable -- Rick has been saying that, I believe.

Yes, I see how important the precise definition of the controlled variable is. As I have considered this at the higher levels it seems to me that the environment is not just the sensed physical environment but also the internal environment stimulated by the activation of previous memories and imagined perceptions (Ed Ford says something like this too). That is why I am still unclear why the one correct name label associated with a particular tone and a particular geometric shape on a particular position on the screen is not a controlled perception. The subject must work his perceptions with memories and imaginations until the association with the least error signal is attained. Or am I all wet?

>..you are studying associations between arbitrary geometric
>patterns (is that right?) and "real" words..

Not quite arbitrary. (See the paragraph above.) Actually my concerns are somewhat constrained: I want to improve the efficiency of training, so my theoretical concerns are placed within an applied framework. This particular task is concerned with a variation of a real training task.

>..could you make the patterns continuously variable, affected by
>slow random disturbances, and look to see if there are any
>identifiable configurations your people control as "the right ones?"
>This would involve you looking for reference levels of
>environmental variables, with the participants setting their own
>references for perceptions of those variables.

I found your suggestions thought-provoking. I will consider them more. At this point it appears that you have identified a higher level type instantiation of the tracking task. It seems to be a good way to give a post-training type of assessment of control after the subject has had a chance to get novel associates under some degree of control. My focus has been more on the learning end--identifying what the learner's level of control is and giving post-performance information that would serve as a disturbance for better control (a good disturbance?). Perhaps the drill could be cast in terms of exploration time of various name associations, followed by the assessment you indicated, then followed by post-performance information such as display of the correct and advice about what may be the problem with an inaccurate or unconfident choice.

On the other hand I wish I could understand better what you were thinking would be the PCT advantage of continuously varying patterns, with disturbances? I see that it does fit nicely with previous PCT methods, but it seems to me that giving the trainee one static set of the three configurations and the choice of name labels (with one correct one) would serve the same purpose. If the subject chooses the correct label he has that name-association under control. If the subject chooses a incorrect name label that differs on one of the three configurations (a different tone for example) then we could assume that the subject has the other two configurations (identical in the correct and incorrect) under control with that name label, but not the tone.

>And now that I think of it, Hugh Petrie talked about similar kinds of testing

Do you mean that your suggestions seem to be similar to the multiple choice testing that Petrie criticized?

Tom Hancock

Date: Wed Jul 21, 1993 6:15 pm PST
Subject: Steamroller play

From Greg Williams (930721 - 2) Rick Marken (930721.0900)

I'm convinced, Rick.

Now on to more important things.

As ever, Greg

Date: Wed Jul 21, 1993 8:13 pm PST
Subject: Perceived insults?

[From Dag Forssell (930721 2050)] Rick Marken (930721.1300)

>I think I know PCT pretty well. "Freedom" is based on PCT. What am
>I going to find out that I couldn't figure out on my own from
>first principles?
>
>>For my part, I sincerely believe that "Freedom" offers you well
>>thought out suggestions on how to do *any* PCT role play *and*
>>insight into how to deal with nonPCTers.
>
>So you were claiming that "Freedom" would tell me something about
>how to deal with nonPCTers that I couldn't figure out on my own
>from knowing PCT.

You are putting a lot of words in my mouth. I have not said that you could not figure something out on your own. You are both extremely smart and knowledgeable of the theories of PCT. It is another question entirely to see evidence that you have in fact figured something out on your own and use it. With the subject at hand, that takes a keen interest in applications and many sleepless nights. My impression is that Ed is way ahead of you in that regard, which is why I think he offers well thought out suggestions.

>Now that you know what I meant by "good disturbance" could you
>tell me if this was a reasonable guess? Maybe there is more -- or
>maybe it is entirely different. Isn't the insight from "Freedom"
>something that could be described in a paragraph or two?

I indicated that I think the term "good disturbance" is non(PCT)sense. I don't think this is a reasonable guess. The answer is different and more PCT kosher. Perhaps also more limited. I think you underestimate "Freedom." If I tried to summarize it in two paragraphs, you would not get any feel for the practicality and range of its wisdom. Your suggestion that it is that simplistic indicates why you have not wanted to read it until now. It has obviously been a waste of your time.

BCP could also be summarized in two paragraphs if you wanted to. How enlightening would that be?

>The problem for the PCTer trying to present the PCT model
>honestly is not one of finding wants that are satisfied by PCT;
>the problem is he existence of all those other wants that make
>certain aspects of the PCT model anathema:

I do not know enough (hardly anything) about the wants of your target nonPCTers to guess about the validity of this.

I look forward to discussing applications and the rest of your post with you in Durango, now that you have read "Freedom."

There is much to agree with about the subsequent comments in your post, which I will table for now.

See you in Durango.   Best, Dag

Date: Thu Jul 22, 1993 5:30 am PST
Subject: Re: Formants, Players

[From: Bruce Nevin ()]   Rick Marken (930721.0900)

> Gary Cziko (930721.0400 UTC)

> >It seems to me that what we are able to do in picking up these formants
> >even with they don't exist (in terms of audio energy) is perceive the
> >speaker's mouth configuration.

> This must be a metaphor, right? All we usually know of a person's mouth
> configuration is its acoustical results.

I think Gary knows this, and is referring to the proposal (floated some time back) that we perceive the speaker's intended phonological gestures in imagination: the gestures that we would have intended had we produced a like acoustic signal. As I said recently, I conjecture that vowels are controlled acoustic perceptions and that consonants are controlled kinesthetic perceptions with acoustic consequences at the margins of adjacent vowels. The hypothesis, then, is that the hearer imagines the intended acoustic perceptions for vowels and the intended kinesthetic perceptions (with acoustic byproducts) for consonants. It is the intended perceptions that are invariant.

You then go on to say something similar (so you need not have supposed that it was a metaphor, right?):

> we perceive speech
> by comparing the acoustical input to the outputs of a model (in our head)
> of the articulatory system.

I suppose you could call it a model. If I hear a clattering sound in the kitchen and I know just which utensil the cat has knocked off the counter, is that imagined perception (including visual image of knife) a model? I think what is involved is something like this general- purpose capacity for intuiting "what could possibly have made that sound?" For speech, the inventory of possibilities is limited, and is organized in a cross-classifying matrix of interrelations of similarity and contrast within which the terms are maximally differentiated, so that the "objects" to be recognized are more familiar than kitchen utensils. (This despite the fact that the "space" for the matrix may be compressed for lenis pronunciation or expanded for fortis pronunciation, or the matrix may be shifted within the available "space" whatever its current size, or two terms in the matrix for one dialect may correspond to a single term in the matrix for another, for one or a very few terms.) I suppose this matrix and the entities in it constitute a model; and I suppose that my memories of what I have in my kitchen constitute a model. I suppose that it's not necessary to use the concept in either case. Just memory of what makes this or that kind of sound when I do it.

> Bill, Tom and I are trying to point out the
> problems with such an approach, but the proof will be in the pudding;
> it will
> be settled (maybe) only when someone invents a perfect speaker
> independent voice recognition system that turns sounds into words as
> accurately as a secretary taking dictation.

Just so. Though to be marketable it would have to be more accurate :-) And that depends crucially on its control of social perceptions. The boss is not always pleased to have to repeat something that he didn't realize was ambiguous or indistinctly said (two different problems).

In your ongoing debate with Greg, you are talking about the intellectual, theory-constructing core of science, and Greg is talking about the politics of science. The debate recurs and is worthwhile I think most of all because lurking under it is a major challenge for applied PCT in an area that we understand very poorly: social relations.

Greg's metaphor about infant PCT not being ready to drive dad's car and being ill advised to try it while thumbing its nose at the cops is a metaphor about politics. Your reply is about truth and the validity of theories. This is how you are talking past each other. Greg, it would maybe help if you kept the two aspects of the problem clearly distinguished from one another, and clearly labelled. Am I wrong?

The political reality is that students and others want careers as well as the pursuit of Truth. Career paths, jobs, grants, publication, steady income, esteem of others, and so on, are controlled perceptions. They are the essence of politics.

Bruce bn@bbn.com

Date: Thu Jul 22, 1993 7:29 am PST
Subject: Re: Higher level control and Education

From Tom Bourbon [930722.0920] Tom Hancock (930721)
>
>>I have been thinking about your discussion, primarily with Rick,
>>about your research design. If you want to modify your research so
>>that it more clearly includes elements from PCT, you need a
>>controlled variable -- Rick has been saying that, I believe.
>

>Yes, I see how important the precise definition of the controlled
>variable is. As I have considered this at the higher levels it seems to
>me that the environment is not just the sensed physical environment
>but also the internal environment stimulated by the activation of
>previous memories and imagined perceptions (Ed Ford says
>something like this too). ...

Perhaps I misunderstood the nature of your original request for sugestions about research designs. My comments were more or less in the form of free associations -- tossing out some thoughts to see if I could better understand your felt needs. My first thought was that, whichever level of pereption *you* (Tom Hancock, experimenter) want to study, you must work in the participants' environment. There, you can either watch, while a person interacts with an environment you do not perturb, or you can perturb it -- disturbing variables you believe might, or might not, be controlled by the person. I was just suggesting some ways you might come up with a new set of variables, or new ways of playing with some old ones.

> That is why I am still unclear why the one
>correct name label associated with a particular tone and a
>particular geometric shape on a particular position on the screen is
>not a controlled perception.

A person might, or might not, make that a controlled relationship. Your job as experimenter is to determine which is the case. Instructing a person to make that (your ideal) become a controlled perception does not make it so -- but you know that.

Assumimg that a person does accept your suggestion, it looks as though, at different times in the procedure, the person might lock onto one or another of what seem to be several possible relationships in the display. So I thought you might want to develop a method that would allow you to identify, with some precision, which part, or parts, of the set of elements the person might be controlling at various times. (Later in your post, it becomes obvious that a person can indeed control different relationships out of the bigger set.)

> The subject must work his perceptions
>with memories and imaginations until the association with the
>least error signal is attained. Or am I all wet?

Not at all, but you need a method for identifying what the person can *do* with that ever-more-complex set of perceived-imagined associations.

>>..you are studying associations between arbitrary geometric
>>patterns (is that right?) and "real" words..
>
>Not quite arbitrary. (See the paragraph above.)

The following is not arbitrary? "... one correct name label associated with a particular tone and a particular geometric shape on a particular position on the screen..." (By "arbitrary," I mean only that the associations were set by the will or judgement of one person, not by an "impersonal law of nature."

Not that nature isn't be arbitrary, As Calvin (of Calvin and Hobbes, the comic strip) recognized when he announced, "gravity os arbitrary.") ..

>>..could you make the patterns continuously variable, affected by
>>slow random disturbances, and look to see if there are any
>>identifiable configurations your people control as "the right ones?"
>>This would involve you looking for reference levels of
>>environmental variables, with the participants setting their own
>>references for perceptions of those variables.
>
>I found your suggestions thought-provoking. I will consider them
>more. At this point it appears that you have identified a higher level
>type instantiation of the tracking task. It seems to be a good way to
>give a post-training type of assessment of control after the subject
>has had a chance to get novel associates under some degree of
>control.

Good. Might the method also be used *during* the training process, to provide evidence of the degree of control the person can achieve, and which variables are in and not in the controlled set?

> My focus has been more on the learning end--identifying
>what the learner's level of control is and giving post-performance
>information that would serve as a disturbance for better control (a
>good disturbance?).

Ah! *This* is the (subliminal?) source of the "good disturbance" that popped up in the discussion between Dag and Rick. Or did they say it first?

> Perhaps the drill could be cast in terms of
>exploration time of various name associations, followed by the
>assessment you indicated, then followed by post-performance
>information such as display of the correct and advice about what
>may be the problem with an inaccurate or unconfident choice.
>
>On the other hand I wish I could understand better what you were
>thinking would be the PCT advantage of continuously varying
>patterns, with disturbances? I see that it does fit nicely with
>previous PCT methods, but it seems to me that giving the trainee one
>static set of the three configurations and the choice of name labels
>(with one correct one) would serve the same purpose.

It was only a suggestion. The things you had been trying didn't seem to be working to your complete satisfaction. (Is there something inherently wrong with continuous variables? I have been told as much by reviewers and editors.) Perhaps I don't understand the real purpose of the training in your study -- what should a person be able to do after succeeding at your task?

> If the subject
>chooses the correct label he has that name-association under control.

I might say the person has *identified* (learned) the association. Whether it is *under control* is another matter -- the one where PCT might help you the most, if it can at all.

> If the subject chooses a incorrect name label that differs on
>one of the three configurations (a different tone for example) then
>we could assume that the subject has the other two configurations
>(identical in the correct and incorrect) under control with that name
>label, but not the tone.

Again, the person associates something, but not everything you would like to see associated. But is it under control?

>>And now that I think of it, Hugh Petrie talked about similar kinds of testing
>
>Do you mean that your suggestions seem to be similar to the
>multiple choice testing that Petrie criticized?

That wasn't what I meant. Hugh did not criticize multiple choice testing; he criticized the way people usually interpret the results of such tests -- as revealing how much of "something" (like "knowledge") a person "possesses." Hugh criticized the idea that such tests provide *objective* information about how much of something a person had absorbed, acquired, assimilated -- or associated. He suggested that the same conventional devices (multiple-choice tests, for example) could be used as tests for controlled perceptions. The example I gave of an adult "playing with" a child was like an example I recalled seeing in his chapter in Ozer's book. So, during my free association, I gave an example of how to "disturb" a person's perceptions of associations, then I recalled that Hugh had given simmilar examples. That was all I meant.

I hope some of this has helped jog your "little gray cells" into new ideas about your research design. ("the little gray cells" -- a favorite phrase of one of my favorite great detectives, one M. Hercule Poirot.)

Until later, Tom Bourbon

Date: Thu Jul 22, 1993 9:05 am PST
Subject: Hancock Memory, No Insult

[From Rick Marken (930722.0900)] Tom Hancock (910793.1100) --

>Actually, I was thinking that the p is the association provided by me
>but that the p*, the subject's perception, is not the same.

I think part of the problem here is that you are studying something that we haven't studied much in PCT -- memory. I think the perceptions that are actually controlled in your experiment are imaginations -- replays of previously experienced perceptions. What you are trying to figure out is something I don't understand at all, which is "how does a person know that an imagined perception corresponds to a pre- viously experienced perception?" The associations you are talking about as perceptions are perceptions -- the subject can perceive that after seeing X you say Y. p* would be the reference for that perception; at first the subject may have no reference for the association -- s/he doesn't care what comes after X. But you have asked thats/he learn to have a reference for this association that is the same what s/he experienced earlier. So after perceiving X-Y s/he is supposed to store a reference value, p*, of X-Y . The next times/he sees X s/he should produce Y to match my reference. So s/he can control the perception of the association by producing Y (instead of Z, etc) when you show him/her X.

Now it's pretty well known that people seem to have trouble storing the appropriate reference levels for associations, especially when there are several different references to be learned. What is also apparent is that people seem to have some idea of whether or not they have stored the appropriate reference for an association. This is the part I don't understand (actually, I don't understand the first part either).

So your study is really about two interesting aspects of memorization; and I don't know of any PCT studies of memory so your study is a very important start. There are probably single subject studies of memory that might have the kind of data you could use as a start at building a PCT model of this process. But here is a suggested version of your study (that could easily be done on a computer) that would give you the kind of data that might provide a basis for modelling:

Have the computer present letter - letter associates (for simplicity). After the first letter the subject types in the associate (which is, of course, a guess the first time) followed by a rating (1-10) of certainty that it is the right associate. Use about 10 letter-letter pairs. Keep running the experiment until the subject is getting the right associate to each letter (and, I presume, rating 10 for certainly each time). Thus, at the beginning of the study the subject cannot control the associations (relative to the references you want them to learn) at all; by the end of the study the subject is controlling the associations perfectly. Every disturbance (first letter) is countered by the appropriate action (second letter) to produce the intended association -- which should also be the association that you consider correct too. Now you can look at the relationship between certainty ratings and whether the actual association was correct or not on each trial FOR EACH ASSOCIATE. So, for one association over trials the sequence of right (R) and wrong (W) associations along with the corresponding confidence ratings might look like this:

W (1) R (7) W (2) W (1) R (8) R (6) W (2) R (10) R (10) R (10) R (10)

Now see what the relationship between right/wrong and confidence ratings over trials looks like for each associate. Maybe the pattern will suggest something about how the the reference memories are laid down. See if there is the same relationship between right wrong and confidence rating over trials for each associate. What kind of relationship is it? If there is a systematic relationship -- and the same one for each associate then maybe the data can be the basis for a first shot at a control model of memory. Actually, this would say something about the development of control -- very hard to model, I think.

I must confess that I don't know, right off the top of my head, how to even approach modelling our ability to "know that we know" -- ie. how we are able to accurately report that what we are rememebering is, in fact, a prior experience. We obviously don't do it perfectly -- people confabulate and make mistakes (" I was SURE I left it there"). But we seem to get it right most of the time -- the subjects will probably be pretty good at knowing whether their answers are right or not. How do they know that? That was you original question, wasn't it? I now have a new PCT answer: "I don't know".

I think the difficulty of your study comnes not so much from the fact that is has to do with higher levels; it's just that it has to do with an aspect of control that we have not studied much -- memory.

So my suggestion is to simplify the task as much as possible, collect data from one subject (to contain the impulse to average over subjects; once you get perfect data from one subject you can try your technique on another subject, and then another ...). Collect at least some good amount of data when the subject is clearly in control (as when the subject is providing the right associate--and a 10 rating -- every time; this might be boring but once the subject is controlling like this you can see how time is involved in the control process -- start speeding up the rate of presentation of the first letter, for example, and see what happens to control then -- maybe this will help you get some idea of what the dynamics of the model should be like).

Keep working on it. We really need some work on the memory aspect of the PCT model. I think your approach is kind of like jumping in the deep end -- but at least you are jumping in.

Dag Forssell (930721 2050) --

I wasn't insulted at all by your post Dag. I'm sorry if I seemed so. I was really just curious -- I thought you were playing a little cat and mouse with me, but I see that you really were not, which is clear from the following:

>I think you underestimate "Freedom." If I tried to summarize it in
>two paragraphs, you would not get any feel for the practicality and
>range of its wisdom. Your suggestion that it is that simplistic
>indicates why you have not wanted to read it until now. It has
>obviously been a waste of your time.

>BCP could also be summarized in two paragraphs if you wanted to.
>How enlightening would that be?

I agree that it wouldn't be very enlightening. But I do think it is possible to give someone a sense of what BCP is about in two paragraphs (or less). That's all I wanted. I think you can give a similar, short description of "Newtonian physics, relativity, evolution, Buddism, etc. If it's done well then people might want to look more deeply into it. If you had to read everything through in order to know it's value we'd never be able to determine how to allocate our limited study time.

If you gave me a two paragraph description of what "Freedom" says about how to do a role play or how to deal with nonPCTers then I would take it for what it is -- a summary. If it seemed like something that might be worthwhile, I would then read the whole thing (which I will do anyway, don't worry).

Best Rick

Date: Thu Jul 22, 1993 12:22 pm PST
Subject: Perception and Imagination in Speech

[From RIck Marken (930722.1230)] Bruce Nevin () --

>You then go on to say something similar (so you need not have supposed
>that it was a metaphor, right?):

The metaphor is that we perceive articulations. For me, a perception is a signal that represents some aspect of external reality:

reality -->|sensory transducers| --> |perceptual function| --> perception

The aspect of reality that is sensed in speech is usually only the acoustic signal. The perception of speech is a representation (defined by the perceptual function) of the sensed acoustical signal.

> If I hear a clattering sound in
>the kitchen and I know just which utensil the cat has knocked off the
>counter, is that imagined perception (including visual image of knife) a model?

Good example. I think the knife is, indeed, a visual image that happens to be associated with the sound (an image that might very well turn out to be the wrong one -- it might be a fork). The sound is "recognized" in terms of its associated (imagined) source -- the image of the dropped knife.

>I think what is involved is something like this general-
>purpose capacity for intuiting "what could possibly have made that sound?"

This may be true for you (a trained linguist) but it sure ain't true for me. When I hear a phomene it's just a phoneme -- there is absolutely no associated imagery (visual or kinesthetic). But when my cat knocks over a vase, I DO instantly cull up various images about what might have caused that sound and, worse yet, what things might look like now as a result. For me, phonemes are just phonemes -- they don't sound like articulations at all. But once you've learned the regular relationship between phonemes and articulations then the sounds probably do start "sounding like" articulation patterns -- becuase you can imagine the articulations. I think linguists can learn how to imagine the articulation patterns that produce speech sounds only after they have developed the perceptual input functions that provide the perceptions whose articulatory causes are to be imagined. In general, I think a control system must develop the ability to perceive phonemes (ie. it must develop the perceptual functions that produce outputs that correspond to the presence of

/d/ /b/ and /g/ etc) before it can learn how to vary the references for the lower level articulatory perceptions that produce those phoneme.

Which makes me think that there may be a new PCT slogan possibility, viz:

"Learning to control begins with learning to perceive".

You can't control the perception of an acoustical variable until you can perceive that variable. So the child's articulatory apparatus may be perfectly capable of articulating the difference between /p/ and /b/ (in terms of voice onset, right?) but the child cannot control that difference until it can perceive the difference between /p/ and /b/ in terms of how they sound. So it can't recognize /p/ and /b/ in terms of articulation (at least when it first learns them -- the arti- culatory perceptions may be incorporated into the recognition process later, once the child has learned to control what it can already perceive) so articulatory perceptions cannot be essential to recognizing speech (though they may become involved after one has learned to control speech perceptions).

Best Rick

Date: Thu Jul 22, 1993 1:11 pm PST
Subject: repeat to Bill P.

[From: Bruce Nevin (Thu 930722 16:49:19 EDT)]

(Bill Powers (930720.1550 MDT)) --

> Your description of the "vowel triangle" suggests a space with
> only two dimensions, employing two "formants". I am unsure about
> just what constitutes a formant in linguistic terms. . . .
> With this approach, the vowel triangle using only two formants
> would be a special case involving only two dimensions. Perhaps
> whoever came up with this triangle tried multiple frequencies and
> found that two were sufficient. But it's possible that by using
> all the frequencies, greater discrimination can be achieved.

The vowel triangle was first described by contemporaries and prececessors of A.G. Bell by tracking the positions of the tongue for different vowels. The vowels [i] and [u] are high, [a] is low, [i] is front, [u] is back (plus lip protrusion/rounding). It is really a vowel trapezoid, as there are front and back "ah" sounds, but I have the symbol only for the front one, [a], on my keyboard. These are considered the "cardinal" vowels at the apices of the vowel space: the tongue cannot be lowered any farther than for [a] and its backed corrollary, nor the jaw opened any wider; the tongue cannot be raised any higher for [i] or [u] without producing turbulance (a fricative consonant sound). Other vowels are distributed around the periphery and within the space defined by the cardinal vowels.

Sometime in the 1950s or early 1960s somone noticed that if F1 is plotted on one axis and F2 on the other of a Cartesian graph in two dimensions (with one of them inverted, I forget which), the plotted positions of the vowels approximate very closely to the articulatory map of the positions of the tongue for those vowels.

There are also other vowels, for example [y] (of e.g. French tu) is a front vowel like [i] with lip protrusion like [u], and the sound represented by [i] with a bar across it (I can't overstrike i and hyphen on the screen) is a back vowel like [u] without the lip protrusion. The Valley Girl pronunciation of the vowel in words like "good" comes close to this. This vowels is not distinctive in English, but it is in e.g. Turkish. Thus, the /U/ of "good" can be unrounded like the Turkish vowel in lax speech (e.g. an unstressed "could" in "he could go too, you know") or in some southern California dialects and it doesn't matter because it's not a difference that makes a difference for other speakers of English. In Turkish, however, it's like the difference between "goody" and "giddy" is for us. These two sounds [U] and [I] are also distinctive in Turkish, each contrasting with barred-i, whereas in English there is only the two-way contrast.

(The convention is to put phonetic symbols in square braces, phonemic between virgules: [i] is a particular sound defined phonetically; /i/ is often thought of as a class of sounds that are not distinctive in a given language but strictly speaking is a logical symbol which together with others represents a set of contrasts in the given language.)

That's the background on the correlation of the articulatory map with the acoustic map. Of course this may be irrelevant for PCT, but that's the history.

> A tracking filter set to a high frequency would be
> excited by all lower harmonics. This would lead to a way of
> accurately perceiving the basic voice frequency, information that
> is NOT contained in the formants. Obviously we need a way to
> detect voice frequency per se in order to recognize and produce
> things like melodies.

We need it to perceive the syntactic intonation contours of assertion, question, and exclamation (which may be nested and are important perceptual cues to word dependencies), and also affective intonation-gestures. We need it together with perception of relative loudness and perhaps other factors to perceive differences of syllable stress (e.g. the verb and noun pronunciations of "export"). Charles Ferguson had some evidence years ago that infants practice intonation contours used in sentences during their babbling long before they have any words to plug into the "sentence intonation" frames. He suggested a way of learning a second language based on these findings--babble nonsense syllables in the other language's characteristic intonation patterns first; then start controlling phonemes and syllables of the other language in the nonsense syllables; then gradually phase in words and phrases.

I think this is most of what I lost yesterday, except that I thanked you a lot, Gary, for your excellent summary of some info from Lieberman on formants. Lieberman also explains why a deep voice is easier to understand than a high-pitched one: there are more data points per formant.

Bruce bn@bbn.com

Date: Thu Jul 22, 1993 1:55 pm PST
Subject: Re: speech analysis

[From Ray Allis (930722.1426 Pacific)]

This was in the mail yesterday.

AI Vol. 3, No. 29
IS July 21, 1993
CS THE COMPUTISTS' COMMUNIQUE

[lots deleted]

OGI Speech Tools is a free toolkit Oregon Graduate Institute's Center for Spoken Language Understanding (CSLU). It includes and X-windows display for speech signals, spectrograms, and phoneme labels; a neural-network package; C library routines; data format converters; and audio drivers for Sun4, DEC/MIPS, and Mac. FTP from /pub/tools on speech.cse.ogi.edu. [Johan Schalkwyk (johans@anago.cse.ogi.edu), comp.speech, 7/17/93.]

Ray Allis - Boeing Computer Services ray@atc.boeing.com

Date: Thu Jul 22, 1993 4:23 pm PST
Subject: PCT in research on a training task

[From Tom Hancock (930722.1630) Tom Bourbon [930722.0920]

I am impressed at how thorough you are with your posts. I will try to do likewise.

>Perhaps I misunderstood the nature of your original request for
>sugestions about research designs. My comments were more or
>less in the form of free associations -- tossing out some thoughts
>to see if I could better understand your felt needs.

The nature of my request was a need to understand how PCT explains the phenomenon I am interested in. I appreciated your post.

>My first thought was that, whichever level of
>pereption *you* (Tom Hancock, experimenter) want to study, you
>must work in the participants' environment.

Yes, PCT has been an immense help to me in seeing more clearly that the experimentor's mindset must not be the primary focus, if science is to advance.

>There, you can either watch, while a person interacts with an
>environment you do not perturb, or you can perturb it --
>disturbing variables you believe might, or might not, be controlled
>by the person.

Yes, and with the phenomenon (up til now) that I have been trying to understand, what I have tried to WATCH about the subject is the following: task response times, certitude rating times, and post- performance in formation study times--all in the context of the subjects rating how certain they are of their response and in the context of being told that they choose a right or a wrong association label. I have thought that at response (in other words at testing) the disturbance has been a test item with a need attached to identify a correct name label (yes, subjects accept that need with varying degrees of vigor--which was the point of some of my original posts along this line). And that the certainty rating is a reflection of the unreduced errror still persistent in their attempt to match the perception in the task with the reference, which reference is a memory of the previous trials with this item. The response time also has appeared to be a measure of the amount of unreduced error, with longer times indicating less well learned items and usually inversely correlating with certainty ratings. And that the post-performance information study time (formerly instructional feedback time) reflects the degree of the subject's attempt to reduce discrepancy--usually the lower the previous certainty (indicating greater error) and also if they are told the response was wrong (which should produce an error signal if the subject is controlling for seeing correct messages) the greater the opposition is in terms of a longer post-performance information time.

>I was just suggesting some ways you might come up with a new
>set of variables, or new ways of playing with some old ones.

Your suggestions are helpful and appreciated.

>> That is why I am still unclear why the one
>>correct name label associated with a particular tone and a
>>particular geometric shape on a particular position on the screen
>>is not a controlled perception.

>A person might, or might not, make that a controlled relationship.
>Your job as experimenter is to determine which is the case.
>Instructing a person to make that (your ideal) become a controlled
>perception does not make it so --
>but you know that.

Yes, I do know that (and as a teacher I experience that), and the my intent is to determine (with the measures indicated earlier) when a trainee is controlling an association. When he is not then he could use a disturbance that might move him into control.

>Assumimg that a person does accept your suggestion, it looks as
>though, at different times in the procedure, the person might lock
>onto one or another of what seem to be several possible
>relationships in the display.

Yes, that is what happens. The subjects who seem to have a high degree of control for getting finished with the drill seem to do this more, and to a lesser extent the subjects who are controlling (my hypothesis!) for seeing correct responses but not so much really understanding why it is correct.

>but you need a method for identifying what the person can *do*
>with that ever-more-complex set of perceived-imagined associations.

Yes, that is what I am trying to pin down.

>By "arbitrary," I mean only that the associations were
>set by the will or judgement of one person, not by an "impersonal
>law of nature."

Oh. I was reacting to my awareness of the instructor's job: he is given a set of objectives which may have been partially arbitrary to start with but when they are in his hands they are sometimes quite defined.

>Might the method also be used *during* the training process, to
>provide evidence of the degree of control the person can achieve,
>and which variables are in and not in the controlled set?

Yes, and I hope to explore that type of task.

>Ah! *This* is the (subliminal?) source of the "good disturbance"
>that popped up in the discussion between Dag and Rick. Or did they
>say it first?

I saw them use that term. It may fall by the wayside, but it fits with my needs: my primary way to help a student is when it seems that he is not controlling for something that is a part of the required course of study and I provide some post-performance information which I hope will be a good disturbance.

>Is there something inherently wrong with continuous variables? I
>have been told as much by reviewers and editors.)

I often prefer to use them, but practically to use continuous variables has sometimes been more effort for me.

>Perhaps I don't understand the real purpose of the training in
>your study -- what should a person be able to do after succeeding
>at your task?

The eventual training objective is that the person will correctly recognize which of multifarious radar systems which has locked onto their plane. My experimental training concern, however is with the correct identification of the name label for each of 27 different combinations of tone, waveform shape, and waveform position.

>> If the subject chooses the correct label he has that name-
>>association under control.

>I might say the person has *identified* (learned) the association.
>Whether it is *under control* is another matter -- the one where
>PCT might help you the most, if it can at all.

You are right and that is what I am interested in. I am convinced that PCT explanations are the way to go with the social sciences; I am disenchanted with inconclusive laundry lists and models that are not fully supported with data and I am impressed with the way students lives change with an application of the principles.

>Do you mean that your suggestions seem to be similar to the
>multiple choice testing that Petrie criticized?

>That wasn't what I meant.

Oh!

>I hope some of this has helped jog your "little gray cells" into new
>ideas about your research design.

Very much so! Thanks. Anything else? I hope to be able to explore along the line you (and Rick M. today) have suggested. At this point I need to draw together the data already gathered into a report(s) that leaves me with no error signals from my conscience.

Tom Hancock

Date: Thu Jul 22, 1993 6:36 pm PST
Subject: Phonemes; Misc

[From Bill Powers (930722.1730 MDT)] Bruce Nevin (930722.1649 EDT)--

Thanks for info about AG Bell and the vowel triangle. Actually this is quite similar to what I have (not very firmly) in mind: to find degrees of freedom in the sound representation that correspond to natural degrees of freedom in the articulators. The motive is to make the construction of a control system as easy as possible, with as little interference between independent systems as possible. Of course to build such a control system completely, it would be necessary to have an accurate computed or physical model of the vocal tract, as Rick suggested. As I don't have such a thing, I'm just concentrating on possible organizations of input functions.

>Sometime in the 1950s or early 1960s somone noticed that if F1
>is plotted on one axis and F2 on the other of a Cartesian graph
>in two dimensions (with one of them inverted, I forget which),
>the plotted positions of the vowels approximate very closely to
>the articulatory map of the positions of the tongue for those
>vowels.

I use a sort of standard groan that goes ee, ih, eh, aah, ah, and uh (lips held wide and constant). I think this brackets the Valley Girl oo although I don't naturally land on it. On my plots, I can see that the higher resonances move consistently downward, while the lower resonances move upward, meeting at approximately "uh". When you then add lip rounding starting with uh, you get aw, oh, oo, with the higher resonances disappearing and the lower ones moving back downward. The "ee -- uh" sequence seems to entail lowering the back of the tongue from almost closed to as far open as possible. So I agree generally with what you say about the F1/F2 plot. Also, given the basic ee-uh dimension and the lip-rounding aw--oo dimension, you can create two axes with unusual (for me) vowels appearing at various points. For example, the umlaut e results from maintaining the ee and rounding the lips fully as in oo. If you make a diphthong starting with aah (as in cat) and rounding the lips, you get the cockney aaoh. If you start with eh, you get the ultrarefined eh- oh. If you simultaneously run down the ee-uh axis and across the uh-oo axis, you get what my kids used to say about anything yucky: eeeoooo, on a falling note.

I won't get this done for the meeting, but it looks as though one properly weighted sum will give one signal that measures the position along the ee-uh axis and a second one will measure the position along the uh-oo axis. Then particular combinations of these two signals should represent one of the possible vowels. After this works for one voice (mine), I can try it with another (Mary's) to see (a) what changes with size of vocal tract, and (b) what might serve as an indicator of that change, to aid in its removal from the perceptual signals.

Incidentally, there's a difference between my set-of-tuned- filters system and a standard specrtrogram. I'm looking at the amplitude outputs of fixed-frequency detectors, whereas on the usual sound spectrogram, amplitude information is subordinated to frequency information. The two are related, of course, but the treatment of the resulting signals is quite different.

Ray Allis (930722.1426 Pacific)--

Being a mere PC user, with Turbo C 2.0 only, I don't think I could run the software you mention. I'll take a look to see if they have any versions for PCs. Thanks for relaying the info.

Tom Bourbon (personal) --

We're too close to the meeting for me to do much with your paper. No need for me to be a joint author.

Best to all, Bill P.

Date: Thu Jul 22, 1993 7:56 pm PST
Subject: nonPCTers and Closed Loop

from Ed Ford (930721:2055)

On turning nonPCTers into PCTers

I think the real key in trying to get someone to take an interest in PCT is not necessarily done by direct selling of PCT. I have found the most effective way is through finding out those reference signals, that is, the various wants, that are important to the people you want to convince, then show them some practical solutions to their problems, and, finally, how PCT, as the theoretical basis, is tied to the practical applications. Once the connection is made and success is achieved (variable has been controlled successfully), then the person's respect, and perhaps acceptance, of PCT begins to grow.

For example, I spend very little time (about 30 to 45 minutes) explaining PCT at the beginning of every workshop, generally enough to then tie what has been said to some very practical idea they can use when working with kids (or, as I did recently, to those on probation when I did a two-day workshop with the Maricopa County Adult Probation Officers). The more practical applications I can draw out of PCT, the more respect they have for the theory. Some could care less, others begin to ask questions about the theory, and some even want to buy not only my book, but are looking for a more indepth understanding of the concept. Most the school districts where I've worked have either had someone write for copies of BCP, LCS, LCSII, Intro to Psych, etc. or they have bought copies for the teachers' library. I might add that I do expand on my explanation of PCT as the workshop progresses. But I never teach anything about PCT unless I can tie it to something practical. However, I've found there is nothing about PCT that can't be shown as a basis for a practical application to some human concern.

To all CSG members: The July edition of Closed Loop will be sent to all paid up members (1993) about Friday, Aug. 13th.

Best, Ed

Date: Fri Jul 23, 1993 2:24 am PST
Subject: Wealthy Little Man/Little Baby?

From Greg Williams (930723)

The following showed up in my e-mail today:

BEGIN INCLUDED MESSAGE:

```
                       Alife Digest, Number 108
                       Thursday, July 22nd 1993

~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
~                   Artificial Life Distribution List                       ~
~                                                                           ~
~        All submissions for distribution to: alife@cognet.ucla.edu         ~
~ All list subscriber additions, deletions, or administrative details to: ~
~                      alife-request@cognet.ucla.edu                        ~
~          All software, tech reports to Alife depository through           ~
~  anonymous ftp at ftp.cognet.ucla.edu in ~ftp/pub/alife (128.97.50.19)  ~
~                                                                           ~
~              List maintainers: Liane Gabora and Rob Collins               ~
~                   Artificial Life Research Group, UCLA                    ~
~                                                                           ~
~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
[...]
```

From: Hugo de Garis <degaris@hip.atr.co.jp>
Date: Mon, 12 Jul 93 15:11:58 JST
Subject: ECAL93 Report, Hugo de Garis, ATR

Dear ALife Digest,
Here is an ECAL93 report. Its a bit late because I was travelling a lot. Hope you can use it.

Cheers,
Hugo de Garis.

ECAL93 REPORT

Hugo de Garis

Brain Builder Group,
Evolutionary Systems Department,
ATR,
Kyoto, Japan.
degaris@hip.atr.co.jp

Firstly, my apologies for the fact that this report is late. I'm usually pretty quick at getting back these ALife conference reports to the ALife network. However, this time I was caught up for a month on a world trip to Europe and the US for the summer conference season and making deals in the US for our new ALife group in Japan.

ECAL93 (The Second European Conference on Artificial Life) was held at Brussels University, Belgium, Europe, on May 24-26, 1993. It was a big event, with most of the world's ALife big names present, e.g. (active in the US) Langton, Kauffman, Ray, Fontana, Rasmusssen, etc. and (active in Europe) Prigogine, Varela, Nicolis, etc). One of the intentions of the conference (according to one of the organizers whom I spoke to), was to show the Americans that "ALife doesn't all happen at Santa Fe", that there has been a substantial ALife effort in Europe for many years, and that it would be a good thing for the two continents to get together. Well, ECAL93 succeeded in doing that, I felt.

[...]

The next ECAL will probably be held in Eastern Europe in 1995.

Rumor has it (actually I got this from a VERY high ALife source), that Rod Brooks aged 2 years in half an hour while listening to Mark Tilden's "Junk Robot" talk at Sante Fe in 1992. Rod likes to be on the cutting edge, but learned during Tilden's talk that he had become main stream. He has subsequently changed direction and is now wanting to build an "upper-body robot" with hand/eye coordination, capable of cognitive function, using a mid level subsymbolic approach. (I hope I got that right, Cynthia?) I called in at MIT just last week, to visit Toffoli's Cellular Automata Machine group. Our "Brain Builder Group" at ATR, wants to build/evolve an artificial brain inside Toffoli and Margolus's CAM8 cellular automata machine.

END INCLUDED MESSAGE

Maybe RB was inspired by Tom Bourbon's presentation in France in '92?

As ever, Greg

Date: Fri Jul 23, 1993 5:20 am PST
Subject: Mouths as Reality

[From Gary Cziko 930723.0206 UTC] Rick (930722.1230) said to Bruce

>The metaphor is that we perceive articulations. For me, a perception
>is a signal that represents some aspect of external reality:
>
>reality -->|sensory transducers| --> |perceptual function| --> perception
>
>The aspect of reality that is sensed in speech is usually only the acoustic
>signal. The perception of speech is a representation (defined by the
>perceptual function) of the sensed acoustical signal.

When I said that when we perceive vowels we perceive the mouth configuration of the speaker, all I meant to do was to point out a perceptual constancy (like size and color constancy, and melody constancy--a melody sounds basically the same regardless of what key it is played in).

We hear [i] (as in "beet") as [i] even when the relative strength of the overtones are quite different (as when the fundamental frequency is varied). The perceptual system "fills in" where it has to.

This makes it easy to speak since we don't have to worry about the pitch of our voice--at least not for vowel recognition.

Even though I don't think we need to know the actual mouth configurations of a speaker to understand speech (many Americans can hear the difference between French [u] (e.g., roue) and [y] (e.g., rue) and can't say the latter at all), look at what can be done with Rick's schematic:

>reality -->|sensory transducers| --> |perceptual function| --> perception

mouth config-->|ears|-->|auditory networks|--> [i]

Rick, what do you have against considering your interlocutor's mouth configuration as reality?

--Gary

Date: Fri Jul 23, 1993 7:19 am PST
Subject: Re: nonPCTers; Wealthy Little ...

From Tom Bourbon [930723.0910]

>from Ed Ford (930721:2055)
>
>On turning nonPCTers into PCTers
>
>I think the real key in trying to get someone to take an interest
>in PCT is not necessarily done by direct selling of PCT. I have found the
>most effective way is through finding out those reference signals, that
>is, the various wants, that are important to the people you want to
>convince, then show them some practical solutions to their problems,
>and, finally, how PCT, as the theoretical basis, is tied to the
>practical applications. Once the connection is made
>and success is achieved (variable has been controlled successfully),
>then the person's respect, and perhaps acceptance, of PCT begins to
>grow.

Ed, your remarks just now have helped me clarify what I think is the biggest difference between the tasks confronting PCT practitioners, on one hand, and PCT theoretician-modelers, on the other, when it comes to "selling" PCT. In your arena, people know they have problems. Accordingly, they seek and accept ideas which, if presented skillfully and sensitively by a provider (Dag, Diane, David, Dick, you ...), seem likely to solve their problem. I am delighted that PCTish providers like you are able to convince people to "buy the product," and that they and you often believe it works.

On the other hand .... The problems that seem most bothersome to behavioral and social scientists seem to be of a different kind. To get a feel for what I mean, you could pick up a few issues of the journal, *Behavioral and Brain Sciences." We have mentioned it many times on the net. Every issue contains a few "target articles" and "commentaries" on the target article. Sometimes there are as many as twenty-five or more commentaries, ostensibly representing many "viewpoints" and "perspectives" on the topic addressed in the target. That seems to be a likely place to seek out the kinds of problems and conflicts in the behavioral and brain sciences.

Well, the problems people air in B&BS are of many kinds, but with hardly any exceptions, the participants hold the same, generally unarticulated and unanalyzed, assumptions and beliefs about the C-->E nature of behavior. People will debate experimental design, or statistical analysis, or general linear models, or competing architectures for neural networks, or who has priority for having said "blah," .... All the while, lurking unseen in the background, accepted by virtually everyone, is an implicit causal model of organisms and their worlds. Sound familiar? That was the thought Bill and I tried to address in "Models and Their Worlds," but none of the scientists to whom we submitted it ever caught on. We were speaking a different language. We were not talking about their problems. At the level we were addressing, the only message to many people in mainstream behavioral and social science is, "you are talking about a fantasy world." There is no way to sugar coat that message, so we said as much in our paper -- a frankness which earned for us the opportunity to publish in the CSG Ghetto Press -- but a mighty fine ghetto is is!

+++++++++++++++++++

Subject: Wealthy Little Man/Little Baby?

>From Greg Williams (930723)

>>...However, this time I was caught up for a
>>month on a world trip to Europe and the US for the summer conference
>>season and making deals in the US for our new ALife group in Japan.

[TB]
Wow! Do you think these people have any problems we could offer to solve? Brain builders, world junkets on the conference circuit, and much more. I suppose someone must suffer for science.

>> Rumor has it (actually I got this from a VERY high ALife source), that
>>Rod Brooks aged 2 years in half an hour while listening to Mark
>>Tilden's "Junk Robot" talk at Sante Fe in 1992. Rod likes to be on
>>the cutting edge, but learned during Tilden's talk that he had become
>>main stream. He has subsequently changed direction and is now wanting
>>to build an "upper-body robot" with hand/eye coordination, capable of
>>cognitive function, using a mid level subsymbolic approach. (I hope I
>>got that right, Cynthia?) I called in at MIT just last week, to visit
>>Toffoli's Cellular Automata Machine group. Our "Brain Builder Group"
>>at ATR, wants to build/evolve an artificial brain inside Toffoli and
>>Margolus's CAM8 cellular automata machine.

>Maybe RB was inspired by Tom Bourbon's presentation in France in '92?

[TB]
Fat chance! He did watch (I think), out of the corner of an eye, while I ran PCT demonstrations, including Little Man, Crowd-Gather, E. coli, and some "social tracking." As I recall, at the time he and his group were trying to get one of their robots to work, while I was talking and "computering" with some some young artists and a group of people who work on cellular automata. (Hmm.)

Our one public exchange came a couple of days later. Let's just say that his comments at that time were not exactly appreciative of PCT, even though he had not heard my formal presentation.

>As ever,

> Greg

(Shouldn't that be: "As wry and sly as ever"?

Until later, Tom

Date: Fri Jul 23, 1993 9:14 am PST
Subject: Articulators, Applicators

[From Rick Marken (930723.0900)] Gary Cziko (930723.0206 UTC) --

>This makes it easy to speak since we don't have to worry about the pitch of
>our voice--at least not for vowel recognition.

Or for consonant recognition either -- since we can still understand speech when it is sung.

>mouth config-->|ears|-->|auditory networks|--> [i]

>Rick, what do you have against considering your interlocutor's mouth
>configuration as reality?

Nothing. I believe in the reality on mouth configurations; I use them all the time. And I think that the discoveries of the articulatory phoneticians are very important and useful. My objections to your diagram above are based on my modelling perspective. The fact of the matter is that what comes into the ears is an auditory signal -- not a mouth configuration. This signal (that sounds like [i]) could have been produced by something other than a mouth configuration -- like a bowed saw or a digital filter. It is important to know how articulations are typically associated with the acoustical signal because that is useful information for building a control model of speech PRODUCTION. But a production model (according to PCT) depends on having a good perceptual model FIRST; that is, we need to design your "auditory network" that produces a signal corresponding to [i] etc. THEN we can hook the output signal of a control system that is trying to perceive [i], for example, to the appropriate lower level systems (ultimately the articulators) that affect the appropriate acoustical variables and produce the intended output. This is what Bill Powers (930722.1730 MDT) described in his last post when he said that his goal was:

>to find degrees of freedom in the sound representation that
>correspond to natural degrees of freedom in the articulators. The
>motive is to make the construction of a control system as easy as
>possible, with as little interference between independent systems
>as possible.

Ed Ford (930721:2055) --

>I think the real key in trying to get someone to take an interest
>in PCT is not necessarily done by direct selling of PCT. I have found the
>most effective way is through finding out those reference signals, that
>is, the various wants, that are important to the people you want to
>convince, then show them some practical solutions to their problems,
>and, finally, how PCT, as the theoretical basis, is tied to the
>practical applications. Once the connection is made
>and success is achieved (variable has been controlled successfully),
>then the person's respect, and perhaps acceptance, of PCT begins to grow.

Tom Bourbon (930723.0910) replies:

>In your arena, people know they have problems. Accordingly, they seek and
>accept ideas which, if presented skillfully and sensitively by a provider
>(Dag, Diane, David, Dick, you ...), seem likely to solve their problem.

>On the other hand .... The problems that seem most bothersome to
>behavioral and social scientists seem to be of a different kind.

Tom goes on to give a maddeningly accurate description of the problem of getting conventional behavioral and social scientists to even LISTEN to (let alone understand) PCT.

I am sure we will discuss this some more at the meeting but let me try to explain why I am not interested in using the successful solution of practical problems as a basis for "selling" PCT. With Tom, I celebrate the skill and sensitivity of the applied PCTers. I think what they are doing is wonderful; I believe in it and support it. I think PCT can help people and that it can make the world a better place. That is one of the reasons why I am so excited about PCT. It's not just an intellectual exercise for me; as someone said in the movie "ET" -- "This is reality, Greg".
The problem with "successful solutions" is that there are too many of them. Behavior mod works for some people; reality therapy works for others; harai krishna works for some; primal scream for others. Not only are there too many successful solutions, there are rarely any failures. Apparent failures can always be attributed to the fact that the person "just didn't get it". So if there was one person (and there must have been ONE) who was not helped by PCT (an unsuccessful application) then this can easily be written off; the person just never really understood PCT, or was resistant to it, etc. But the same can be said for the failures of other applications -- the behavior mod program didn't work because it was not applied properly, the reality therapy didn't work because the therapy was not done properly, the harai krisna chanting didn't work because the person didn't believe hard enough, the primal scream didn't work because the person didn't scream loud enough.

So, again, while I think the application of PCT is extremely worthwhile and valuable, it is not interesting to me as a basis for "converting nonPCTers". I have no doubt that it works -- that you can convert people to PCT with successful applications. But I feel that many (not all) of the "conversions" based on successful application of PCT will be fairly superficial; PCT made them feel good so they go with it. But maybe I'm wrong; maybe the "successful applications" converts are just as deep as the "scientific testing" converts (the latter group being MUCH smaller than the former). If so, great. What I personally prefer are converts who really understand the PCT model so that they can accurately communicate it. But maybe that's what you get with people converted by successful applications. I dunno. I guess I just prefer the science and, thus, would rather "convert" the scientific nonPCTers.

Best Rick

Date: Fri Jul 23, 1993 9:22 am PST
Subject: Correction

[From Rick Marken (930723.1000)]

I said (to Gary "The Body" Cziko) --

>THEN we can hook the output signal of a control system that is trying to

>perceive [i], for example, to the appropriate lower level systems
>(ultimately the articulators) that affect the appropriate acoustical
>variables and produce the intended output.

Intended output??? Did I say THAT??

I meant INPUT. The control system produces an intended INPUT.

Geez. Bodes ill for the figure contest.

Best Rick

Date: Fri Jul 23, 1993 9:35 am PST
Subject: Re: Correction

From Tom Bourbon [930723.1218] Rick Marken (930723.1000)]

>I meant INPUT. The control system produces an intended INPUT.
>
>Geez.

Now you've done it! From now on, in perpetuity, the world will be plagued with confusion over exactly what it is that a control system controls! See what happens when you start to worry about your public image and about selling your ideas? You start you talk like everyone else.

>Bodes ill for the figure contest.

Was that supposed to be "bods" ill?

Later, Tom

Date: Fri Jul 23, 1993 9:39 am PST
Subject: Memory and modeling

[From Tom Hancock (930723.1000)]

Your posts are encouraging me to keep pursuing this line of work. Thanks. I wish I could interact with you some more on your post of the 19th and then on the one of the 22nd.

Marken (930719.1230)
>Based on what I know of your research, let me just suggest that
>what you might be interested in is the subject's ability to control
>the relationship between a confidence rating and an association
>they make between a configuration and a name. Is this right?

Yes, perhaps. I have thought that that relationship manifests itself in two components: 1. the previous associates are compared to the present time perception for the best match, 2. the degree of the match is indicated by the confidence rating. What do you think of this? I am not sure if we would say in that case that confidence is controlled. Unless perhaps the confidence is a concomitant measure of the control by a higher level system which determines the degree (or vividness) of the match which is needed.

>For example, the subject is shown a configuration and gives a name. This
>name is either right or wrong -- something that the subject might
>know if he or she has previously been told the name that goes with
>the configuration.

Right, they know it with varying degrees of certainty--which do correlate with actual performance, with most subjects being slightly overconfident.

>Anyway, the subject's response to the configuraton
>can either be right (R) or wrong (W). The subject is then to give a
>rating of his or her certainty that the response is R or W. The con-
>trolled variable might be some measure of the association between
>a binary variable -- the subjects response (R or W) -- and a continuous
>variable -- the subject's rating of certainly (0.0 to 1.0?).

In light of your subsequent post I wonder if you still think this may be true. The possibility is plausible and intriguing to me.

>The problem with this controlled variable is that part of it is imagined -
>- the subject's perception of whether his/her answer is R or W. Subjects
>can't control the relationship between their rating and whether they were
>"really" R or W, can they?

My understanding is that subjects can increase the likelihood of being right, by engaging in more memory search, screen scanning or mnemonic building. A subject who does so should also have a higher level control for more distinct associations or for increasing the likelihood of correct responses.

(Rick Marken (930722.0900)

>I think part of the problem here is that you are studying something
>that we haven't studied much in PCT -- memory. I think the perceptions
>that are actually controlled in your experiment are imaginations --
>replays of previously experienced perceptions.

Right.

>What you are trying to figure out is something I don't understand at all,
>which is "how does a person know that an imagined perception
>corresponds to a previously experienced perception?"

I have thought that the way a person knows is by the error signal. If there is a large error signal in the match of the imagined perception to the reference then he will report not being certain. But if possible matches have been accessed and the signals from the incorrect ones have been completely suppressed and there remains a good match, then the subject knows (or thinks he knows) that he has found the right one.

>The associations you are talking about as perceptions are perceptions --
>the subject can perceive that after seeing X you say Y. p* would be the
>reference for that perception; at first the subject may have no reference
>for the association -- s/he doesn't care what comes after X. But you have
>asked thats/he learn to have a reference for this association that is the
>same what s/he experienced earlier. So after perceiving X-Y s/he is
>supposed to store a reference value, p*, of X-Y . The next times/he sees X
>s/he should produce Y to match my reference. So s/he can control the
>perception of the association by producing Y (instead of Z,
>etc) when you show him/her X.

Yes, that is the way I see it with one addition. The subject does not just say Y automatically (at least at the learning stage). The subject mediates between the associates--typically by building some verbal category connections. So when X is perceived, S searches for the bridge or mnemonic. When it is found and other similar mnemonics are rejected (signals inhibited), then the S says Y.

>Now it's pretty well known that people seem to have trouble storing the
>appropriate reference levels for associations, especially when there
>are several different references to be learned. What is also apparent
>is that people seem to have some idea of whether or not they have
>stored the appropriate reference for an association. This is the part I
>don't understand

The research with neural networks (e.g. Rumelhart) bears certain similarities to PCT and gives us a hint at understanding this: a disturbance is sensed (a new item), it perturbs the network of memories, constraints (reference signals?) to the activation of the memory network will affect what continues to be active, until a pattern of activation remains, which pattern is the subject's interpretation. My thought is that the activation for that interpretation that remains is monitored and if there are still some persisting signals from other interpretations then the subject may be less sure about the appropriate reference, and vice- versa.

>There are probably single subject studies of memory
>that might have the kind of data you could use as a start at building a
>PCT model of this process.

Your PCT schema if more fully developed than mine (you have many more constraints guiding your imagination to PCT-correct interpretations), so do you have any thoughts about what kind of data I would be looking for?

>But here is a suggested version of your study
>(that could easily be done on a computer) that would give you the
>kind of data that might provide a basis for modelling...

I like your suggestion and I believe I will write a proposal partly along that line for continuation funding from the Air Force. Maybe I could send you a copy--if you are not getting overloaded by me.

>So my suggestion is to simplify the task as much as possible, collect
>data from one subject (to contain the impulse to average over subjects;

>once you get perfect data from one subject you can try your technique
>on another subject, and then another ...). Collect at least some good
>amount of data when the subject is clearly in control (as when the
>subject is providing the right associate--and a 10 rating -- every time

Rick, I will propose to do as you have said--but at present I do have data where this has already been done. It is not from a methodology which is as simple and clear cut as you have suggested. But I have looked at response patterns for one subject who has been consistently correct with one item. My problem has been interpreting this in terms of PCT

>So, for one association over trials the
>sequence of right (R) and wrong (W) associations along with the
>corresponding confidence ratings might look like this:
>
>W (1) R (7) W (2) W (1) R (8) R (6) W (2) R (10) R (10) R (10) R(10)
>
>Now see what the relationship between right/wrong and confidence
>ratings over trials looks like for each associate.

*****
Your hypothetical data appears similar to what the subjects actually do. In addition I have been trying to include the response time, and post-instructiona l information time which is associated with each trial. But so far I am not sure what how it all fits with a precisely articulated model. As presented in earlier posts, the response times decrease as certainty increases; the post-instructional information study time increases when there has been an incorrect and when certainty decreases; and as the subject tends to get all corrects and certainty ratings of 100% then the response time and the study time reach a minimum. In my attempt at modeling, I have considered that these minimum times, along with 100% certainty are indicative of that subject's reference for that item. Yes, this reference seems to be similar within subjects and between high level correct items.

Rick, in light of our discussion to this point, do you have any additions (or subtractions) on what the dynamics of the model should be like?

I am still wondering what are some of the basic PCT principles on varied response times? Such as: crossing more levels takes more time, or if there is poor control it takes more time, or reorganization at higher levels takes more time?

>the subjects will probably be pretty good at knowing whether their
>answers are right or not. How do they know that? That was you original
>question, wasn't it?

Yes, sort of. I was questioning how one might model a situation where there is understanding and where there is varying degrees of understanding or perceived meaning. To me, this is an extension of my research concerns into the important everyday concerns of humans.

Tom Hancock

Date: Fri Jul 23, 1993 10:06 am PST
Subject: CSG -- Here I come.

[From Rick Marken (930723.1100)] Tom Bourbon (930723.1218) --

>Now you've done it! From now on, in perpetuity, the world will be plagued
>with confusion over exactly what it is that a control system controls!

I knew it! Drat!

Ah well. What does it matter anyway? Control input? Control output? Big deal.

Best Rick

Date: Fri Jul 23, 1993 12:13 pm PST
Subject: real/imagined mouths

[From: Bruce Nevin (Fri 930723 15:34:08 EDT)] Rick Marken (930722.1230)

Gary's question is going unanswered, so I'll send this without change. For starters, here's the same question in different words:

> The aspect of reality that is sensed in speech is usually only the acoustic
> signal. The perception of speech is a representation (defined by the
> perceptual function) of the sensed acoustical signal.

Why confine yourself to the point of view of a hearer who has never spoken?

> >I think what is involved is something like this general-
> >purpose capacity for intuiting "what could possibly have made that
> >sound?"

> This may be true for you (a trained linguist) but it sure ain't true for me.
> When I hear a phomene it's just a phoneme -- there is absolutely no
> associated imagery (visual or kinesthetic).

Rick, are there not many perceptions of which ordinarily you have no conscious awareness? And is it not the case that you can become aware of such perceptions by practice at directing your attention (training)? The idea is that this is a specialization of the same general-purpose capacity. In other words, I am arguing (again! still!) that there is nothing special about language, and that it can be modelled in the same ways as other things that are easier for us to understand. How would you model the case of recognizing what fell in the kitchen? Suppose you had only 40 objects in the kitchen that the cat could possibly knock off the counter, that each made characteristically different sounds (though if they fall a short way or softly it's harder to tell them apart), that you had a number of years of daily practice guessing what had fallen and getting indirect confirmation or disconfirmation of most of your guesses, all your waking hours every day, and that guessing right mattered a lot to you because doing so served many many other purposes for you. Phonemes are like that.

> But once you've learned the regular relationship between phonemes
> and articulations then the sounds probably do start "sounding like"
> articulation patterns -- becuase you can imagine the articulations.

Rick, you don't have to have training as a linguist or a phonetician to learn the relationship between speech sounds and what you have to do to produce them. The training is about naming them, categorizing them, representing them by funny marks on paper, and so on, but none of that is necessary for modelling what speakers and hearers do unconsciously. A child who talks to you has indeed "learned the regular relationship between phonemes and articulations", or she would not be able to pronounce words so that you could recognize them.

> In general, I think a control system must develop the ability to
> perceive phonemes (ie. it must develop the perceptual functions that
> produce outputs that correspond to the presence of /d/ /b/ and /g/ etc)
> before it can learn how to vary the references for the lower level
> articulatory perceptions that produce those phoneme.

This surprises me. It does because I endorse and agree with your proposed slogan

>"Learning to control begins with learning to perceive"

And perceiving speech includes kinesthetic perception for the speaker.

My understanding of what infants do in babbling is that they are learning to control their perceptions of the sounds that they make. This trial and error process continues through all the stages of learning to speak and understand. The child finds that what she perceived as a repetition of some word older people use is not perceived as a repetition by them. There's some difference between "yes" and "less" that she doesn't yet control. This results in recurrent error until they reorganize their system of speech sound-differences so as to distinguish such words in a way that adults recognize. The result is a system that has more phonemes, perhaps differently organized, after the reorganizing. It seems to me pretty likely that the child procedes by trial and error with her own pronunciations until she hits on something that no longer produces error in her interactions older people.

It turns out that many of the differences between consonants correspond to easily differentiated places of contact (or near-contact) between the tongue and parts of the oral cavity. It also turns out that these places are areas within which variation in the position and contour of the tongue make little acoustic difference. For these and other reasons I believe that the child learns to control a perception of touching the tongue to one certain part of her mouth or to another as the way of distinguishing between /t/ and /k/. My experiment with masking noise did not run long enough but suggested that these kinesthetic perceptions provide an "anchor" that vowels lack when acoustic perception is lost.

Gary has talked about how even the formants for vowels, which we might think of as pretty stable and reliable acoustic cues, cannot be directly perceived if there happen not to be any harmonics at the frequencies of the formants. (Recall that the vocal tract filters out all harmonics except in the bands that we call formants.) For consonants all of the acoustic features that have been identified have turned out to be dispensible, for any given cue you can recognize the consonant without it. Put this together with the visual cues overriding the acoustic cues in the experiment I recounted, and you get the idea that the hearer usually gets the vowels fine, and grasps at any kind of available evidence in order to determine the consonants. For the speaker, the most reliable perceptual differences for the t-ch-k contrast are those nice, safe, stable, error-resistant places for the tongue to touch. The same differentiation of tongue places correlates with the d-j-g, n-ng, s-sh-x (x as in "Bach") contrasts. When we produce those sounds, the kinesthetic perceptions of placing the tongue in the right areas is among the perceptual cues that we got it right. These perceptions are much much easier to control than the acoustic perceptions associated with those consonants. The acoustic perceptions are difficult to construe as a "consistent result" produced by "variable means". When we hear a word (or a nonsense syllable), why should our memory and imagination of "what could possibly make that sound" not be relevant?

For other phonemic distinctions, this is not so necessary, because there are much more obvious and reliable acoustic cues. Thus for nasalization (b-m), the sounds of spirants and fricatives (s, sh, x, th, and their voiced counterparts), and, as you mentioned, voice onset time (VOT) differentiating voiced-voiceless (b-p). Chinchillas can be trained to differentiate VOT in experimental situations. Possibly they could be trained to differentiate t-k in clearly articulated syllables, that is, in a situation providing a stable, reliable, invariant acoustic cue in the "chirp" between consonant and vowel. A stable VOT differentiation is available in most of ordinary spoken English; there is no such reliable invariant for t-k (or p-t, if you can't see the lips) in the world outside the laboratory. I question whether a control system can be trained to distinguish the distinctions b-d-g as they occur in played-back recordings of ordinary speaking, if that control system lacked the means to produce the same distinctions and some kind of test that their renditions are correctly perceived by other speakers of the language. So there's an empirical test. So far as I know, no-one has checked it out. Maybe Martin knows of some work like this.

Now it could be argued (as Bill has) that we resolve an indistinct consonant at higher levels of word choices and their meanings, like a fill-in-the-blanks letter puzzle. This does not account for how people perceive nonsense syllables (without meanings) in experiments where what had been thought of as the acoustic cues for the consonants were not present. Such experiments indicate that there is stuff going on at lower levels. I suggest only that the "what could have made that sound?" capacity to imagine oneself repeating what is heard may play a part.

> You can't control the perception of an acoustical variable until you
> can perceive that variable. So the child's articulatory apparatus
> may be perfectly capable of articulating the difference between /p/
> and /b/ (in terms of voice onset, right?) but the child cannot control

> that difference until it can perceive the difference between /p/ and
> /b/ in terms of how they sound. So it can't recognize /p/ and /b/ in
> terms of articulation (at least when it first learns them -- the arti-
> culatory perceptions may be incorporated into the recognition
> process later, once the child has learned to control what it can already
> perceive) so articulatory perceptions cannot be essential to
> recognizing speech (though they may become involved after one
> has learned to control speech perceptions).

An infant does not learn /p/ and /b/. It learns to approximate words. Control of the sound-distinctions develops only as it becomes necessary to differentiate between words that had been approximated in the same way. Then those sound-distinctions start to be controlled perceptions in their own right, and all sounds perceived as intended as speech are mapped on to them, however ill or well they fit. Hence foreign accents. In languages of India there is a 3-way VOT distinction,
b-p-p^h , where b and p^h are like English b and p, and intermediate p
is like the second consonant in "spell"--neither, but somewhat like
h
Spanish p (though not so tense). Actually, they differentiate their b and p by greater VOT difference than that by which we differentiate our b and p, creating "space" for the intermediate, unaspirated p. So when they speak English, their b sounds a little too fully voiced, and their p sounds extra strongly aspirated. An infant in an English-speaking family may make all of these sounds, and the range between them, without differentiation. Then it turns out that control of that heard delay before voicing of a vowel begins after a consonant, and control of devoicing between a vowel and the silence of a following stop consonant, serves to make only one distinction common to many words. Every b/p-like sound is categorized on one side or the other of this distinction. The child makes sounds far enough to one extreme or the other of the VOT range, or she is not understood. Has the child learned p and b? No, this distinction applies concurrently with many other distinctions to make what we with our alphabetic presuppositions interpret as p-b, t-d, s-z, and so on, so you could say that by learning to control VOT in the English-speaking way she has learned all those phonemes at once, except of course she has to control other perceptions so as to differentiate p-t, etc.

(Bill Powers (930722.1730 MDT)) --

Sounds good (so to speak). Your "ee" with lips rounded is what phoneticians write [y], the sound in French tu (or rue) or German u-umlaut. If you produce i-y-u (ee, ee with rounding, oo, changing smoothly) you can actually hear F2 descending toward a constant F1. Sounds like some sort of Sci-Fi movie sound effect. These things are easier to hear with your ears plugged with water, say in the bathtub. Family members may think strange things, but it's too late for that. We're long since over the brink.

Working with relative amplitude signals from overlapping fixed-frequency filters is kind of like the eyeglasses for lip reading. It is very similar to a sound spectrograph. The original sound spectrograph runs a narrow filter variably down from the top of the range to the bottom, synchronized with a tape loop of the speech sample. On each pass, the filter output burns a trace in sensitive paper. The same could be done with a battery of narrow-pass filters covering the range, and something like this is what Martin described as happening in the ear. Seems to me that the differences are in the number and width of the filters and in their overlapping in your system (and in the ear, I think?), or those are the significant differences. The darkness of the trace on the paper is the amplitude. Both systems display both amplitude (is there sound energy from an unfiltered harmonic present?) and frequency. Am I right?

It's terrific, what you're doing. I'm disappointed that I won't be there to play with it with you. I've met Chuck and Clark and Avery, I wish I could join all of you and meet more of you in the flesh. Hopefully next year.

Too much. Gotta run.

Bruce bn@bbn.com

Date: Fri Jul 23, 1993 12:59 pm PST
Subject: Sound Spect

[From Bill Powers (930723.1430 MDT)] Bruce Nevin (930723.1534) --

>Working with relative amplitude signals from overlapping fixed-
>frequency filters is kind of like the eyeglasses for lip
>reading. It is very similar to a sound spectrograph. The
>original sound spectrograph runs a narrow filter variably down
>from the top of the range to the bottom, synchronized with a
>tape loop of the speech sample. On each pass, the filter
>output burns a trace in sensitive paper. The same could be
>done with a battery of narrow-pass filters covering the range,
>and something like this is what Martin described as happening
>in the ear. Seems to me that the differences are in the number
>and width of the filters and in their overlapping in your
>system (and in the ear, I think?), or those are the significant
>differences. The darkness of the trace on the paper is the
>amplitude. Both systems display both amplitude (is there sound
>energy from an unfiltered harmonic present?) and frequency. Am I right?

My sound spectrograph is just as you describe, except that it uses a battery of (about 100) programmed filters and displays the amplitudes at each frequency and time as a degree of white. My fixed-filter thingy is just another spectrograph, but with only a few frequencies (10 at the moment), and the display is like a series of 10 oscilloscope tracings showing amplitude as deflection above a line. The latter gives a lot more information about the sound signal -- not just "present or absent", but HOW MUCH sound energy is present at each instant.

>It's terrific, what you're doing.

That remains to be seen. It may be a big waste of time, converging right back to what's already being done. Unlikely that I'm the first person to play around with these variations. Maybe by next year's meeting we'll be able to noodle around with it together.

Best, Bill P.

Date: Fri Jul 23, 1993 3:29 pm PST
Subject: Practical solutions

[From Dag Forssell (930723 1520)] Rick Marken (930723.0900)

>..I am not interested in using the successful solution of
>practical problems as a basis for "selling" PCT.

How about "theoretical problems" or just "problems."

Problems = error signals -- existing or created by you (good disturbance??) are the only way to get the attention of another control system, right?

>...What I personally prefer are converts who really understand the
>PCT model so that they can accurately communicate it. But maybe
>that's what you get with people converted by successful
>applications. I dunno. I guess I just prefer the science and,
>thus, would rather "convert" the scientific nonPCTers.

"people converted by successful applications" How about people converted by successful resolution of their problems - whether practical, theoretical or otherwise.

Your own conversion to PCT must have resolved some problem you had. Why else would you have been curious enough to read the book?

Rick, I think you are denying the laws of nature as PCT describes them, when you try to get people without first testing what their "problems" are. I can sympathise with your denial of what PCT teaches us, but I think you are kidding yourself. How long are you going to bloody your nose trying before you recognize that you are dealing with autonomous living control systems? I'll confess I have bloodied mine enough, but slowly I am wising up.

Best, Dag

Date: Sun Jul 25, 1993 4:57 pm PST
Subject: TESTS OF CONTROL - RKC

FROM: Bob Clark (930725.08:25 pm EDT) Bill Powers (930720.0815 MDT)

SOME vs EVERY -- at least that's "agreed." Good.

CONTROL -- SYSTEM; VARIABLE

WTP asks and answers:

W>What constitutes a control system? It is a system in which

W>a. Physical actions are regularly related to informational inputs
W>from the environment. An "informational" input is one that
W>affects signals inside the system in a unidirectional way. A
W>"signal" is a low-energy variable that can alter the states of
W>high-energy variables (neural or chemical signal, enzyme).

W>b. The informational inputs depend directly and continuously on
W>effects of the same physical actions, as well as on independent
W>influences.

W>c. The gain around this loop is substantially more negative than
W>-1 at all frequencies below some finite limit.

In a: By "Physical actions" I think you mean changes in variables that exist in the environment of the system. And that these variables are, at least in principle, subject to detection and measurement by physical instruments. Could one of these "external variables" (is that acceptable terminology?) be considered "controlled variables" if it passes The Test?

In c: You refer to "this loop" without specifying the nature of the "loop." Perhaps "loop" should be specifically defined. You have also omitted the requirements for a loop to be "closed."

Of course I know what you mean in general, but some of the "Fly Trap" discussion suggests that these words -- "closed loop" -- as used in control theory, may not be entirely clear to some people.

Again in c: You refer to the "gain around this loop" without indicating that it must be "power gain" rather than "force" or "voltage" or other equivalent. I think this has come up in other situations, and I think that you agree with "power" vs "force." I make this suggestion to improve clarity.

W>In order for an organism NOT to be a control system in any
W>regard, the only existing closed-loop relationships similar to
W>those above would have to have very low or positive loop gains.

True, but could it be that some organism (presumable very simple in composition) might have no closed-loop systems at all? Seems improbable, doesn't it. Perhaps someone will be interested in performing the required experiments.

W>I think it would be impossible to find any living organism in
W>which one or more closed loops with negative gain cannot be found.

This seems to be an arbitrary assumption. It may well be valid. But I see no reason to include this assumption. More useful, I think, is to leave the question open, for experiment to reveal. The implied experimental studies of variables that may turn out to be "controlled variables" could well be interesting in many ways.

W>So the only real question remaining is the magnitude and

W>sign of the loop gain involved.

The remainder of your paragraph is no problem, assuming that the "gain" is a "power gain," and that the nature of the controlled variable has been identified. The latter may be a "key" question, however. And, of course, the failure to verify one proposed controlled variable does not eliminate other possible candidates.

W>I repeat it because you occasionally forget it, or at least fail to apply it:

Unfortunately, Bill, I sometimes fail to make myself clear. Here you were referring to my remarks about phototropism, observing that many plants do not show this effect. My concern here was still with application of the word, EVERY.

W>It is not logical to use the lack of one type of control system
W>to "prove" that some plants contain no control systems at all.

Surely, Bill, you did not think this was my argument? Rather, my point was that proving that some plants DO have control systems does not prove that they ALL do. And I have stated that I accept that some plants have such systems, and so do many more. But, if that is an important matter, experiment is needed. It is not, in my opinion, sufficient to point out the possibility that a given variable could be controlled variable. Over the years a lot of experimental data has been gathered that could provide a useful guide for experimental search for controlled variables and their associated systems.

B>>For a control system, presence of a closed loop is a "necessary
B>>condition," but not a "sufficient condition."

W>The "sufficient condition" is completed by specifying that the
W>loop gain is greater than 1 and negative.

Yes, but this needed to be pointed out. Perhaps some such statement should have been included, but it seemed unnecessary at the time.

Rick's post, Rick Marken (930718.1800), notes the negative sign, and the eventual return to the original state -- after the disturbance is removed, eaten up or whatever! The original state is some kind of "equilibrium condition." Such a "condition" would include the continued operation of any control systems it may have. In addition, the "physical action" required in "a." above should be specified and any opposing output actions should be detectable during the test disturbance. To say it another way, the time needed for the opposing action to occur and be detected should be short compared to the observation time.

W>In the case of the Venus Flytrap, I am not convinced that the
W>trigger response is a closed loop in itself, because the
W>immediate effect on the touch comes too late to affect the
W>trigger response.

Your discussion in terms of a hypothetical output function suggests variables that might be considered for experiment.

Perhaps this is a good place to point out that: I really don't much care whether the Fly Trap action involves a controlled variable or not. There are other variables in that plant that might turn out to be controlled variables. I have more interest in the experimental procedure required to establish the presence/absence of controlled variables and control systems.

The people studying plants have been primarily concerned with improving their productivity, crops -- or destroying them, weeds. They have reached many of these goals without need for control systems concepts. However it is conceivable that further gains could be achieved if control system concepts and methods were applied. As long as they reach their goals, why should they bother about ideas that they find diffcult, esoteric and unnecessary? If we want to affect them, we must show them that we can help them improve their results.

W>Plants operate on a much slower time-scale than animals.

Certainly the case for many plant systems. The relatively rapid action in trapping the insect suggests that this might be an unexpectedly fast system. Experimental study of prospective controlled variables would be appropriate.

W>I'm not in a position to state the loop gain of this control system,
W>but I'll wager that it is both negative and large.

The "wagering argument" is irrefutable. It must either be accepted as proof -- or ignored.

B>>I suspect that seeming control systems are, in fact, "balance
B>>of force" systems (like the "ball in a bowl"), or "one way
B>>systems" tending to bring the "controlled variable" to an extreme value.

W>I couldn't disagree more. I don't believe that any such systems
W>exist in organisms except as components of control systems:

The "belief argument" is equivalent to the "wagering argument."

Both arguments can only be resolved by suitable experiment.

W>You're offering the traditional concepts developed before anyone
W>knew of control systems.

Is it NECESSARILY the case that ALL of them are wrong, incorrect, incomplete, or what? As long as the results are acceptable, why should these ideas be rejected by those who are using them? Of course, we think that application of Perceptual Control Theory would lead to better results. Fine -- that is for experimental demonstration.

W>All phenomena explained in such traditional ways need serious re-investigation.

Yes, certainly. Until the results start coming in, I'll postpone decision on these questions.

Thanks for your response to my request for a repeat of the "controlled variable" questions.

I now have them.

It would be interesting to apply them carefully to the Fly Trap.

VARIABLE (engineer/physicist/mathematical version)
vs
VARIABLE (population sample, statistical inference, null hypothesis)

The term "variable" used throughout Feedback Control Theory is the engineer/physicist/mathematical version. When I was working with McFarland I had occasion to learn about the version used in statistics. I think the reason I found this very difficult was because of the drastically different treatment of "variables." To me, this was a very different and unfamiliar use of "variables."

With the e/p/m version, variables are conceived as having continuing identity combined with variations in magnitude. Frequently, but not always, the magnitude is considered to vary with time as an independent variable. If relations among these variables exist, they are expected to be invariant during the experiment, even though the variables themselves may vary. A relatively small amount of variation in the relationship is expected, usually attributed to limited accuracy of measurement. Statistical methods are sometimes used to improve the effective accuracy of the measurment.

The variables used in statistics are population characteristics that remain constant while data are being gathered. Relations among these variables are established by calculating correlatio■_H @Ç I think that a lot of communication problems result from this "mixed use" of "variable."

See you Tuesday! Bob Clark

Date: Sun Jul 25, 1993 8:00 pm PST
Subject: Re: dudidu

When you get tired of communing with your articulators, there is a very fascinating treatment of all this in a PhD thesis at CCRMA (Stanford's computer music lab). The SPASM system implements a fairly complete model of the vocal tract. It runs on a NeXT, in real-time on the outboard DSP chip. I attended a talk by the author, Perry Cook.

```
ccrma-ftp.stanford.edu:
        pub/Theses/PRCThesis.ps.Z
                PhD thesis,
        pub/SPASM.tar
```

software

It doesn't speak. It sings.

--
Lance Norskog thinman@netcom.com
Data is not information is not knowledge is not wisdom.

Date: Mon Jul 26, 1993 3:41 am PST
Subject: expectations

From Bruce Bill Powers (930723.1430 MDT)

> >It's terrific, what you're doing.

> That remains to be seen. It may be a big waste of time,
> converging right back to what's already being done. Unlikely that
> I'm the first person to play around with these variations.

What's terrific (my perception) is that you are experiencing directly what some of the variables are. We'll be able to carry discussion forward from shared referents. And even if techniques and identified variables did converge with other people's findings, that would hardly be a waste of time. Convergence of independent lines of research is confirmatory, and just because of the PCT perspective and modelling a different account will emerge. But I think it is likely that you will find something new to pull out of the acoustic/articulatory hat.

> Maybe by next year's meeting we'll be able to noodle around with it together.

That's my hope.

Bruce bn@bbn.com

Date: Mon Jul 26, 1993 6:16 am PST
Subject: Variables and Methods

[from Gary Cziko 930726.1340 UTC]

Bob Clark (930725.08:25 pm EDT) observed:

>VARIABLE (engineer/physicist/mathematical version)
>vs
>VARIABLE (population sample, statistical inference, null hypothesis)
>
>The term "variable" used throughout Feedback Control Theory is the
>engineer/physicist/mathematical version. When I was working with
>McFarland I had occasion to learn about the version used in statistics.
>I think the reason I found this very difficult was because of the
>drastically different treatment of "variables." To me, this was a very

>different and unfamiliar use of "variables."
>
>With the e/p/m version, variables are conceived as having continuing
>identity combined with variations in magnitude. Frequently, but not
>always, the magnitude is considered to vary with time as an independent
>variable. If relations among these variables exist, they are expected
>to be invariant during the experiment, even though the variables
>themselves may vary. A relatively small amount of variation in the
>relationship is expected, usually attributed to limited accuracy of
>measurement. Statistical methods are sometimes used to improve the
>effective accuracy of the measurment.
>
>The variables used in statistics are population characteristics that
>remain constant while data are being gathered. Relations among these
>variables are established by calculating correlation coefficients and/or
>calculating probabilities that the Null Hypothesis for the proposed
>relationship is invalid. It is expected, and usually found, that a
>relatively large amount of unexplained, "random," "noise" is involved.
>
>I think that a lot of communication problems result from this "mixed
>use" of "variable."

I agree. You are pointing out a major argument made by Philip Runkel in his 1990 book _Casting Nets and Testing Specimens: Two Grand Methods of Psychology_. He usefully refers to the "population" use of variables as the method of relative frequencies ("casting nets") while the "engineering/physics/mathematic" use of variables is testing specimens.

Another related distinction made in the social/behavioral sciences is that between "cross-sectional" and "repeated-measures" designs. Cross-sectional refers to a variable varying over a population or sample of individuals (relative freqencies). Repeated measures refers to observing the same same individuals under at least two conditions.

However, repeated-measures data is still usually analyzed using group statistics in which all traces of individuals are washed away. Now when I am on a dis-sertation committee I insist that if repeated measures have been used, that we look at each individual. I argue that a statistical test can reveal that subjects perform and yet 55% of the subjects actually did better under condition A!

Some "quantitative" researchers I know of have indepepently (of PCT) come to the conclusion that even repeated measures are pretty useless since the data is much too lumpy (like trying to figure out the story of "Gone With the Wind" by seeing only the opening and closing frames of the film). Bob Siegler of Carnegie-Mellon uses what he calls the "microgenetic" method to examine how children discover and use new strategies for solving math problems. It involves continuous study of children while they grapple with math problems. It is interesting that his microgenetic method of children's reorganization have lead him to the conclusion that there are "evolutionary" components of variation and selection involved in such reorganization. He compares traditional repeated-designs methods for investigating children's cognitive development as "child can do A; then a miracle occurs; child can now do B."

I have found that while I have a very tough time selling PCT to colleagues (for all the reasons that Dag Forssell and Rick Marken have pointed out plus others), I CAN get some people to follow arguments about research methods. The arguments lead from (a) "independent" groups research designs to repeated measures to (b) individuals within repeated measures to (c) increasing the number of repeated measures (along with reducing the group size) to finally (d) the method of specimens that PCT uses. --Gary

Date: Mon Jul 26, 1993 9:18 am PST
Subject: Re: Articulators, Applicators

>From Oded Maler (930723) Rick Marken (930723.0900)

* The problem with "successful solutions" is that there are too many of them.
* Behavior mod works for some people; reality therapy works for others;
* harai krishna works for some; primal scream for others.

But Truth is One!

(and this works for others.. :-)

--Oded

Date: Wed Jul 28, 1993 5:30 pm PST
Subject: two notices

[Avery Andrews]

2 indications of what others are up to:

From @CSLI.Stanford.EDU:bug-sdl-request@weber.ucsd.edu Thu Jul 29 07:32:15 1993
Date: Wed, 28 Jul 93 16:13:26 -0400
From: steve%cougar@gte.com (Steven D. Whitehead)
To: bug-sdl@weber.ucsd.edu
Subject: CFP AAAI Spring Symposium

PRELIMINARY
CALL FOR PARTICIPATION
"Toward Physical Interaction and Manipulation"
AAAI SPRING SYMPOSIUM SERIES
STANFORD CALIFORNIA
MARCH 21-23, 1994

We are delighted to invite contributions for the 1994 AAAI Spring symposium: "Toward physical interaction and manipulation" to be held on the campus of Stanford University on March 21-23, 1994.

SYMPOSIUM DESCRIPTION:

The range and scope of practical robotics applications depends critically on the ability of robots to physically interact with their environments. Current applications are highly specialized, and typically they involve carefully controlled, well understood workspaces with little or no sensory feedback. Construction costs and inflexibility limit the economic viability of these systems. The general manipulation skills of humans and other animals contrasts starkly with the current capabilities of robots. From threading a needle, to opening a door, to catching a ball, to moving a sofa, we engage our environments in myriad ways. Unlike most current robots, we rely upon rich sources of sensory feedback to cope with uncertainties in our varied world.

The purpose of this workshop is to draw together researchers from a range of disciplines to study the principles of physical interaction and manipulation. The goal is to consider theories, paradigms, and ontologies for both natural and artificial systems, and to develop generally useful concepts, architectures, and algorithms for building and describing them.

The approach is to select in advance a set of tasks that range in difficulty and span a number of research issues. Each prospective participant is to develop conceptual designs for one or more of these tasks prior to the workshop. It is acceptable for designs to be speculative, as we encourage creative solutions. However, the aim is to examine tasks in detail and sketch complete systems. At the workshop, selected designs will be presented, discussed, and compared in an attempt to reach a more general understanding. By analyzing a range of tasks, we aim to broaden our perspective, identifying common themes and useful design principles. The rationale for this format is that participants will be well prepared for the discussions by thinking in detail about some of these tasks in advance. The list of candidate tasks follows:

- make a cup of coffee
- fry and serve an egg
- prepare buttered toast
- play catch
- insert and play a video tape
- vacuum/mop the floor or mow the lawn
- dig a hole/trench
- (un)lock a door with a key
- open, pass through, and close a door
- feed someone using a fork, knife, spoon, cup, etc.
- retrieve a screwdriver from the toolbox in the garage.
- fold clothes
- move large objects (boxes, chairs, furniture)

These activities involve a range of skills and will most likely require a range of mechanisms. They can be characterized by their requirements for:

- real-time dynamics
- ballistic vs. servo control
- timed control
- position/orientation/velocity/force control
- tool usage & action at a distance
- multiple temporal phases

- sensor modalities (e.g., visual, haptic)
- compliance
- constraints on the workspace/environment

Participants should attempt to characterize their tasks and designs according to these (and other) features to facilitate comparison.

SUBMISSION & PREPARATION: Potential participants should submit a short description of their background and research interests along with designs and analyses for individual tasks. To improve the depth and quality of the designs, participants are encouraged to work in teams, especially in collaborations that combine complementary expertise. Of course, demonstrations of working systems, including simulations and videos, are encouraged. Send submissions to either:

Steven Whitehead
GTE Laboratories Incorporated swhitehead@gte.com
40 Sylvan Rd. phone: (617) 466-2193
Waltham, MA 02254 FAX: (617) 890-9320

or

David Coombs
Natl Inst of Stds and Tech (NIST) coombs@cme.nist.gov
Robot Systems Division
Building 220, Room B-124 phone: (301) 975-2865
Gaithersburg, MD 20899 USA FAX: (301) 990-9688

ORGANIZING COMMITTEE:
Emilio Bizzi, MIT; Jon Connell, IBM Watson; David Coombs, NIST, co-chair, (coombs@cme.nist.gov); Ken Goldberg, USC; Rod Grupen, UMass; Stan Rosenschein, Teleos Research; Steven Whitehead, GTE Labs, co-chair, (swhitehead@gte.com);

IMPORTANT DATES:

Submissions due:	October 15, 1993
Notification of acceptance:	November 15, 1993
Final registration deadline:	March 1, 1994
Spring symposium:	March 21-23, 1994

From @CSLI.Stanford.EDU:bug-sdl-request@weber.ucsd.edu Thu Jul 29 05:20:44 1993
From: brd@cs.cornell.edu (Bruce Randall Donald)
Date: Wed, 28 Jul 93 15:12:14 -0400
To: bug-sdl@weber.ucsd.edu
Subject: Special Issue

I have just finished editing a special issue (actually, three issues in one volume) of Algorithmica on Computational Robotics. The issue is out now, and if you're interested you might want to read or order it.

{\em Algorithmica}, Vol.~10, Nos.~2/3/4, (ed. B. Donald). Special issue on {\em ``Computational Robotics: The Geometric Theory of Manipulation, Planning, and Control.''} (Springer-Verlag: New York) Aug/Sept/Oct (1993), pp.~91--352.

---Bruce