CSG_9111

Date: Thu Oct 31, 1991 1:28 pm PST
From: jbjg7967
Subject: Pribram's pontifications

[from Joel Judd[

Awhile back mention was made of K. Pribram's "new" book (_Brain and Perception_). A couple of people said they had seen it (Bill was one I think) and someone (Chuck?) quickly posted some tidbits from the intro. Since then there has been no comments about it.

I tried to get through as much as I could [understand], and would like to mention some of what seemed important claims, and get responses from others more knowledgeable than myself. In general, I was surprised how many of his references dated to the 50s and 60s. I even went to the index in Powers' BCP and sure enough, there are comments here and there to some of the same things Pribram is still discussing (eg. TOTE). In general, Bill was either accepting of Pribram's findings (eg. limbic involvement in perception of sequence) or sure that if a certain hypothesis panned out, it wouldn't controvert PCT. Finally, most of the book is devoted to visuospatial work, and since there are several imminently qualified people on the net, I'd be grateful for some translations of the highly technical discussions.

Pribram starts out with a premise of "reciprocal connectivity between hierarchically ordered neural systems, by means of which processing leads to a selection proceedure in which input is matched against a resident microstructure (genetically or experientially produced memory). The result of the match acts like a set point on a thermostat (or homeostat) to instruct further processing" (p.4). Unfortunately he goes on to discuss the set point as a "set of attractors," but the outline of a control process seems to be there. One of the results of "reciprocally acting" systems is that self-determination increases the larger the amount of experience "stored...in a top-down fashion."

Chapters 3-4 discuss the transformations that take place in the visual system, especially those occurring initially--from lens-retina and the inner-outer plexiform layers. Basic to Pribram's discussion is a resolution of the "grain problem" ("either a) percepts are constructed from simpler elements or b)regions in the brain directly resonate to configurations already present in the input to the senses. He says the solution is to call sensory processes "transformational": initial processing involves "reordering a distributed, enfolded array that forms the content of the sensory input" (p.48). Is he using "transformation" in the same sense as CT explains? The discussion which follows is where I begin to lose track of the argument. Pribram differentiates between "optical image" (the transform of the pupil/lens), and retinal receptor (more complex transform processes). Of interest in the course of these processes is to be able to determine what then are the "critical features" of a precept. His answer: "...that which remains invariant across processing stages." These are found by placing constraints on processing in a situation which can be described using physics or engineering terminology. "The limits within which the subjective report of the ensuing perceptual experience remains unchanged provide an indication of the invariant features of that situation" (p.48). Is this returning to the problem of finding objective environmental features, or something more subtle (or too confusing to tell)?
Chapter begins discussion of higher level processing, specifically object forms. This seems to be good evidence for the postulated change in order of levels 3-4 discussed in August. Bill, would you read this chapter and report back to the class tomorrow? The quote beginning the chapter is by Helmholtz (1909):

> Thus, by our movements we find it is the stationary form of the table in space which is the cause of the changing image in our eyes. We explain the table as having existence independent of our observation because at any moment we like, simply by assuming the proper position with respect to it, we can observe it.
> The chapter deals with terms like "constancy" and how movement is essential to

perceiving the "invariants" previously mentioned, at least that's my naive interpretation of it. If the physics of this could be explained in layman's terms, I'd be much obliged.

What's also intriguing is that he hints that phoneme construction can be interpreted in much the same way as object-form perception (citing an OLD--1969--Liberman et al paper and the Motor Theory of Speech), and also talking about a "center of gravity" determination of vowel construction, citing one Syrdal (1985). Bruce, Martin--any clues to the connection here?
Well, I'll quit with something from the summary to chapter 6 that sounds like a pseudo-control system:

The execution of actions is controlled by targets composed of images of achievement. Images of achievement are constructed by a two-part process: learning what to do and doing it skillfully. Skill is DEFINED by a continuous movement toward a target...Skill is ACHIEVED by an optimization proceedure (learning what to do) that eliminates "error," defined as movements that do not accord with the least action principle in perceptual processing. (160)

Date: Fri Nov 01, 1991 5:33 am PST
From: CHARLES W. TUCKER
Subject: ON SOCIAL CONTROL,,,AGAIN

CSG-EM17 FROM CHUCK TUCKER 911101.0700

RE: Changing conceptualizations of 'social control':Classical vs. Modern

I am gathering materials for my course next semester on "control and social control" and I have been bothered by the current conceptualizations of 'social control' that ignore individual as selves or self-regulating organisms. This has concerned me so that I intend to put together some of the materials on this NET and with the assistance of all of you who care to join present a paper at the SSS meetings entitled "The Myth of Social Control." The other day University of Chicago Press had a sale and I purchased a book that I had seen but was not in my library: Morris Janowitz. 1976. SOCIAL CONTROL OF THE WELFARE STATE and I found this statement:

"'Social control' refers to the ability of a social group or a society to engage in self-regulation (Park and Burgess, 1921; Turner, 1967 cited here). The obverse of social control is coercive control. . . . In the intellectual tradition of sociology, social control transformed sociological inquiry at the turn of the century from a speculative enterprise to an empirical research effort (his review of 1975 cited here). The notion of social control carried with it philosophical implications - concern with higher moral principles - the basis on which social control would be created and the ethical principles and goals to which social groups would aspire. Suddenly, in the 1930s, under the impact of the Great Depression, sociologists departed from these formulations of social control as they became fascinated and preoccupied with issues of power. The result was not an effective sociological perspective for analyzing the complexities of the social (p. 9) structure and its relation to political power. Instead there emerged an emphasis on oversimplified 'power' theories often grounded in gross economic determinist arguments. The idea of social control was temporarily transformed into a pejorative

term which came to mean conformity and social repression. But the contemporary problems of the welfare state focus attention on the underlying character of social order and the issues of social control. . . . Throughout this essay, the term 'social control' will be used in its 'classical' and enduring sense, although this is not the fashion in which it is used by many writers (p. 10)."

I also suspect the influence of another deterministic theory (I will let you chance a guess here) and the failure of the notion of "internalization" also had something to do with this transformation. I plan to trace this conceptual history out and will likely (surprise, surprise!) propose that there be a return to the "classical" concept of social control and use other words (e.g., power, force, coercion, and the like) with the

Date: Fri Nov 01, 1991 7:49 am PST
From: jbjg7967
Subject: development of control

[from Joel Judd]

Developmental psychologists or anyone who might know:

Forgot to add a related question yesterday. I've been reviewing my huge neuroanatomy textbook (Schwartz and Kandel, $50 in 1988) lately looking for evidence of control systems, since I took the class before my recent conceptual reorganizations occasioned by you all (and you know who you are). Anyway, I was wondering, given the recent discussion on afferent:efferent pathway ratios and descending control, what if any developmental evidence exists, say, for certain pathways to be predominantly afferent and then show development of efferent pathways over time, how much descending control might be inborn (eg. the visual orienting control from the superior colliculus) etc. In other words, what developmental evidence is currently available?

Date: Fri Nov 01, 1991 8:41 am PST
From: Martin Taylor
Subject: Re: Catching fly balls

[Martin Taylor 911101 10:45]
(Gary Cziko 911031.1110)

Gary asks:
> I can see how playing in a domed stadium or outdoors
>with a partly cloudy sky could provide a background field against which to
>judge movement of the ball. But this may not be possible outdoors under a
>perfectly clear or completely overcast sky. Is catching fly balls more
>difficult on a clear or overcast day outside?

As a long-time cricketer, I can answer definitively "Yes". A day with a clear blue sky almost brings a (cricketing) outfielder to a state of terror for just this reason. There is no background against which to judge the ball. I have nearly been killed under such conditions because of misjudging a ball hit very high. Only in the last few milliseconds did I realize that the ball was falling behind me instead of in my hands. It could have been on top of my head just as easily.

Martin Taylor

Date: Fri Nov 01, 1991 11:04 am PST
From: CHARLES W. TUCKER
FROM CHUCK TUCKER 911101.1337

. . . appropriate definitions for these other processes. COMMENTS ?

Thanks,

Gary

Date: Sat Nov 02, 1991 3:16 pm PST
From: TJ0WAH1
Subject: Neural signals

[from Wayne Hershberger]

Gary Cziko:

I consider myself an old hand, even an expert, at eye pressing, but I must admit that it never occurred to me to do it while walking about. Yet, it is such an obviously clever thing to try, I can't imagine why I never tried it. My compliments!

Rick Marken:

You said something some time ago about a symposium at WPA. I would be interested in participating if the time were right, but I can't remember when WPA meets. As for your more recent remark, "the purposeful behavior of organisms IS control," I would like to chorus, "Amen."

Martin Taylor:

>I wonder if you are aware of the 1974 thesis of my colleague
>R.G. (Bob) Angus: Attentional Focus and Saccadic Suppression
>(York University, Toronto). His experiments supported the
>interpretation that there was no (or little) actual attenuation
>of vision before or during the saccade, but a compensation for
>the predicted change of visual direction, beginning as much as
>200 msec before the actual movement began.

No, Martin, I'm not. Could you please fill me in. What were his methods? Matin, Matin, & Pola (1970) reported finding "monotonically increasing shifts...with stimuli presented as early as 240 ms before the saccade," but this gradual trend is partly an artifact of their method, Fechner's method of constant stimuli. This psychophysical method is insensitive to discrete shifts of retinal local signs, and will even misrepresent a discrete shift as a continuous one:

> For instance, suppose, for sake of argument, that the shift in retinal local signs that attends a saccade occurs in a discrete, stepwise fashion, and that the latency, but not size, of this step varies from trial to trial. Successive trials of repeated stimulation of the same retinal locus at the same relative time (e.g., 20 msec prior to eye movement) will yield a bimodal distribution of apparent visual direction, one mode comprising the effects of the trials on which the stimulus precedes the shift and the other comprising the effects of the trials on which the stimulus follows the shift. The central tendency of the distribution as a whole, customarily taken to represent the true visual direction or local sign of the retinal signal, may be observed to depend heavily upon the relative frequencies of the two types of trials, which, in turn, depend heavily upon when during the perisaccadic interval the stimulus is presented. In general, the later the stimulus occurs in the interval, the more frequent the postshift trials are likely to be and, hence, the greater the apparent shift in local signs, even though the actual shifts are all of the same magnitude whenever they occur (Hershberger, 1987, p. 40).

We know now (from Jordan's dissertation) that my example's use of a moment 20 msec before the saccade was too conservative; it should have been 80 ms, which is the time at about which the shift occurs, discretely and completely.

Bill Powers:

I have just downloaded the oculomotor paper. Thanks. I'll start chewing on it over the weekend but it may take me some time to digest it. I've got some things on my plate which have been demanding my immediate attention and it may take a while--please be patient with me.

Bill Powers (911026.1100)

>As I discovered with the eye model, and later with the Little
>Man dynamic model, strong feedback and tight control don't
>necessarily produce the effect you expected. The reason isn't
>that strong feedback doesn't produce tight control, but that the
>system isn't controlling what you thought it was. In the arm
>model, the tightest control isn't on position, but on
>acceleration. The next loop, the dynamic stretch loop, controls
>velocity tightly, but not position. The combination of stretch
>and force feedback results in very low gain for position
>control. You have to get all the way to the third level (or
>fourth?) before position control is much good, and even then you
>have to add the visual loop to make it really tight. Maybe this
>is where model-based control should be introduced.

I was trying to pick your brain about this in Durango this Summer, but there were too many demands on your time at that time. I would very much appreciate your expanding on this theme now that you are snowed in for the duration. I think that you are probably correct, this IS where model-based control should be introduced.

Bill Powers (911027.1200)

>neural signals are not ABOUT the world of experience; they ARE
>the world of experience.

Almost, but not quite. Any claim that neural signals are ABOUT the world of experience is undoubtedly false, as you say, precisely because the reverse is the case: my world of experience is ABOUT my neural signals--and about the optic array in the ambient light, and about the stuff with resting mass which is said to give structure to the ambient light. It is not true that my experiences comprise neural signals. Rather, to quote an authority you may recognize, "this world presents itself...in three dimensions, stereo sound, and living color, chock full from edge to edge of continuously-present smoothly changing noise-free colors, shades, objects, motions, relationships, and operations in progress."

According to a coherence theory of truth, the perceptual objects comprising our experience (phenomenal world) may be said to depend upon (be about) the conceptual objects we "construct" (neurons, photons, electrons, input functions) to the degree and only to the degree that these conceptions (models) account parsimoniously for the perceptions in question.

Bill Powers (911028.0900)

>A configuration-level input function does not RECEIVE
>configuration signals; it GENERATES signals that indicate
>presence of a specific configuration.

I think you've got it right here. It doesn't generate the configuration, it generates the signals (in that part of the putative physical domain called neurophysiology) that indicate the presence of a specific configuration (in that part of the putative physical domain called the physical environment).

Perception is NOT imagination.

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu

Date: Sun Nov 03, 1991 8:52 am PST
From: Oded Maler
Subject: Re: Multilevel control & dimensionality

[From Oded Maler 911103]

(Reply to Bill Powers (911028.0800))

This how I interpret your claim:

Given n elementary one-dimensional control-systems that can somehow interact with the world such that their preceptions of it will match their reference signals, then it is possible to construct from them several m-dimensional control systems each of them controlling for some linear combination of a subset of the basic input. In fact, the controlled virtual variables of the higher level systems need not be a linear function but monotone in each input variable. Later you restricted these "input functions" to be linearly-independant (otherwise the pair of functions x and -x will be an immediate counter-example). Can you try to characterize more precisely what type of systems of input functions can operate concurrently on such an n-dimensional space? Are you sure that just using the signs is sufficient?
--Oded

Date: Mon Nov 04, 1991 5:45 am PST
From: POWERS DENISON C
Subject: Misc comments

[From Bill Powers (911104.0700)]

Gary Cziko (911031) --

>How do I know if the angle of the ball is growing at a steady rate?
>You've provided an "objective" variable to control, but PCT has to do
>better than than this. What is the "subjective" variable that is controlled.

Martin Taylor et al have provided some practical answers. The subjective variable has to be the apparent rate of rise of the ball relative to a coordinate system anchored to the stadium or the horizon. You need a patterned backdrop to judge the angular velocity of the ball if your peripheral vision isn't wide enough to include the horizon and the ball at the same time (for all but the highest fly balls the normal visual acceptance angle should suffice). The approximation doesn't have to be especially precise; if it's slightly off, you will be running at slightly the wrong rate, but the angular speed of the ball will then depart from constancy enough to notice and you will adjust your running rate. As the

ball approaches, angular velocity errors imply smaller and smaller position errors and become easier and easier to correct (although it's true that you have less and less time left).

Epistemology buffs note: the subjective perception of angular velocity has to approximate the "actual" angular velocity (measured against some scale) if this method of catching is to work. This is one of those indirect pieces of evidence that makes me suspect (though I can't prove) that there's a Boss Reality out there.

Joel Judd (911031) --

Re Pribram:

>Pribram starts out with a premise of "reciprocal connectivity between
>hierarchically ordered neural systems, by means of which processing
>leads to a selection procedure in which input is matched against a
>resident microstructure (genetically or experientially produced memory).
>The result of the match acts like a set point on a thermostat (or
>homeostat) to instruct further processing" (p.4). Unfortunately he goes
>on to discuss the set point as a "set of attractors," but the outline of
>a control process seems to be there. One of the results of "reciprocally
>acting" systems is that self-determination increases the larger the
>amount of experience "stored...in a top-down fashion."

While Pribram often comes up with interesting ideas, they don't seem to have any explanatory power. For example, in the bit above he says that "the result of the match acts like a set point on a thermostat..." but he doesn't say how the RESULT of the match, rather than one of the inputs to the process, can act like a set point. Explaining this would be critical if the model is to mean anything.

I always get the feeling that Pribram is fuzzy about just how things like that work, as if he has some smart electronic technician in the background whose explanations of mechanism he is relaying without actually understanding them. How do you match "input" against a "microstructure?" One is a signal and the other is a thing or an abstraction. What is "reciprocal connectivity?" There are so many kinds, with such radically different implied properties, that this generalization actually could be true of wildly different kinds of systems. It says nothing. What is this scale of self-determination? What kind of fashion is a "top-down fashion?"

Pribram's explanations seem to me more metaphorical and anatomical than functional. His papers are like beautiful colored brochures that show you the package of the product and tell you how much it will do for you if you buy it, without ever presenting the product itself. Even his concept of "slow potentials," which is a germ of a good idea, goes nowhere because he can't say what underlying processes are going on that would be affected in some specific way by sub-threshold interactions among neurons. I think that Pribram has never managed to cross the line between abstract generalization and modeling. He's always telling you the class to which something belongs without telling you what difference that membership makes. So the brain operates like a hologram. Fine. But what does the brain's hologram DO? What is the brain FOR?

>Basic to Pribram's discussion is a resolution of the "grain problem"
>("either a) percepts are constructed from simpler elements or b)regions
>in the brain directly resonate to configurations already present in the
>input to the senses. He says the solution is to call sensory processes
>"transformational": initial processing involves "reordering a
>distributed, enfolded array that forms the content of the sensory input" (p.48).

You see? Either perceptions are veridical or they are constructions (which just about covers the possibilities). The solution is to classify them as "transformational," as if finding a taxonomical slot for them tells you which choice to make. If the transformations

are one-to-one, then the perceptions are veridical. If the transformations generate outputs that are arbitrary functions of the inputs, then they construct rather than report a reality. So the solution isn't a solution at all --the basic question is untouched. Clever.

You quite rightly call the following a "pseudo-control system":

"Skill is DEFINED by a continuous movement toward a target...Skill is ACHIEVED by an optimization proceedure (learning what to do) that eliminates "error," defined as movements that do not accord with the least action principle in perceptual processing." (160)
This smacks of Rosenbrock's model of "policy" as a control process. It's still computation of output, with goals in the environment. And of course the "least action principle in perceptual processing" is a complete bluff. What is the least-action minimum of a function that recognizes a triangle? Have you calculated any perceptual Hamiltonians lately, Dr. Pribram?

No matter which of a dozen mutually-contradictory models of brain function might ultimately prove to be correct, Pribram will be able to point to his writings and say "See? I got there first." I sometimes wonder if this isn't the explanation of most writings of this type -- an attempt to get lucky with a generalization without any idea of why it is or isn't appropriate. By the way, if you haven't done so yet, count the number of self-citations in Pribram's book. This will tell you whose voice is loudest in his ear.

Tom Bourbon (911030) --

Swamped with students? Sounds interesting. Why not put some of the progress reports you sent me on the net? I'm sure that everyone will be fascinated.

Chuck Tucker (911101) --

We will certainly pitch in via the net and do what we can to help on the paper for the SSS meetings (Southern Sociological Society, I presume?). Put me on the list to receive drafts.

Janowitz' adoption of the statement that "'Social control' refers to the ability of a social group or a society to engage in self-regulation" doesn't quite seem to capture the CT view (at least mine). Just making it the obverse of coercive control doesn't make it right. But perhaps you're just showing that one sociologist has recognized coercive control as a major feature of the standard view of social interactions.

Joel Judd (911101) --

>... what if any developmental evidence exists, say, for certain pathways
>to be predominantly afferent and then show development of efferent
>pathways over time, how much descending control might be inborn...?

Good question: do we begin mostly with "instinctive" perceptions or "instinctive" actions? Another question that (understandably) hasn't been raised in developmental psychology or ethology is whether we begin with instinctive reference signals, too. Of course an instinctive reference signal wouldn't mean much without a perception to go with it ...

Wayne Hershberger (911102) --

>I would very much appreciate your expanding on this theme [model-based
>control] now that you are snowed in for the duration.

I won't have much to say until I can incorporate this idea into a working model. Other work to do first. But it's on the list.

>It is not true that my experiences comprise neural signals. Rather, to
>quote an authority you may recognize, "this world presents itself...in
>three dimensions, stereo sound, and living color, chock full from edge
>to edge of continuously-present smoothly changing noise-free colors,
>shades, objects, motions, relationships, and operations in progress."

I was not describing what the world IS, I was describing how it APPEARS. This appearance is the world we directly apprehend. The CT model, based on the appearances of neurology and physics, leads me to conclude that this world is, physically, a collection of neural signals, although it does not look that way (that is, it does not look the way neural signals do when we visualize them on an oscillogram using electronic means of sampling limited aspects of neural activity).

I know that you insist that there is no other world than the world of appearances (or that we don't need to consider one in explaining perception). When we aren't concerned with explaining, but only with experiencing and living (no theoreticians required), I agree with you. Reality is precisely what we experience. But to add "and nothing more" is to assert what we can't know directly, and to deny, for no good reason, the implications of physics, chemistry, and neurology, all of which claim to represent a world of immense detail that is inaccessible to our senses. You have never answered my comments about physics and so on.

From me:

>>A configuration-level input function does not RECEIVE
>>configuration signals; it GENERATES signals that indicate
>>presence of a specific configuration.

From you:

>I think you've got it right here. It doesn't generate the
>configuration, it generates the signals (in that part of the
>putative physical domain called neurophysiology) that indicate
>the presence of a specific configuration (in that part of the
>putative physical domain called the physical environment).

I don't know how configuration-perception works, but I think I do have a workable model of how sensation-perception works, and configurations are derived from sets of sensations (in experience as well as in the model). So if sensations are arbitrary patterns, then so are configurations.

Given any two intensity signals x and y, you can construct a sensation signal from the expression ax + by (a more complex expression could be used, but this is sufficient). The coefficients a and b are variable over a continuum of values. Hence the referent of a signal s computed as s = ax + by will change as a and b are varied (with or without any change in x and y). Each different way of weighting the basic intensity signals will create an apparent sensation that relates differently to other perceptions. If your view, translated down a level, were correct, then there would be only one combination of a and b that would provide a veridical sensation, while all other combinations would create illusions of nonexistent sensations. In order for the correct pair of a and b values to be discovered, there would have to be a way to compare the signal s with its referent, the "true" perception of the sensation, without going through the perceptual transformation s = ax + by.

But in the brain, the signal s is all there is: there is no independent channel through which the brain can know what sensation "actually exists." The only way I can see to arrive at a model consistent with experience, while maintaining that a brain is essential for experience, is to propose that the signal s IS the experienced sensation; that it is s we experience, and not its cause, a sensation comprised by x and y. As Oded Maler puts it, using an expression I have also used in the past with the same intended meaning, s

represents a "virtual" variable; it is an ordering imposed by the brain on the two variables x and y. There is no necessary correspondence between s and the underlying reality, although it is possible that the function ax + by expresses a significant relationship between x and y on the one hand, and other unseen variables on the other hand which form parts of our model of the putative physical reality.

The crucial difference between our views, as nearly as I can comprehend it, is in my assumption of another reality that is not part of the world of direct experience. If there is no such separate reality, then of course all that exists in nature is experience as it appears to us. Neural signals, physics, chemistry, and all such conceptions are just that: conceptions and nothing more. They are simply ways of ordering our experiences and have no significance beyond that. They do not refer to anything unexperienced.

I can understand that conclusion as a conclusion, but I can't accept it as a fact. It is simply another conjecture. Its truth, then, comes down to the evidence we have for and against it, and to how we reason about that evidence.

I think that control phenomena provide us with evidence that there is a universe beyond the limits of human perception; that this universe imposes its properties between our actions and their perceived results; that there are independent agencies in this universe that are capable of disturbing our control actions without our being able to detect the causes of the disturbances. I think we learn about these properties only indirectly, and as conditioned by the kinds of perceptual systems we have and do not have. I think we have to infer the nature of the disturbing agencies and the properties of the world, by building models that would, if they were true representations of the unseen world, explain how our experiences are related to each other. I do not think it is likely that we have arrived at models that just happen to capture every significant entity outside us, every significant functional relationship among those entities. And I do not think it is likely that the world of direct experience exhausts the degrees of freedom that really exist in the universe around us.

So that is my basis for accepting, as the most reasonable hypothesis, the existence of a real universe apart from our perception of it, and for denying, on the basis of the same reasoning, that our perceptions are likely to be veridical renditions of that universe.

>Perception is NOT imagination.

I agree. Imagination is, however, a subset of perception. Some of our experiences are generated inside the brain and do not depend on the current external state of affairs, even though they may sometimes give a convincing imitation. But the rest do depend on something outside. In neither case, however, do perceptions without the aid of reason give us a picture of what is really causing them -- however inadequate the picture.

Oded Maler (911103) --

I realize now that my use of a two-level system introduced an irrelevancy -- the argument doesn't depend on having two levels. You wonder about nonlinear functions; I used linear ones because I can't solve the equations for nonlinear cases.

>Can you try to characterize more precisely what type of systems of input
>functions can operate concurrently on such an n-dimensional space? Are
>you sure that just using the signs is sufficient?

I can try, but I don't know how impressed by my attempt a real mathematician will be. The easiest way for me to generalize this approach is to use simulations instead of analytic algebra. I'll lay this out mostly in words, and perhaps you can help me find a more rigorous way of getting to the point.

Suppose we have a function of several environmental variables, the function producing one perceptual signal of varying magnitude. This perceptual signal is compared with a

reference signal, generating an error signal. The first question to consider is whether an output function can be constructed such that the perceptual signal is guaranteed to approach the value of the reference signal.

I think that a more general argument than what follows can be constructed, but I'll not attempt it: here I will assume that the response of the environmental variables to outputs of the control system is proportional or at least globally monotonic, and does not involve any time functions (the environment is critically damped).

Because the perceptual signal depends on a set of input variables, altering any given input variable can have only one of three effects on the perceptual signal: to increase it, to decrease it, or to leave it unaffected. We can therefore connect the output of this control system, through separate paths, to each of the environmental variables and adjust the sign of the effect to assure that the perceptual signal will be changed, via each path, in the direction toward the value of the reference signal. We can omit connections that have no effect either way. This means that the error signal will have to have an effect, after transmission through each complete loop, of contributing to a decrease in the error signal.

There is already a negative sign in the loop, introduced by the standard design in which the perceptual signal is subtracted from the reference signal. The remaining sign of effects around the loop must be positive. So we calculate the partial derivative of the perceptual signal with respect to each variable, and use its sign to determine the required sign of the effect of the one output signal on each environmental variable.

To bring the error to zero, we now add a time-integrator to the output function, so that as long as the error is non-zero, the output will continue to increase in the corresponding direction. This is what I meant by using a "simulation." It's really a method of steep descent.
The system will then approach the zero-error condition with time, provided that the overall relationship of output to perceptual signal remains monotonic or at worst has only local reversals. Within this constraint, the relationship can be nonlinear. The output will simply continue to change until the error is zero.

It's important to take the control system's point of view here. The objective of control is not to achieve any particular combination of values of the environmental variables; it's to make the one perceptual signal match the reference signal. It doesn't matter what the individual values of the external variables are, as long as the perceived "virtual variable" reaches its reference state.

Now it should be clear why no weightings of the output signals are needed. Equal weightings mean that the output affects all the environmental variables by equal amounts, although not in the same directions. If we used weightings, the variables would be affected by different amounts. In either case, there would be some overall gain factor by which the perceptual signal depends on the output signal. The fact that the output is affecting the perception by many paths instead of just one makes no difference in the outcome except in how fast the final state will be approached -- some paths may contribute only weakly to error correction. But the output will continue to change until there is no error to cause further changes -- it doesn't matter what the external overall amplification factor is.

Now we can introduce a second system sensing the same collection of environmental variables. A different input function is needed -- the second system must not sense the same function of the variables or a function differing only by a constant of proportionality, or there can be no solution.

The second system is designed exactly as the first one is, except for the form of the input function. And the result will be that the second perceptual signal will match the second reference signal independently of what the first system is doing.

I suppose that the constraints on the nonlinearities of the two input functions would now be more complicated. We have to consider whether there are values of the two outputs that can cause the two perceptual signals to reach all values in the required ranges (eventually) by continuous or piecewise paths. I don't know how to specify the constraints. If the two input functions are orthogonal there will be no problem, but pairs of systems like these can work even with a considerable degree of non-orthogonality. I do know of one limit: the disturbance of each system's perception caused by the other's output has to be less than some critical amount to avoid a positive-feedback interaction. I think the limit is an effect that is 50 percent of the system's own effect on its own input (with two systems). More generally, positive feedback paths created by interactions must be weaker than the local negative feedbacks.

You can see that the questions become complex when multiple systems are modeled. It would be nice if someone could come up with a general statement of the requirements for multiple systems to work this way, without output weightings. WITH output weightings, borderline cases could be made to work better. But my point is that the weightings are not likely to be critical, and, I think, can usually be dispensed with. This is certainly true of all the linear systems I've worked with, but that won't satisfy a mathematician.

The most important aspect of this analysis is the idea that these systems are controlling for a subjective condition, not an objective one. They are only trying to control their own perceptual signals. The state of the environment that actually exists when all errors are corrected, in terms of the detailed environmental variables, is immaterial. That objective state matters only insofar as it affects the organism -- and if it does matter, the organism will acquire another control system to keep that aspect of the situation, as perceived, under control. Yet another function of the same set of environmental variables. So in the end the objective state of the external world never does matter to the system. It matters only to the observer of the system.

Date: Mon Nov 04, 1991 6:40 am PST
From: Bruce E. Nevin
Subject: language, information, code

[From: Bruce Nevin (911104 0752)]

I have been out sick since Wednesday afternoon. And we just had "the blizzard of '78 minus the snow." An 11/6 deadline is at risk just now, and I have a lot to catch up on, while still unfortunately being ill.

(Martin Taylor 911030 14:10)--

Martin, I have read your response to the code/conduit business quickly, will try to get back to you first. Yes, I am operating in a different linguistic environment, one in which it is normal to make overt and to challenge the prevalent and customary (and almost always tacit) metaphors of language functioning as code, of language as having a code-like structure, and of language functioning as a conduit for conveying meaning from one consciousness to another.

I think the short of it is that natural language is different from formal "languages" in important respects, but much of what has been presented as linguistic theory has to do with formal "languages."

I put the latter term in quotes because of the metaphorical use of the term for formal systems such as programming "languages" and other offspring of mathematical logic. In particular, in the semantics of formal languages you do have a prior metalanguage external to the formal "language" for expressing its semantics and its syntax. This prior metalanguage depends upon the background vernacular of shared natural language and ultimately devolves to it for its information content.

It is in this context that the coding metaphor and the conduit metaphor have become so deeply established as to pass unnoticed. Informal extension of the meaning of "code," of the sort that you are surprised is not universal, seems to me to follow from an intuitive recognition that the relation between, say, Morse code and English is inadequate for metaphoric analogy to the relation between English and the information or meanings we express with and understand in English. But there is no explicit definition of this extension to the meaning of the term "code" and its congeners encode, etc., and that is I think because it is not even recognized to be metaphorical, it is presumed a fact about the relation of language to meaning. We do not explicate the shift of meaning of "king" from "absolute monarch" to "titular (figure)head of state" either. (Such semantic shift is characteristic of natural language, though not of formal languages.)

The only thing we have for representing the information in natural language is natural language itself. Even when we set up some system of "formal semantics" (read: semantics of a formal "language" into which sentences of a natural language are somehow translated), it turns out to depend upon information expressed in our shared background vernacular--a natural language.

Harris has determined a "least grammar" of a natural language (and more fundamentally a methodology for doing so), in which the combinability of primitive elements is as unrestricted as possible, and the description itself imparts no additional restrictions. Every difference of form (in these terms) corresponds directly to a difference in information. By formal manipulation, the obscuring effect of grammatical reductions can be removed from a text, and the text transformed so that this correspondence is direct and obvious. This transform of a text is a representation of the information in it.

Other theories of language presuppose a prior and external metalanguage. The grammatical machinery employed in other theories of language imposes restrictions that belong to the grammatical description and not to language, and obscure the correlation of form with information. Indeed, other theories of language do not seek such a correlation, because they instead seek a correlation of natural-language expressions with expressions in a formal "language" serving as semantic metalanguage. In some versions (varieties of semantic features) these expressions are alleged to "be" the meanings in some sense; in others (familiar also in AI) the mapping to expressions in a formal "language" is because of the transparent relation of the latter to its (typically truth-functional) semantic interpretation.

I hope this clarifies my perspective. I will try to attend more carefully to your message and others later this week.

Bruce
bn@bbn.com

Date: Mon Nov 04, 1991 9:52 am PST
From: Oded Maler
Subject: Re: Misc comments

(From Oded Maler 911104)

(Bill Powers 911104.0700, last paragraph):

>The most important aspect of this analysis is the idea that these systems
>are controlling for a subjective condition, not an objective one. They
>are only trying to control their own perceptual signals.

This aspect is the only one I understood already from the beginning. I find it reasonable, although the problem of *correspondence* between some perceived combinations of sensory signals and what happens to the organism in the "real" world, is sometimes under-emphasized. Even if the organism behaves according to simple perecption-based rules, in

order to show that this behavior achieves something in some external physical environment, you can get involved in very complex models which are very hard (if not impossible) to analyze.
I'll try to respond to the rest later.

--Oded

Date: Mon Nov 04, 1991 3:40 pm PST
From: marken
Subject: APS Update

[From Rick Marken (911104)]

Three people (besides myself) have expressed an interest in participating in a symposium on perceptual control theory at the American Psychological Society Annual Meeting in San Diego, Ca. But I have no committments yet. I have to submit the proposal for the symposium (or debate or whatever) by December 6. I have to submit names with the proposal. So, this is the LAST call. I must know by the end of this week (Nov 8) who is willing to have me put their name down as a participant in the PCT symposium. If the proposal is accepted, those whose names are listed would have to arrange to be in San Diego from June 20 - 22. As I said in the original post, I could provide lodging -- but no transpostation to San Diego and no food or convention registration costs.

I would like Tom Bourbon to be in on this, if possible. I already have "maybe"'s from Joel Judd, Wayne Hershberger and Bill Powers. Anyone who knows PCT and considers themselves a "scientific psychologist" (that's what APS is about) is also welcome to volunteer.
If I get a sufficient number of people committed to this symposium, then I will write up the proposal post haste. I should say that I would like the topic of the symposium to be something that goes right at the jugular of the "scientific psychology" establishment. I want to deal with questions like "what is behavior?", "what is the implicit model of human nature assumed by current psych methodology", "what is control?", "what is missing in non-PCT applications of control theory to behavior?", "why attractor models don't control", and "what does PCT imply about how to do psychology scientifically". It's alot of stuff - but it's all about foundational issues. The title might be something like "What if it really doesn't work that way? The nightmare scenario for the life sciences" or "A scientific revolution in psychology: They don't call it 'revolution' for nothing". Well, actually, I plan to come up with a much nicer title than these -- but I do want the symposium to get right down to the nitty -gritty. I also want it to be fun. So, let me know ASAP if you can commit to going to the meeting if, by some strange chance, the proposal is actually accepted.

Best regards

Rick

Date: Tue Nov 05, 1991 3:05 am PST
From: Oded Maler
Subject: Configuration control

(From Oded Maler 911105)

This as an attempt to rephrase the lower-levels of the model, and to make the underlying assumptions and abstractions more explicit.

The starting point is the existence of n elementary control systems, each of them can make its input match its reference by outputting a well-defined error signal which is the difference between input and reference.

The following points should be noted:

1) There is no notion of time in this model (this can be rephrased by having an infinite gain). The transduction of the environment variable into an input signal, the computation of the error signal, its transduction back to the environment, and its influence on it are all assumed to take no time. (Note that it is not necessarily a negative criticism against the model, such an abstraction is underlying many "respectable" mathematical disciplines).

2) The *range* of values of input-reference pairs which the system can handle is not mentioned explicitly. However, it is clear that in realistic systems this range is bounded. This might be meaningful when we move up to multi- dimensional systems, because some trajectories for correcting a "virtual" signal may go outside the controllability region of the elementary inputs. Also, concerning the previous note, if "real" time is taken into account, it is clear that it might take longer to correct a larger error signal.

3) The inter-dependence among the elementary systems, their inputs and outputs is ignored. It is assumed that each of them can achieve its goal concurrently and independently of the other (this also relates to the previous point concerning ranges).

4) It is assumed that "perception is not imagination", that is, that there is some more-or-less real world at least at the level of basic sensations. The ends of the sensory-effectory organs of the basic control systems are causally connected to something out there (out of the boundaries of the control systems, it can be inside the body of course) and this something influences the input signals and is influenced by the ouptput signals of each elementary control system. This last assumption is not necessary for the model itself but it is necessary for its analysis in relatively-objective terms. I don't want to get into this aspect, but just want to make sure that what we mean generally by "controlling the perception" is less trivial than, say, closing out eyes, or hallucinating (at least at the lower-levels).

Now given such a set of elementary systems, you want to control some more complicated perecptual signals which are computed from the elementary ones. As long as you don't take propagation delays, computation time, wiring limitations etc. into account, that is, the "input functions" are as timeless and "non-constructive" as the elementary systems and signals, it makes no difference if you consider two-levels or a more refined hierarchical decomposition. Because if two higher-level control systems are in conflict with each other, this is inherent in the way their respective input functions are influenced by the elementary variables. It is not that I am against hierarchical decomposition, on the contrary, but with the current underlying model, (still at what you call the "configuration" level), everything that can be controlled hierachically can be controlled with two-levels (and this is most evident in the case of linear combinations).

So, assuming the existence of ideal, timeless, and unbounded-range elementary control systems, I believe it can be proved that such and such systems each controlling some function of the inputs can operate according to the principles you mentioned, given that their inter-dependence is such and such (I believe that for the linear case, some "real mathematicians" can do it rather easily).

I have to stop now.

--Oded

Date: Tue Nov 05, 1991 11:36 am PST

From: marken
Subject: configuration control

[Fom Rick Marken (911105)]

Here are some quick comments on Oded Maler's (911105) attempt to make the underlying assumptions and abstractions of the hierarchical PCT model more explicit.

>1) There is no notion of time in this model (this can be rephrased by having
> an infinite gain).

Not really. Effects propogate around the loop at some rate. This is handled in my spreadsheet simulation by adding only a portion of the error determined output to the output variable on each iteration. We usually do ignore transport lags -- but they are important (in real control systems) and can be incorporated into the simulations very easily.

>2) The *range* of values of input-reference pairs which the system can handle
> is not mentioned explicitly. However, it is clear that in realistic systems
> this range is bounded.

Correct.

> This might be meaningful when we move up to multi-
> dimensional systems, because some trajectories for correcting a "virtual"
> signal may go outside the controllability region of the elementary inputs.

This is also a problem for simple control systems. I you cannot produce, say, outputs that can counter the prevailing disturbance, then you cannot control the input variable.

> Also, concerning the previous note, if "real" time is
> taken into account, it is clear that it might take longer to correct a larger
> error signal.

There are relative timing considerations in the design of hierarchical control systems. Powers has already mentioned at least one of them, viz. higher level systems must operate more slowly (correct error more slowly) than the lower level systems on which they depend. This also has implications for the relative gain possible for systems at different levels: higher level systems must have lower gain than lower level systems. Maybe this is why we can control our posture better than our principles.

>3) The inter-dependence among the elementary systems, their inputs and
> outputs is ignored. It is assumed that each of them can achieve its goal
> concurrently and independently of the other (this also relates to the
> previous point concerning ranges).

This is not really true of PCT. In a hierarchical control system of the type Powers proposes there are well defined inter-dependencies between systems in terms of perceptions, outputs and even mutual effects -- the environmental effects of one control system can influence the input to another control system in the hierarchy -- and vice versa. These inter-dependencies are all nicely handled by the individual control systems. How-ever, there is a dependency (which Bill mentioned) between the nature of the perceptual connections and output connections of hierarchically related control systems -- these connections must preserve negative feeback.

>4) It is assumed that "perception is not imagination", that is, that there
> is some more-or-less real world

Totally real -- its nature is just "more-or-less" inferrable from experimental tests of models of what that reality is like (ie, from scientific behavior).

> The ends of the sensory-effectory organs of the basic control systems are
> causally connected to something out there (out of the boundaries of the
> control systems, it can be inside the body of course) and this something
> influences the input signals and is influenced by the ouptput signals of
> each elementary control system.

Exactly.

> what we mean generally by "controlling the perception" is less trivial
> than, say, closing out eyes, or hallucinating (at least at the lower-levels).

You bet. It's that pain-in-the-ass real world out there, between my outputs and my experience, that makes it so much more interesting to actually control my perception of myself doing a perfect forehand smash than to imagine myself doing it.

>As long as you don't take propagation delays, computation time, wiring
>limitations etc. into account, that is, the "input functions" are as timeless
>and "non-constructive" as the elementary systems and signals,
>it makes no difference if you consider two-levels or a more refined
>hierarchical decomposition.

Do you mean that, if these things are taken into account, then hierarchical decomposition is not going to work? But we do take these into account -- and the model works. Again, my spreadsheet works with three levels -- intensities, sensations and relationships. And it includes all but the "propagation delay" considerations that you mention above. I bet it would work fine with them in too.

> Because if two higher-level control systems are in conflict
>with each other, this is inherent in the way their respective input functions
>are influenced by the elementary variables. It is not that I am against
>hierarchical decomposition, on the contrary, but with the current underlying
>model, (still at what you call the "configuration" level), everything that can
>be controlled hierachically can be controlled with two-levels (and this is
>most evident in the case of linear combinations).

I guess I don't understand the above paragraph. Are you assuming that hierarchical control systems will necessarily have conflicting elementary systems if they are not "instantaneous" or linearly decomposable?

>So, assuming the existence of ideal, timeless, and unbounded-range elementary
>control systems, I believe it can be proved that such and such systems
>each controlling some function of the inputs can operate according to
>the principles you mentioned, given that their inter-dependence is such and
>such (I believe that for the linear case, some "real mathematicians" can
>do it rather easily).

Maybe that is what you mean? My guess is that this is not true. I also believe that a mathematical proof would be difficult or limited (like the Minsky/Papert proofs that perceptrons could not detect connectivity). I think the best approach is to do what the current perceptron people (now called neural networkers) are doing -- just forge ahead building models of the capabilities that you know exist in the nervous system. If the models work, then great; if not, try something else. With control theory, it is the organizational principle that is most important -- not necessarily the details of the model. Though, to the extent that the details work, they serve as an existence theorm for plausible control structures.

Hasta Luego

Rick

****

Richard S. Marken USMail: 10459 Holman Ave The Aerospace Corporation
Los Angeles, CA 90024 Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

Date: Tue Nov 05, 1991 12:04 pm PST
From: Martin Taylor
Subject: Re: Configuration control

[Martin Taylor 911105]

A (possibly) simple technical question.

For linear transformations such as rotations and translations in a multidimensional space, it is the case that any sequence of transformations can be replaced by a single linear transformation that takes the initial configuration into the same final configuration as the sequence. This will be true whether the sequence of transformations involved manipulations of lower-dimensional independent subspaces or manipulations of the total space.

In perceptron-like neural network structures, the power of multiple levels comes from the fact that the individual layers perform non-linear operations (they saturate). I hypothesize that the power of the hierarchic control system model comes also from the non-linear effects of the elemental control systems (apart from considerations of simplicity of design and of explanation).

Is is true that the effect of any hierarchy of linear control systems can be emulated by some one-level multidimensional linear control system that incoporates all the degrees of freedom embodied in the elements of the hierarchy?

Martin Taylor

Date: Tue Nov 05, 1991 12:47 pm PST
From: marken
Subject: correction

[From Rick Marken (911105b)]

I said (911105):

> higher level
>systems must have lower gain than lower level systems. Maybe this is why we
>can control our posture better than our principles.

Sorry, it's the other way around -- higher level systems (which are slower) can have HIGHER gain than lower level (faster) systems. If there is a difference in our ability to control lower (posture) vs higher (principles) level variables, the required difference in gain with level would predict poorer control for the lower level variables. This is too bad since my guess is that we do control lower level variables better than higher order ones. It is the reason , I think, why consciousness seems to spend most of its time monitoring the higher levels of control. Ah well. Maybe this is worth a little research.

Regards

Rick

****

Richard S. Marken USMail: 10459 Holman Ave The Aerospace Corporation
Los Angeles, CA 90024 Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

Date: Tue Nov 05, 1991 2:42 pm PST
From: Martin Taylor
Subject: Re: correction

[Martin Taylor 911105 16:20]
(Rick Marken 911105b)
>
> my
>guess is that we do control lower level variables better than higher order
>ones. It is the reason , I think, why consciousness seems to spend most of its
>time monitoring the higher levels of control. Ah well. Maybe this is worth
>a little research.
>

I agree with Rick's intuition, but propose a different solution--that the reorganization needed to develop an accurately functioning high-level control system takes a long time, and may never actually be completed.

Feedback at these higher levels can involve much delay, so that it becomes a very difficult long-range problem to discover which direction to alter the reference signals for lower-level control systems. In addition, the environment may be changing (think of Seymour Papert trying to ride a bicycle after the inverting prisms were removed from his spectacles), so that what was a helpful action last year is now an unhelpful one.

At any level, there are many control systems working in parallel, all supported by many of the same elemental control systems at the level below. It will not normally be the case that all the upper ones require the same direction of change in the percepts of the lower ones. An upper one cannot tell whether a lower one has actually moved in the direction the upper one required, because it is only an influence on the lower one, along with all its peers that also rely on the same lower one. Only statistically can the upper one determine whether the changes it asks of the lower one were in the direction it really needed, and the gathering of these statistics might take a very long time.

[PLEASE, don't get sidetracked by the obvious factthat the statistics are not computed numerically. That red herring gets dragged across too many tracks.]

I suspect that many of the references driven by the requirements of high-level systems are inappropriate to the satisfaction of those requirements, and that the environment is not stable enough to allow a precise determination of when this is happening. Add to that the likelihood of conflict, and you get the intuited situation, in which the lower level control systems work more precisely than the higher ones.
Martin Taylor

Date: Tue Nov 05, 1991 2:45 pm PST

From: Martin Taylor
Subject: Re: Very old stuff

[Martin Taylor 911105 1700]
(Bill Powers 910724; yes, July)

I've been trying to clear up a mess of old files, and came across this. Way back in the Light Ages around the summer solstice, Bill said:

>I'm still worried, however, about configurations of configurations,
>sequences of sequences, and so on. The only real answer concerning these
>possibilities is to get down and look. If you can genuinely experience a
>configuration of configurations (not sneaking in any relationships, and
>not perceiving the different "recursions" in parallel), then so be it. So
>far, my conclusion is still *no recursions.*

We had exactly the same problem for a little while in developing the Layered Protocol theory of communication. What we realized was that we were making a type-token error. If a particular protocol appeared to be supporting itself, it wasn't, and couldn't be, because it could be in only one state at a given moment. What was actually happening was that a separate instance of a protocol type was supporting another instance of the same type (e.g. for a silly example, one might type words by spelling them "double-u, oh, are, dee, ess", in which case a word-level protocol would be supporting another word-level protocol).

Could something of the same kind be happening with configurations of configurations? There need be no recursion, inasmuch as different elemental control systems control the configuration and the configuration of configurations.

But then again, maybe after 4 months, this is no longer a conceptual problem.

Martin Taylor

Date: Tue Nov 05, 1991 3:35 pm PST
From: Martin Taylor
Subject: Re: language, information, code

[Martin Taylor 911105 17:45]
(Bruce Nevin 911104 0752)

I await your further comment, but I disagree with the thrust of what you have so far written, in these statements:

>
>The only thing we have for representing the information in natural
>language is natural language itself.
>
>Other theories of language presuppose a prior and external metalanguage.

I do not see why a theory of language needs any more than a description of the circumstances in which the language is used and the circumstances in which it is used. There seems to me to be zero probability that a formal description can encompass the uses or the observed structure of language. I know there exist many theories of language that treat "language" as if it were an object having a finite description. That is not a view with which I can sympathize. To me, language is a means of affecting the knowledge or behaviour of another party. To some extent the user of the language takes advantage of conventionalized forms, but to a greater extent the communicative process is one of

situation-specific negotiation. I cannot believe that natural language (and still less that formal language) can describe all these possible negotiations, and that is what would be required "to represent the information in natural language".

If all you are talking about as natural language is the relationship of word types in structured phrases or sentences, then you are talking about a greatly impoverished skeletal version of the fully fleshed substance of language. I can imagine that natural language would suffice to describe that skeleton, but I withold judgment as to whether it does so.
A theory of language must, as a minimum

>"explicate the shift of
>meaning of "king" from "absolute monarch" to "titular (figure)head of
>state" either. (Such semantic shift is characteristic of natural
>language, though not of formal languages.)"

Or rather, it must indicate how a listener knows what the speaker refers to in saying "the king of rock singers" or "the king of Ruthenia". (Or why the listener does not know, in some circumstances). Language use is full of partial references, structural associations (e.g. using a register normally inappropriate to the situation, for special effect--as Bill's "I be with you there, Dude" (partial misquote, I expect) in response to a posting (of yours?) a few months ago on the formal properties of language), unrecognized metaphors (well, that's what most of language is, isn't it), false identifications (the use of the word "is" and its relatives usually signals this)...

How can we represent language? I don't know how we represent it to talk to each other about it, better than with natural language, but internally we use it rather than represent it. We don't use a simulator of language, any more than we compute numerically the statistics we use. We use language itself--a process, not an object.

Martin Taylor

Date: Wed Nov 06, 1991 7:55 am PST
From: POWERS DENISON C
Subject: Model; d of f; metalanguage

[From Bill Powers (911106.0800)]

Oded Maler (911105) --

I think you've captured the model as described in the algebra accurately. I'm delighted that you think there is some reason to go along with my generalizations at least for the linear case. Thanks for the intellectual effort. It's hard to read someone else's equations.

In the "real" model, some of the limitations of which you speak are removed -- although I'm unable to write the equations for all 11 levels!

In particular, don't confuse the algebraic simplifications for a working dynamical model of the system. As I think I mentioned, the algebra is valid under the assumption that the differential equations that describe the complete system have steady-state solutions (i.e., the system is dynamically stable). The result is a quasi-static picture of the behavior, under the assumption that nothing changes fast enough to invalidate the algebraic relationships. For me, at least, trying to handle the dynamics gets in the way

of seeing the basic relationships --I get lost in the complexities of the dynamics and lose the most revealing relationships. I also run out of mathematical petrol.
If you include dynamic properties of the variables x and y (such as mass and viscosity), of course, you then have to introduce the dynamics of control in order to show how the system is stabilized. If you just assume that the system is stabilized, these considerations don't have to be introduced. That's the basis for my algebraic treatment.

So time is not really missing from the model -- it's just suppressed. If the inputs change over time, the signals inside the system also change over time. The simultaneous algebraic solutions show the end-points of those changes. The perceptual signals will remain close to the reference signals if the changes aren't too rapid. That is almost a definition of the range of conditions over which the system can maintain good control. I didn't include disturbances acting directly on the controlled variables (I can do that if you want, or perhaps you could try doing it yourself to get familiar with the manipulations). Disturbances can be functions of time, as can reference signals, without altering the algebra of the system's internal operation. This takes us one step closer to a model with full dynamics.

As you say, the algebra gives no indication of the range of values of the variables. There is, of course, a limited range in the real system. Neural signals have a maximum impulse rate, and furthermore they can't go negative. To make a control system that can control a variable THROUGH zero, it's necessary to have at least pairs of systems operating in balance, with comparators for positive errors being physically separate from comparators for negative errors. Such pairs can be reduced to equivalent single systems with bidirectional signals. The nonlinearity of muscle response (and sensory nerve response) creates some interesting opportunities for gain variation through setting the opposing reference signals so as to create a small amount of conflict at "zero" net reference signal. I haven't tried to construct a working model with this amount of detail, but it's certainly possible. Also, in real systems there is a limit to the size of external disturbances that can be counteracted by changes in output -- real muscles have limited strength.

The interdependence of the separate systems is only partially ignored. If the perceptual functions are not completely orthogonal, the control processes in one system will disturb the inputs of the other system, and of course there is strong interaction at the (unweighted) outputs. The other system will simply alter its output to cancel such disturbances, just as if they came from the environment. In the equations I posted, the input functions are not orthogonal and the unweighted outputs certainly have effects that are not optimally aimed. Yet the algebra shows that independent control is achieved. If you look at the equations prior to the approximation achieved by letting loop gain become very large, you will see that both virtual variables depend on both reference signals. But with high loop gain, the effect of the unwanted reference signal is made arbitrarily small.

In another sense, however, I do ignore interactions -- possible influences of one perceptual function directly on the other. In principle that kind of interaction can be handled by redefining the perceptual functions so that they depend slightly differently on the inputs from lower levels. They already receive their inputs from a common set of lower-level signals, so an effect that goes from the lower level, into one perceptual function, into the other perceptual function (the interaction) and thence to the other perceptual signal is equivalent to a direct effect through a modified perceptual function in the other system.

Yes, this model accepts the reality of the physical model of the environment. It thus employs two kinds of models. One model is the ordinary physical model of an environment, the other is the proposed model of the neuromuscular system. It's models all the way, defined in terms of human perceptions (mine).

The higher-level systems are not limited to the constraints on the lower-level systems. I did mention the possibility of uncontrolled signals reaching higher level input functions.

But it is also possible for new KINDS of variables to become controlled. For example, a higher-level system could institute a temporal pattern of reference-signal changes in the existing lower-level systems, and perceive in terms of a dynamic function of the lower-level variables, such as rate of change or (for oscillatory variables) phase or amplitude. Even with the elementary system depicted by the equations I posted, it would be possible for a higher-level system, through adjustment of reference signals, to make x-y increase at one speed while x+y decreases at a different speed (for a while). Or the system could make a sine-wave change in the difference between x and y match a sine-wave change in the sum. Or the sum could be made to vary at twice the frequency of the difference, or to vary in a sine-wave 90 degrees out of phase with the difference, or to have twice the amplitude of the difference changes without regard to frequency. This would amount to control of new stationary functions of the two "virtual" variables. Also, it would be possible to create an infinite number of different patterns of changes -- events -- in which the elements are successive reference-values of x-y and x+y in arbitrary combinations. At even a higher level, a nonzero value of x-y could be treated as a dot and a nonzero value of x+y could be treated as a dash, so messages could be sent in Morse or binary code. There is no need to limit the higher-level input functions to the same types encountered at the lowest levels, weighted sums. In fact, that is the basis of my hierarchical model: a change of type of variable with each successive level.

Conflict is not limited, therefore, to the terms of the lowest level. Two higher-level systems might control different temporal patterns of the virtual perceptions. These patterns could be compatible or incompatible, quite independently of whether there is conflict at the level of weighted sums. The content of two multiplexed Morse-coded messages might prove to be contradictory. And so on.

Rick Marken (911104) --

Count me in on the symposium. Mary and I can stay with her sister and my brother-in-law, who happen to live in the same place in San Diego. How much will the registration cost?

Martin Taylor (911105) --

>Is is true that the effect of any hierarchy of linear control systems >can be emulated by some one-level multidimensional linear control system >that incoporates all the degrees of freedom embodied in the elements of >the hierarchy?

This must be true, with respect to spatial dimensions at least. If we consider more generalized dimensions, however, I'm not so sure. In the physics game, there's a tendency to take the higher levels of perception for granted -- to treat them as the backdrop against which the lower levels of behavior are studied. Physics doesn't do much with "shape," for example. Only a few analytical forms are ever considered, together with abstract measures such as moment of inertia. The shape of am amoeba is almost impossible to characterize by any analytical method -- yet anyone can recognize a photomicrograph of an amoeba. So even at the configuration level of perception there is more going on than would be evident in the spatial degrees of freedom of the environment. I'm not sure that shape recognition can be accomplished, fully anyway, using only linear combinations of weighted summations. Depth perception alone almost certainly has to involve nonlinear functions.

As I indicated to Oded above, after all the weighted summations are done, there is still the dimension of time to consider. This dimension is an independent variable at the intensity and sensation levels, with respect both to disturbances and reference signals. To make temporal functions into controlled variables, another level is needed that perceives temporal functions, as the lower levels do not. I don't include here the internal time-dependencies, but only what is represented by a steady perceptual signal. When we perceive velocities, a steady perception means a constant velocity -- constant change of the lower-level variables. At the configuration level, change doesn't exist in perception; the control systems are always looking at the present-time error and trying to make it small, even if the reference signal is changing. That change itself is not

represented in perception: a static perception means a static arrangement, which is what a configuration is.

I am, by the way, getting very tempted to rescind the switch of the transition level to somewhere below configurations. The differentiation and integration I spoke of may be just the internal mechanism of the sensation-level input functions. When we see a continuous static configuration, the microchanges that maintain it are not evident in subjective perception, even though they exist. I don't think they are under control. I think I want to return the transition level to its former position, just above configurations.

That aside, there are many more "dimensions" to consider once we free ourselves of the idea that perceptions are all linear combinations of lower levels. As I said above, an "event" introduces a new type of variable, not strictly a time variable nor a space variable. Relationships take us outside the realm of analytic functions: "between", for example, or "inside," or "less pungent" or "noisier." And categorizing brings us to the brink of logical variables, with the dimension of sequence or ordering next, and rules following that. I don't see how such things could be expressed as one big linear combination of intensities. It's as though the lower levels are blind to aspects of the world that can be captured only by higher perceptual functions that pay attention to temporal variables, temporal patterns, and relationships, classifications, and ordering of such things. And more.

Bruce Nevin (various) --

I'm having second thoughts about "language is perception", too. It seems to me, for example, that Harris's partial orderings can't be recognized without using a mode of perception that isn't expressed in words. I'm having doubts that language has any metalanguage at all: it's just language -- as Martin says, a process instead of a thing, or as I would say, signals handled by a process that is not represented in the signals. That which does language is not the same thing as the language that it does. Even to recognize the difference between two different word orderings requires sequence perception, which can't be expressed in words but can only be done. In other words, I think that a "metalanguage" is being asked to do a job that only a neural perceptual function can do.
Best to all

Bill P.

Date: Wed Nov 06, 1991 8:59 am PST
From: cam
Subject: Re: Configuration control

[Chris Malcolm 6/11/91]

[Martin Taylor 911105] aks the interesting question:

>Is is true that the effect of any hierarchy of linear control systems >can be emulated by some one-level multidimensional linear control system >that incoporates all the degrees of freedom embodied in the elements of >the hierarchy?

And if so, is it necessarily the case that the hierarchical system is the more economical implementation?

Date: Wed Nov 06, 1991 9:41 am PST
From: Joseph Michael Lubin
Subject: saturation versus thresholding

[From Joe Lubin 911106.1100]

[To Martin Taylor 911105]

> In perceptron-like neural network structures, the power of multiple levels
> comes from the fact that the individual layers perform non-linear
> operations (they saturate).

It is more accurate to say that their computational power comes from a thresholding operation by which subspaces are partitioned. While it is possible for neurons to use their saturation properties to nonlinearize their subspaces in useful ways, there must always be a decision made which requires a (soft or hard) threshold. Threshold-linear neurons (without saturation but with a threshold) have decent computational power.

Date: Wed Nov 06, 1991 9:54 am PST
From: Oded Maler
Subject: Re: configuration control

[From Oded Maler 911106]

(A reply to Rick Marken (911105))

> I guess I don't understand the above paragraph.

Let me try to re-explain. Remember that I'm still at the configuration level (and I won't climb up until I'm sure I understand it..) and I consider 2 alternative decompositions of this level (I should call it super-level because I consider its decomposition into sub-levels). In one case there are just two levels, sensations and configurations which are, say, linear combinations of sensations. In the alternative case there are some intermediate levels between the two, with the usual exchange of reference signals, input etc. My trivial claim is that if a pair of configuration control systems cannot reach their goals simultaneously (for some pair of their reference signals) in the first case, the same pair of configurations cannot be controlled for in the second case. (I'm not sure whether the converse is true - someone less ignorant than me in linear algebra could check, but it shows that unless you put some additional details into the model (constraints on the connectivity, computation time and space, etc.) everything that can be controlled by a hierarchy can be controlled by a 2-level hierarchy. [And I'm *not* talking yet about temporal patterns, events which are qualitatively different perceptions]

--Oded

>Maybe this is why we
>can control our posture better than our principles.

I loved this sentence, although it might not be true for all of us..

Date: Wed Nov 06, 1991 10:10 am PST
From: Oded Maler
Subject: Collapsing Hierarchies [Was: Re: Configuration control]
[Oded Maler, Ibid.]

[Chris Malcolm 6/11/91]

>[Martin Taylor 911105] aks the interesting question:

>>Is is true that the effect of any hierarchy of linear control systems
>>can be emulated by some one-level multidimensional linear control system
>>that incoporates all the degrees of freedom embodied in the elements of
>>the hierarchy?

>And if so, is it necessarily the case that the hierarchical system is
>the more economical implementation?

I believe (depending on the asumption) that the answer (to Martin's question) is yes. You can express the highest-levels as linear combinations of the lowest-ones, and get a two-level (i.e., one-level) system that has at least the same range of control.

As for your question, the answer depends on what you call "economy". Comutation times, fan-in? wiring, propagation delays? evolutional plausibility?

This reminds me of a paper by Y. Abu-Mustafa from Caltech, were he claimed that although every Boolean function can be computed on a basis of gates with fan-in 2 (that is, low-connectivity, "high"-hierarchicality), in order to learn (generalize from examples) there is an advantage in high connectivity as exhibited by neurons. His arguments had to do with the entropy of the environment, but, as usual, I didn't follow the details.
--Oded

Date: Wed Nov 06, 1991 11:05 am PST
From: Martin Taylor
ubject: Re: Dimensionality

[Martin Taylor 911106 13:00]
(Bill Powers 911106.0800

>Martin Taylor (911105) --
>
>>Is is true that the effect of any hierarchy of linear control systems
>>can be emulated by some one-level multidimensional linear control system
>>that incoporates all the degrees of freedom embodied in the elements of
>>the hierarchy?
>
>This must be true, with respect to spatial dimensions at least. If we
>consider more generalized dimensions, however, I'm not so sure.
>...
>
>As I indicated to Oded above, after all the weighted summations are done,
>there is still the dimension of time to consider. This dimension is an
>independent variable at the intensity and sensation levels, with respect
>both to disturbances and reference signals.
>
>
>That aside, there are many more "dimensions" to consider once we free
>ourselves of the idea that perceptions are all linear combinations of
>lower levels. As I said above, an "event" introduces a new type of
>variable, not strictly a time variable nor a space variable.
>Relationships take us outside the realm of analytic functions: "between",
>for example, or "inside," or "less pungent" or "noisier." And
>categorizing brings us to the brink of logical variables, with the
>dimension of sequence or ordering next, and rules following that. I don't
>see how such things could be expressed as one big linear combination of
>intensities.

I asked the question because most of the technical discussion to date has been based on linear systems. It is not clear to me whether all linear control systems, specifically INCLUDING those whose degrees of freedom are distributed over time, can be emulated by a single-level multidimensional linear system. Remember that "linear" can include time delays. Basically linear means that the concept of superposition applies--effects can be summed and differenced--and this seems to be fundamental to many of the specific comments that have been made: reference signals for a low-level control system are a weighted sum of the signals asserted from the next higher level, for instance.

Representation of the control system behaviour in Laplace Transform notation should take care of the time dimension, and such transforms are legitimate for linear systems, so it should not matter whether the degrees of freedom are distributed over time, space, or both.
In neural network studies, it is shown that any LINEAR network (e.g. one in which the nodes form a weighted sum of their inputs) can be emulated by a single-layer network. The power of multilayer neural nets depends on the nonlinearities that ensure that the effects of the successive layers cannot be seen as simple rotations of a basis space. My question is whether the power of a hierarchic control system likewise depends on the nonlinearities that are usually dismissed in the ongoing discussions in this group.

In the last quoted paragraph, Bill seems to be very clear that the nonlinearities are indeed important. This is also my intuition. As soon as you partition the perceptual space into categories, as you must whenever you choose between behaviour possibilities such as "fight or flight", you inevitably lose linearity. At the cusp, some reference level splits into two--successful fight or successful flight. Either may be chosen to satisfy the higher reference (perhaps happy survival or some such). Language multiplies this need to categorize. Planning becomes an issue, which is not the case in a basically linear system. It is never assured that behaviour that initially moves the perception nearer a reference level is behaviour that can move the perception TO the reference level.

So, if non-linearities are acknowledged to be an important aspect of the hierarchic control system, my question becomes moot. Nevertheless, it retains some interest. Does anyone have a definitive answer?

Issues related to planning and feedback should be a major thread in future discussion, if they are tractable.

Martin Taylor

Date: Wed Nov 06, 1991 11:50 am PST
From: Martin Taylor
Subject: Re: saturation versus thresholding

[Martin Taylor 911106 13:30]
(Joe Lubin 911106 11:00)
>
>
>> In perceptron-like neural network structures, the power of multiple levels
>> comes from the fact that the individual layers perform non-linear
>> operations (they saturate).

>It is more accurate to say that their computational power comes
>from a thresholding operation by which subspaces are partitioned.
>While it is possible for neurons to use their saturation
>properties to nonlinearize their subspaces in useful ways, there
>must always be a decision made which requires a (soft or hard)

>threshold. Threshold-linear neurons (without saturation but with
>a threshold) have decent computational power.

That's a very restrictive notion of the functions of neural nets. The need for a decision is imposed from outside, in the experimenter's choice to use the net as a classification device. Nns are quite good at that, but I think that to treat them as primarily classification devices is to ignore most of what they do best. Typically, there is no need for a decision until there is a need for behaviour. The nets usually do not control behavioural choices. What nets do well is to bind the perceptual space into topological neighbourhoods, in contrast to rule-based systems which treat the perceptual space as being built up from a myriad of unrelated entities.

Actually, I don't care which kind of non-linearity you introduce. But you must have some to do more than some kind of generalized rotation of the perceptual space, such as a Fourier transform.

Martin Taylor

Date: Wed Nov 06, 1991 12:36 pm PST
From: Bruce E. Nevin
Subject: fast MRI and imagination

[From: Bruce Nevin (911106 1432)]

I read an account in Monday's Boston Globe (pp. 28-29, in the Health/Science section) of speeded-up MRI technology (magnetic resonance imaging) developed at Mass. General Hospital here in Boston.

They get an image every 1/20th sec, as opposed to 1/5min. This is not fast for neural events, of course, but it registers increase in blood volume in areas where neural activity is increased. The work reported Friday in _Science_ involved seven healthy people perceiving visual stimuli (checkerboard of lights). John Belliveau is the principal author.

"When the lights were on and flashing . . ., in the far back of the brain . . ., blood volume increased by mroe than 30 percent, and the rest of the brain remained unchanged." (Dr. Bruce Rosen, MGH radiology, a co-author)

A particular point of interest:

> Earlier studies using another type of scanning technology called positron emission tomography (or PET), which produces fuzzier images, suggest that "when you think about a dog versus seeing it, that same part of the brain seems to light up," but additional areas are called into action as well, Rosen said.

That last clause seems to me to have important implications for the "imagination connection." PCT as presently formulated would anticipate that the same areas of the brain (actually less, and not more) would be activated, and that other areas of the brain would not be activated. Perhaps this has to do with something extra that is required to get into imagination mode as opposed to handing up sensory input from (in this case) the retina.

Bruce

PS--Life has not got less busy. I put Sarah in the hospital yesterday with pneumonia. She's much better today, but the night before was rough. My Wednesday deadline has slipped to Friday. Whoopdedoo. Back later.

Date: Wed Nov 06, 1991 1:07 pm PST
From: Gary A. Cziko
Subject: Recycling Perception

[from Gary Cziko 911106]

I just received a copy of "Campus Recycling News" where I learned that my campus's recycling efforts last fiscal year saved 20,000 trees.

One of the headlines says "Close the Loop--Use Recycled Products."

Since recycling is now "in" and politically correct, perhaps PCT could profit from this "close the loop" mentality. How about "Behavior: The Recycling of Perception"? Sounds downright ecologically friendly, doesn't it?
--Gary

Date: Wed Nov 06, 1991 2:06 pm PST
From: marken
Subject: Re: configuration control

[From Rick Marken (911106)]

Oded Maler (911106) says:

>Let me try to re-explain. Remember that I'm still at the configuration level
>In one case there
>are just two levels, sensations and configurations which are, say, linear
>combinations of sensations. In the alternative case there are some
>intermediate levels between the two, with the usual exchange of reference
>signals, input etc. My trivial claim is that if a pair of configuration
> control systems cannot reach their goals simultaneously (for some pair of
>their reference signals) in the first case, the same pair of configurations
>cannot be controlled for in the second case.

I think this is probably correct. I cannot think of how to prove it right off hand -- but it has been my experience when modelling equivalent two level and single level systems.

>but it shows that unless you put some additional details into the
>model (constraints on the connectivity, computation time and space, etc.)
>everything that can be controlled by a hierarchy can be controlled by a
>2-level hierarchy. [And I'm *not* talking yet about temporal patterns,
>events which are qualitatively different perceptions]

I agree and would go even futher. It is likely that everything (I would say any particular variable-- not necessarily "everything") can be controlled by a 1 level hierarchy (even temporal patterns). PCT is not committed to a hierarchical structure, let alone the particular hierarchical structure proposed by Powers. The hierarchical structure is a hypothesis about how control systems might be structured. It is based on some reasonable evidence including 1) the anatomical structure of the nervous system 2) phenomenal observation of experience (perceptual experience seems hierarchical) 3) some behavioral experiments -- Powers' arm movement demo, Powers and my mirror reversal study and 4) modelling efforts -- Powers' little man and eye control models.This is all very circumstantial and not too compelling. What is needed is some combined research/modelling efforts which can test whether ot not a hierarchically structured model is needed.

I should make a point here that Bill Powers has made many times but fits nicely in this context. Bill has never claimed that his hierarchy should be taken as revealed truth. The hierarchy is mainly a pedagogical device; it helps people (like me) understand that a control system organization can IN PRINCIPLE do everything that we consider "behavior". The hierarchy shows that a control organization can walk, talk, fall in love, get married and go to church. All of these "behaviors" can be understood as controlled perceptions (of sensations, configurations, transitions, relationships, programs, principles and system concepts). But this is all hypothesis. What we need are smart people like you (Oded) who can actually build a system that might be able to perceive and control a "relationship" or a "program". These are very non-linear, non-analytic percetual com-putations. If you can build a machine that imitates a system controlling some complicated variables (like relationships or principles) and if you can do it with a non-hierarchical architecture, then that's evidence that maybe hierarchy is the wrong hypothesis. All PCT has to say is that, when modelling control, try to understand the perceptions being controlled. Of course, to model this you also have to understand how system outputs can influence these perceptions (via the environment). But the main problem of modelling complex control, from a PCT point of view, will be modelling the processes that perceive the controlled result. This will be tough -- but not hopeless. What is hopeless (from the PCT point of view) is trying to model complex strategies for producing outputs to cause the controlled result.

Best Regards

Rick

****

Richard S. Marken USMail: 10459 Holman Ave The Aerospace Corporation
Los Angeles, CA 90024 Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

Date: Wed Nov 06, 1991 2:07 pm PST
From: Martin Taylor
Subject: Re: fast MRI and imagination

[Martin Taylor 911106 16:40]
(Bruce Nevin 911106 1432)
>
>A particular point of interest:
>
> Earlier studies using another type of scanning technology called
> positron emission tomography (or PET), which produces fuzzier
> images, suggest that "when you think about a dog versus seeing
> it, that same part of the brain seems to light up," but
> additional areas are called into action as well, Rosen said.
>
>That last clause seems to me to have important implications for the
>"imagination connection." PCT as presently formulated would anticipate
>that the same areas of the brain (actually less, and not more) would be
>activated, and that other areas of the brain would not be activated.

Check Lassen, Ingvar and Skinh|j, Brain function and blood flow, Scientific American 1978, 239, 62-71 (Don't know which month). They have fascinating pictures of the different areas of the brain involved with various activities such as reading or mental arithmetic. There are later papers as well, but I don't have references.

Martin Taylor
PS I intend the symbol "|" to represent the Danish slashed "o". This probably is not the conventional symbol, but that's what it means in "Skinh|j" above.

Date: Wed Nov 06, 1991 3:42 pm PST
From: Gary A. Cziko
Subject: Bickhard and encodingism

[from Gary Cziko 911106.1650]

Bill Powers:

I got a reply from Mark Bickhard which he said I could share with you. In fact, he said:

>A conversation with Powers would be very interesting to
>me - please feel free to forward that little comment of
>mine. I have some reservation about getting involved in
>a time demanding network, but might consider that. I am,
>however, leaving town tomorrow and returning Nov 18, so
>I could not be much of a conversational participant until
>after that - and maybe that is the time to find out more
>about your network as well.

Here is first reply.

>Thank you for your note. Yes, I have read Dr. Powers
>book and other works, and, I think, have even met him
>once at a Gordon conference. I agree that the
>interactive model has strong convergences with the
>"behavior as control of perception" model, but I would
>egocentrically state the difference somewhat
>oppositely from Dr. Powers. In particular, I think that
>construing behavior as subserving the control of
>perception is not far off in a strictly functional sense,
>but it is simply wrong and highly misleading from an
>epistemological sense. If a system is supposed to be
>controling perception, how is it to know what those
>perceptions are? Without further modeling, the default
>answer is encodingism, and it is clear by now that I at
>least find that unacceptable. If, on the other hand, it is
>recognized that the only way to recognize perception -
>or anything else - is indirectly (implicit definition or
>differentiation) via reaching particular strictly *internal*
>states, then we have behavior subserving control of
>internal states, not of perception. The inputs, in this
>view, serve only, and logically can serve only, the
>function of participating in various internal state
>changes. They are not epistemically directly
>accessable. In this view, it is not so much that behavior
>is another form of perception, as Dr. Powers suggests,
>but, instead, that perception is another form of
>interactive behavior (Bickhard and Richie, 1983).
>
>If it is argued that perceptual inputs can be recognized
>by some sort of match process, then I ask "match of
>what, and how is it done?" If its supposed to be an
>epistemic match, then how does the system know what
>it's receiving, and how does it know what it is even
>trying to match - in other words, how do you solve the
>problem of representational content? If no epistemic
>claim is made, but instead only some physical level,
>neural presumably, process subtracts firing

>frequencies over determinate axons (or some such test
>for match), then 1) what the system is controling is not
>the inputs, but the internal states that result from such
>purely functional (nonepistemic) "matching processes".
>Once it is understood that interaction must be in the
>service of such control of internal state (since internal
>state is the only functionally accessable aspect of a
>system), then it is also clear that such simple matches
>are not the only manner in which useful controls of
>internal states could occur. In fact, abstract machine
>theory has an entire mathematics devoted to exploring
>what sorts of input patterns can be functionally
>"recognized" (though again only in an indirect
>functional sense). None of these require or involve
>control of perception in any direct sense, including in
>their interactive versions.
>
>Mark

Perhaps if you and other CSGnetters could generate some interesting reactions, we could entice Bickhard's to join the net. I realize that probably only Bill Powers has seen his work, but others may still be able to react to Bickhard's comment here. I will forward comments to him from the net.

Let me close by saying that I find that Bickhard has made some powerful arguments against the type of "symbolic encodingism" on which much of cognitive science seems now to be based and provides insight into why the type of hierarchy we use in PCT should in fact develop. I think we would have much to learn from his interaction with us, as well as he from us.

--Gary

Date: Wed Nov 06, 1991 9:41 pm PST
From: RLPSYU08
Subject: Re: fast MRI and inagination

From Tom Bourbon --
In nearly every case, reports of brain areas that "light up" in particular tasks, or under particular "challenges," present difference images -- images in which only the areas that exceed some threshold amoount of difference between control and challenge conditions are shown. All "areas" in which the difference does not exceed the threshold are blacked out in the image. Whatever else they might show, difference images do not show "the parts pf the brain that perform the task" under consideration.

Right Writer (Rite Riter?) and PCT:

An undergraduate student, eager to show me how his new software checked for grammar, usage and other things, showed me a passage in which the program had marked his text. The program inserts remarks, marked by <<*. When I read the material, I saw something he never suspected was there.

"Bourbon and Powers (1991) demonstrate <<*REPLACE demonstrate by SIMPLER show or prove?*>> the results when two popular theories of behavior (stimulus-control theory and cognitive plan theory) are modeled <<* PASSIVE VOICE: are modeled *>> and simulated: both models fail in the simulation. <<* IS THIS SENTENCE TOO DIFFICULT? *>> The perceptual control theory model <<*
IS THIS AMBIGUOUS? *>> ...!
Give up folks -- even a dumb software package puts
us down!

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402

Date: Thu Nov 07, 1991 7:27 am PST
From: POWERS DENISON C
Subject: Config control; modeling

[From Bill Powers (911107.0630)]

Chris Malcolm, Joe Lubin, Martin Taylor, Oded Maler (911106) --
Martin asked:
>Is is true that the effect of any hierarchy of linear control systems
>can be emulated by some one-level multidimensional linear control system
>that incoporates all the degrees of freedom embodied in the elements of
>the hierarchy?

And Chris commented:
>And if so, is it necessarily the case that the hierarchical system is
>the more economical implementation?

The question and its answer are not relevant to a model of a real behaving system. It is possible to trace out at least some hierarchical loops in real nervous systems: therefore the correct model will be some sort of hierarchy. It doesn't matter whether, after the fact, the same functions could have been accomplished with a one-level system. They are not accomplished that way. If a system is organized as a hierarchy, it will have different failure modes from one that is not; it will have different time-delays, and different partitionings of time-delays. Its evolutionary development will be different. Its stages of growth will be different from that of a one-level system. I don't think that our models should be driven by those mathematical concepts to which the brain has become accustomed (although I obviously don't disapprove of mathematical models). As Rick Marken said with a wisdom that becomes more apparent as a function of time: phenomena first; theory second.

Joe Lubin says
>While it is possible for neurons to use their saturation
>properties to nonlinearize their subspaces in useful ways, there
>must always be a decision made which requires a (soft or hard) threshold.

I'm going to need some help with what you mean by nonlinearizing subspaces in a useful way, Joe.

Is it "decisions" that we need to model? Such an assumption would have to be shown pertinent to understanding what the brain does before I could accept it. I think that this concept threatens to take us back to a digital model of the nervous system, in which the supposed all-or-none action of neurons led naturally to the idea that behavior consists of either-or choices between alternate actions.

People clearly do make choices, but that is not all that they do. In fact, in my model choice-making would occur in one sense at the category level, and in another sense at the program level. That's 2 out of 11 classes of control processes. If I am correct that there are classes of control processes, and that each higher level introduces a new class, then for each class we need a new computing principle -- one model is not going to explain the whole brain. The type of computing process needed at the level of spinal control of the stretch/tension "reflex" does not need "decisions". Brain-stem processes can probably get by with linear or nonlinear weighted summations at the level I term "sensation" control, with no choices necessary. At the configuration level we need a process that can extract invariants relative to many different transformations: 3-space rotations, translations,

expansions, and figure-ground contrasts. These processes must produce a perceptual world in which configurations are continuously present and continuously transformable, just like the world that our perceptions actually present to us. Configuration control requires no decisions.

As far as I can see, the first level where something like a decision is needed is the level at which we draw boundaries around groups of lower-level percepts and create classes. But even here, decisions have to be made only near the boundaries: for the interior, all that is needed is an OR operation, which in a continuous world can be approximated by addition and normalizing (I think). Class-membership is experienced on a continuous scale, not either-or: given perceptual situations can be better or worse exemplars of a class. The membership establishes the identity of the variable, but the goodness of fit establishes its magnitude.

Decision as a choice between equally-valid alternatives, it seems to me, becomes relevant only when there is conflict -- when choosing either alternative has disadvantages. This implies that one higher-level control process can't correct its own error without inducing error into another one of comparable loop gain. If the alternatives aren't equally valid, or equally disadvantageous, then no choice needs to be made. If a choice does have to be made, the systems in conflict can't make it: there must be a higher-level arbiter that uses its own unitarity criteria of selection as the basis for setting the lower-level goals -- without conflict. If the higher-level system is missing, the conflict simply persists until the external situation changes. The nearest that a "leaderless" set of systems can come to automatic conflict resolution is through lateral inhibition, which creates hysteresis.

I would like to see some discussion of why "decisions" are important to model, and assuming they are considered important, what phenomena they are needed to explain.

Martin, you say

>Basically linear means that the concept of superposition applies--
>effects can be summed and differenced--and this seems to be fundamental
>to many of the specific comments that have been made: reference signals
>for a low-level control system are a weighted sum of the signals
>asserted from the next higher level, for instance.

I agree with your definition of "linear," but your example is unfortunate. If you simulate the situation I described to Oded in algebra, you will find that the superposition of reference signals from the two higher-level systems does NOT have to be linear for the system to work. Certainly some kinds of nonlinearity would be ruled out, but curvatures in the addition-subtraction process are OK if they're not too extreme. All that the curvatures can do is alter the effective weighting and the loop gain as a function of signal amplitude. It is easy to show that over some range, the output weightings can change without altering the final result, perceptual signal = reference signal.

Also, even in perceptual functions, linearity is not essential (it isn't even likely, given Weber-Fechner's or Stevens' "law"). This, too, can be demonstrated in simulation (the analytical approach rapidly becomes impossible). If reference signals are recordings of past perceptual signals, then the store of reference signals will contain exactly the same nonlinearities (relative to physical measures of the outside world) that the perceptual signals have. Linearity loses its meaning in connection with reference signals, because it has to be judged relative to the addressing scheme, or at least to the organization of output functions at the next higher level.

Martin, you're saying that Joe's concept is a "very restrictive notion of the functions of neural nets." So is yours, if you mean to rule out Joe's concept. While agreeing with me about decisions in general, you still say "What nets do well is to bind the perceptual space into topological neighbourhoods, in contrast to rule-based systems which treat the perceptual space as being built up from a myriad of unrelated entities." How does this

explain the fact that we do employ rule-based systems (language, mathematics, etc.) and that we do categorize and frequently have to resolve conflicts (i.e., make decisions)? I point to the early part of this post: one computational model is not going to cover all the types of perceptual variables we can control at all levels of organization.

You're going to have to give me some help with binding perceptual space into topological neighborhoods. Is that a classification of something that could possibly be illustrated more intuitively?

Mark Bickhard (indirectly, 911106 via Gary Cziko) --

I have defined perception as the existence of a signal in a perceptual neural pathway, not in terms of an objective external world. I consider the experienced external world to be a model in a brain, as opposed to the "real" world which is a hypothetical construct. So we are still in complete agreement. The brain has no way to verify that its perceptions actually represent something of the same form in an external reality, although it is permitted to brains to conjecture that there is some sort of correspondence.

Tom Bourbon (911106) --

Very cogent comment: the thresholding in the data conceals activities in other parts of the brain. "Whatever else they might show, difference images do not show 'the parts of the brain that perform the task' under consideration." Ah, me, will anyone listen? Have you noticed how simplistic cause-effect descriptions are in brain research (of course you have)? In Science this week there's more on those studies: a "photic input" creates "activity" in the visual cortex. All intervening stages of computation are gaily skipped over, as if they simply relay the effects without modifying them. Even since Huebel and Weisel, people have talked about "cells" that are sensitive to things like orientation of lines. How much lower-level processing is required before there can be a signal related in that way to retinal patterns? One thing is damned sure: the cell where the signal is measured isn't doing that computing.
Time to send -- great stuff on the net these days.

Today is Mary's birthday.

Best to all

Bill P.

Date: Thu, 7 Nov 1991 09:51:06 PST
From: marken@AERO.ORG
Subject: Bickhard and encodingism

[From Rick Marken (911107)]

Bickhard via Gary Cziko (911106.1650)

>the interactive model has strong convergences with the
>"behavior as control of perception" model, but I would
>egocentrically state the difference somewhat
>oppositely from Dr. Powers. In particular, I think that
>construing behavior as subserving the control of
>perception is not far off in a strictly functional sense,
>but it is simply wrong and highly misleading from an
>epistemological sense.

What is the "interactive" model??

NB. PCT does not "construe behavior as subserving the control of perception" -- behavior IS controlled perception. Even when behavior is part of the means used to control

perception (as when muscles are tensed and relaxed to make a forehand smash) the muscle tensions are themselves controlled perceptions. The only behavior that can be construed as subserving the control of perception are the uncontrolled outputs of the control system -- which, in nervous systems are efferent neural signals (the "error" signals of the control systems).

Related to this: Here is a "Great quote" I overheard Bill Powers make recently:

Someone said to Bill something like "So and so says that 60% of understanding behavior is a matter of understanding perception".

Bill said "I think so-and-so is about 60% right".

I loved it!

> If a system is supposed to be
>controling perception, how is it to know what those
>perceptions are?

What in the world does this question mean??!!?? Perceptions are perceptions. Just look (and listen). They are presumably representations of a hypothetical "real world". The system does not need the know "what those perceptions are" in order to control them; it controls perceptions (in the model, perceptual signals that are functions of external reality). The only control systems that have epistimological problems are the ones that are trying to understand the nature of control systems.
HAPPY BIRTHDAY MARY

Regards

Rick

Date: Thu, 7 Nov 1991 14:25:40 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Bickhard and encodingism

[(S)Martin Taylor 911107x]
(Rick Marken 911107)
>
>Someone said to Bill something like "So and so says that 60% of understanding
>behavior is a matter of understanding perception".
>
>Bill said "I think so-and-so is about 60% right".
>
You mean 36% of understanding behaviour is a matter of understanding perception?

No--don't answer that. It's the silly season.

Martin Taylor

Date: Thu, 7 Nov 1991 15:07:01 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Config control; modeling

[Martin Taylor 911107 14:30]
(Bill Powers 911107.0630 -- Bill, you get up MUCH too early) >
>
>Martin, you're saying that Joe's concept is a "very restrictive notion of
>the functions of neural nets." So is yours, if you mean to rule out Joe's
>concept.

No, of course I don't want to rule him out. I just want to point out that SOME kind of nonlinearity is needed. He seemed to say that there was only one kind that mattered.

>The question

[about whether a hierarchic linear system can be emulated by a one-level system. MMT]

>and its answer are not relevant to a model of a real
>behaving system. It is possible to trace out at least some hierarchical
>loops in real nervous systems: therefore the correct model will be some
>sort of hierarchy. It doesn't matter whether, after the fact, the same
>functions could have been accomplished with a one-level system. They are
>not accomplished that way. If a system is organized as a hierarchy, it
>will have different failure modes from one that is not; it will have
>different time-delays, and different partitionings of time-delays. Its
>evolutionary development will be different. Its stages of growth will be
>different from that of a one-level system.

If this is true, then the answer to my question is "No, it cannot be so emulated." Then, the issue I was trying to raise can be addressed. In what way is a linear hierarchic control system more powerful than a flat multidimensional one? If it is not, and is more complex, then the real system would probably not have evolved that way.

But, as I said in the same posting (I think), if nonlinearities are accepted as being there anyway, the question becomes moot. One just has to be careful that the nonlinearities do not invalidate any claims based on superposition. Sometimes they may, sometimes not.

>"What nets do well is to bind the perceptual space into topological
>neighbourhoods, in contrast to rule-based systems which treat the
>perceptual space as being built up from a myriad of unrelated entities."
>How does this explain the fact that we do employ rule-based systems
>(language, mathematics, etc.) and that we do categorize and frequently
>have to resolve conflicts (i.e., make decisions)? I point to the early
>part of this post: one computational model is not going to cover all the
>types of perceptual variables we can control at all levels of organization.
>
I had no intention to make exclusive categories. Neural nets often make very good classifiers, and the nets I am working with (that exhibit chaotic behaviour) categorize very well and sharply, because they show catastrophic rather than sharply bounded changes of behaviour. Most neural network classifiers simply have assigned thresholds or use the strongest output, or something like that. I think, as you do, that it is important that the system be self-energized so that the control input only has to affect a much more powerful process, rather than having to actually drive it (see your discussion on attractors versus control systems a few days ago). My underlying claim is that the success people have had with making neural nets do classification obscures the fact that they are even better at doing useful transformations. The natural unit on which a rule works is a category. The natural unit on which a neural net works is a region of similarity in some structure of its input and/or output.
>
>You're going to have to give me some help with binding perceptual space
>into topological neighborhoods. Is that a classification of something
>that could possibly be illustrated more intuitively?
>
One man's intuition is another man's absurdity, but I'll try. Usually with a neural net, a small change in the input results in a small change in the output. This is in sharp contrast to a rule-based system, in which there is no a-priori connection between the change of input and the change of output, and indeed the very concept of "smallness" of

change of input is hardly valid for that reason. In a rule-based system, based on words, for example, it makes a big difference whether I say "we can now do it" or "we can not do it". Is the change from "now" to "not" big or small? It's a non-question, isn't it?
Non-linear effects in the network may alter the magnitude of the relative change, and in a complex network one can get quite similar outputs from inputs that are in many ways substantially different, such as translationally invariant outputs from inputs in which most of the pixels have changed from black to white or vice-versa. The binding of perceptual space makes " A" close to "A ", and at that level the various placements of "A" form a topological neighbourhood. At a higher level, one could imagine a net deciding that a penguin was like a fish, because in many ways its underwater actions are like those of a predatory fish. The same net might decide that the penguin was like a bird. The penguin is in the same topological neighbourhood as fish and as bird. The net would have bound it there.

I (mis?)read into your comments that you think I think a hierarchic linear control system SHOULD be represented as a flat multidimensional one. That's wrong. What I am trying to get at is the potentially great increase of power provided by the nonlinearities. In your comments you point out many of these as if they were counter to my position. I want to bring out into the open the implications of the non-linearities,

>the fact that we do employ rule-based systems
>(language, mathematics, etc.) and that we do categorize and frequently
>have to resolve conflicts (i.e., make decisions)?

The interesting points are the bifurcations, the places where choices must be made--something no linear approximation can simulate.
Martin Taylor

Date: Thu, 7 Nov 1991 22:39:50 GMT
From: cam@AIFH.ED.AC.UK
Subject: Re: Config control; modeling

[From Chris Malcolm]

Bill Powers (911107.0630) writes:

> Chris Malcolm, Joe Lubin, Martin Taylor, Oded Maler (911106) --
>
> Martin asked:
> >Is is true that the effect of any hierarchy of linear control systems
> >can be emulated by some one-level multidimensional linear control system
> >that incoporates all the degrees of freedom embodied in the elements of
> >the hierarchy?
>
> And Chris commented:
> >And if so, is it necessarily the case that the hierarchical system is
> >the more economical implementation?

And Bill replied:
> The question and its answer are not relevant to a model of a real
> behaving system. It is possible to trace out at least some hierarchical
> loops in real nervous systems: therefore the correct model will be some
> sort of hierarchy. It doesn't matter whether, after the fact, the same
> functions could have been accomplished with a one-level system. They are
> not accomplished that way.

You forget that some of us are interested in building robots. Consequently we are interested in how biological systems do it, as clues to good design ideas. But there are some things (such as wheels, or the von Neumann computer) which biology eschews, but which

are nonetheless useful in artificial animals. Very often nature has done it the most economical way, but sometimes screws it up, such as the inverted retina in the vertebrate eye.

If I had a _proof_ that hierarchical control systems were more economical than flat multiple ones, it would be a lot easier to get research funding!

Date: Thu Nov 07, 1991 7:35 am PST
From: Revised List Processor
EMS: INTERNET / MCI ID: 376-5414
MBX: LISTSERV@vmd.cso.uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
CC: POWERS Bill
EMS: INTERNET / MCI ID: 376-5414
MBX: CZIKO@uiucvmd.bitnet
Subject: Your removal from the CSG-L list

Dear subscriber,

As of Thursday, November the 7th of 1991, you have been removed from the CSG-L distribution list (Control Systems Group Network (CSGnet)) by POWERS Bill <CZIKO@UIUCVMD>.

Virtually,

The LISTSERV management

Date: Thu Nov 07, 1991 7:40 am PST
From: g cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: g-cziko@uiuc.edu

TO: * Dag Forssell / MCI ID: 474-2580
Subject: Re: Vacation

>Please stop mail or whatever the command is.

I cannot set the option to "nomail" for you since only an individual subscriber can do this. I can, however, delete you from the list, so this is what I did.

To get back on when you return, you should send the following message to LISTSERV@VMD.CSO.UIUC.EDU

subscribe csg-l FORSSELL Dag, city, state, etc.

To stop mail temporarily you should send the following command to the same LISTSERV

set csg-l nomail

and then to turn mail back on when you return

set csg-l mail

Have a good trip.--Gary

==========Gary A. Cziko Telephone: (217) 333-4382 University of Illinois FAX: (217) 244-0538
Educational Psychology Internet: g-cziko@uiuc.edu (1st choice)
210 Education Bitnet: cziko@uiucvmd (2nd choice)
1310 South 6th Street N9MJZ
Champaign, Illinois 61820-6990
USA

Date: Fri, 8 Nov 1991 07:05:12 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: Test; config control; linear transform

[From Bill Powers (911107.1000)]

We seem to be getting more deeply into modeling, which is satifying to me. Some loose ends from earlier posts:

Oded Maler (911104) --

>... the problem of *correspondence* between some perceived combinations
>of sensory signals and what happens to the organism in the "real" world,
>is sometimes under-emphasized. Even if the organism behaves according to
>simple perecption-based rules, in order to show that this behavior
>achieves something in some external physical environment, you can get
>involved in very complex models which are very hard (if not impossible)
>to analyze.

This is why we are still at the level of stick-wiggling models. However, using the test for the controlled variable, it should be possible to pin down just which effects of behavior are relevant, saving a lot of trouble chasing after non-controlled effects. This is equivalent to seeing what the system is perceiving and what the reference condition is. Then the modeling problem becomes that of trying to guess how a variable of that type is perceived. This requires going up the levels one at a time, because the elements of a higher perception are those of the lower type. You have to understand the lower-level computations first, because they don't need to be done again at the higher level.

(911105) --

>This [range of control] might be meaningful when we move up to multi-
>dimensional systems, because some trajectories for correcting a
>"virtual" signal may go outside the controllability region of the
>elementary inputs.

A position control system will head directly toward the reference position without regard to intervening obstacles. The trajectory is whatever results from the way muscles are used on the way to the goal -- it isn't controlled at the configuration level. Normally, however, the actual configuration equals the reference configuration, so trajectory planning is really planning how to change the position reference signals. I think that "planning" is the wrong word to use at these low levels, however (an aside to Martin -- we can get into that later). Outputs are not planned -- not if you want to deal with disturbances.

I think the problem about trajectories is related to the reason for development of higher levels.

>Also, concerning the previous note, if "real" time is taken into
>account, it is clear that it might take longer to correct a larger

>error signal.

In a linear system, it doesn't. If the error is larger, the adjustment to lower-level reference signals is larger so the motion is faster. A step-change in reference signal will produce an asymptotic approach to the final state that simply scales up in amplitude (including any instabilities). In a nonlinear system of the most usual kind (loss of response-slope at higher signal values, or saturation), this is no longer true: larger moves take longer.

An experiment for which Rick can provide the reference shows that human arm position control systems are quite linear. Start with both hands in front of you on the edge of a table. Pick a near target on the table for the left hand and a far target (four or five times as far away) for the right hand. Then QUICKLY place the forefinger of each hand on its target. The reason for doing this QUICKLY is so that the reference position snaps to the target position before the hand has moved significantly. You will find that the fingers arrive at their targets at essentially the same time. The far finger certainly does not take four or five times as long to travel. This is indirect evidence that the position control systems are reasonably linear.

In our tracking experiments, we use a linear control-system model (even when the environmental connections are made nonlinear). This linear model produces simulated motions that differ from the real ones by only 3 to 5 percent of the range of the real movements. So making the model nonlinear could gain us only a few percent in accuracy. The best model, by the way, does include a delay (about 0.16 sec).

Martin Taylor (911105) --

>For linear transformations such as rotations and translations in a
>multidimensional space, it is the case that any sequence of
>transformations can be replaced by a single linear transformation that
>takes the initial configuration into the same final configuration as the
>sequence. This will be true whether the sequence of transformations
>involved manipulations of lower-dimensional independent subspaces or
>manipulations of the total space.

"Linear transformation" is in the eye of the beholder. If you start with a point (say a corner of a square) on the retina at coordinates x,y, and want to rotate it (with the square) by some angle, the new coordinates x', y' are

$$\begin{vmatrix} x' \\ y' \end{vmatrix} = \begin{vmatrix} x \\ y \end{vmatrix} \begin{vmatrix} \text{Cos(angle)} & \text{Cos(angle)} \\ \text{-Sin(angle)} & \text{Sin(angle)} \end{vmatrix}$$

Sine and cosine functions are certainly not linear; they aren't even single-valued. Translations in x and y are linear. But when you do BOTH, you can't get away from the nonlinear functions even if you use different coordinate systems. A rho-theta coordinate system gives you linear transformations in angle, but nonlinear ones in translation.

I should have mentioned earlier that while eveyone seems to be referring to our subject as "configuration control," weighted sums apply in my model only to sensation control. See p. 105 of BCP. We have really been discussing what I think of as sensation control.

Configuration control can't be done by linear weighted sums alone (I

think) because extraction of invariants must use things like the rotation matrix above. In stereo vision, it's clear that we can recognize "a square" that is tilted away from us so its sides no longer appear of the same length on the retinas. As the square changes its orientation in 3-space, we still recognize it as the same figure. We can read letters on a page that is not orthogonal to the line of sight. Consider the letter "H". If we see a big H spelled out on a football field by a marching band, its projection onto the retina might actually look more like an "A" than an "H". But we see it as an "H".

One way this might be done:

Suppose there are three control systems, internal to the perceptual function. These systems adjust the properties of a neural network; they don't produce behavior.

One of them applies a translation to move the attention to the centroid (or something) of the figure, or vice versa. The amount of the translation-signal becomes a perceptual signal representing the position of the figure relative to the direction of looking.

The second one applies an x-y or radial scaling factor centered on the direction of looking to shrink or expand the figure to match a reference-figure that is the hypothesized "standard shape". The amount of scaling signal becomes a perceptual signal representing the absolute size (on the retina) of the figure.

The third control system applies a rotation angle about the center of attention to make the orientation of the figure match that of the "standard shape." The signal that adjusts rotation becomes the perceptual signal representing the orientation of the figure.

The total error signal in these three control systems indicates the degree of mismatch. If the total error signal is small enough, the identity of the figure is considered recognized. The other three signals then indicate variable attributes of the figure. The three control systems operate simultaneously and independently of each other. In three-dimensional shape recognition, there would have to be three size systems and three rotation systems. If the translation system sent its result to the derotation system and the result of that went to the scaling system, the latter two systems could be two-dimensional. I think that this model is buildable as a simulation. I think, too, that in making it work we might learn that the configuration level is really several levels -- it would be fun to try to indentify them in experience once we see what's required.

Note. The Little Man in C (major) is coming along. Getting the pointer formats to work for the coordinate transformations was a bitch. Now at least the 3D presentation with variable points of view is done. Next the control systems.

If there are any C gurus out there, I would really appreciate a tutorial on the syntax for typing variables. It still doesn't make sense to me -- I get it right by trial and error.

I use this typedef:

```
typedef struct {int x3d,y3d,z3d;} polytype3d[100];
```

and polyg is

polytype3d *polyg

The pointer notation I had to use was

(*polyg)[i].x3d = -halfside.

Why did I have to put the parentheses there? I sort of get the idea that *polyg refers to the array of structures; empirically I found that (*polyg[i]) refers to the ith array of structures and not to the ith structure in the array -- but what's the general rule? There must be some nice mechanical way to parse this stuff so I don't have to use my feeble brain.

C certainly is dumb -- it let me pick the ith array even though the arrays weren't themselves part of a defined array -- it just took the next bunch of memory, which I discovered by accident using the debugger. I had a row of polytype3d variables defined on one line, and suddenly the debugger gave the wrong name to the array I was looking at, or thought I was looking at. Light bulbs. No wonder the display looked so funny.

==========

Date: Fri, 8 Nov 1991 15:57:09 +0100
From: Oded Maler <Oded.Maler@IRISA.FR>
Subject: Re-inventing the wheel is a real fun

[From Oded Maler 911108]

It's hard to control a multi-party communication cuncurrently with your future position, your posture (especially during hang-over) and a lot of other things (including a new mailer).

So without referring explicitly to particular recent posts, I'll try to summarize my point of view in an unorganized manner. I`m not sure I'll be able in the future to participate in the discussion with the same degree of intensity, and I apologize in advance for any unresponded message.

The basic model of PCT is a hierarchy of servo-mechanisms. Each level of the hierarchy controls a different type of perceptual variables. Since the essence of the feed-back paradigm of goal-directed behavior lies in the notions of order (in the mathematical sense, lesser/greater) and of distance, my first concern was to see how these notions are applied to the second level variables, i.e., configurations which are some spontaneous combinations of elementary intensities. I accept it by now that certain configurations can be controlled this way and the assumpption employed in order to prove it or not worse than those usually used in various models of mathematical control theory [when these assumption are violated, you may have Parkinson disease].

As a side sociological remark, I will repeat that what most people call Control Theory, is mathematical engineering discipline, which has rather progressed a lot since the times of N. Wiener. For example, the recent proceeding of the ACC (American Control Conference) is a 9-volume, several K-pages, collection of hundereds of papers, dealing essentially with the same problem of making such-and-such environment stay within certain boundaries, given such-and-such information sources etc. Now, I am not a control-theorist in this sense, and I cannot read most of these papers, and I know that sometimes the mathematics there is just an

effective barrier against strangers, and I know that that the awareness of most people in such big communities to the origins of their discipline and to the really interesting questions is very low, etc. But still, one cannot ignore the fact, that those questions of controllability, observability and realizability that we have been discussing recently, rather naively, are the main concern of a very large community, whose models are very precise and whose standards of rigor are high.

Coming back from this side (and, really unnecessary) remark, I'll try to state the major problems associated with the higher-level variables. The next levels introduce some primitive notion of time/memory. As I understand it, the level above the configuration level may impose a time-varying reference signal. How does this level operate? Clearly, the peceptual variable it controls is a function of both a limited history of configuration variables and some internal clock. I think I can imagine how it works in principle. The servo mechanism can still work here but the virtual inpput [I moved from Suntools to X windows which seems to have a larger sampling rate from the keyboarddddddd] signals, are partly coming from the inside (clock, memory [e.g., integral of configuration values over the recent temporal window]). The problem here is, what kind of references are given to this type of systems. I guess it is not a look-up table of arbitrary functions, but rather a set of parametrized functions (e.g., sinusoidal waves), and the reference signals can be the parameters...

Ooff, I'll stop climbing up because I see it will take rather a long time to arrive to the level in which I can consider myself (a humble automata-theorist) knowledgable - the level of discrete sequences and state-transitions. In fact I was working recently on trying to extend methods for specifying and verifying isolated software systems, into methods for reasoning anout hybrid systems consisting of digital programs controlling continuous environments (I don't recommend for reading - formal logic can be as disgusting as differential equations or matrices). This is one of the reasons I was concerned about implicit timing assumptions.

What I want to add before closing, is that it is really important not to take the levels too literally. For example, if you look at a basic sense organ in a small space-time-scale you can detect inside it, maybe a small detector of temporal patterns.
Another example consider inter-personal and inter-cultural differences. For an inter-personal example, when I make a coffee in an unknown environment, I first take the coffee and the sugar from wherever they are and put them in my favorite position, while other people with better low-level control just launch the spoon to where the ingredients are. So the same functional task is performed by different "levels".

It's enough.

--Oded
==========

Date: Fri, 8 Nov 1991 07:57:05 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>

Subject: flat vs bulky models

[From Bill Powers (911108.0800)]

Martin Taylor (911107) --

I think I get up too early, too. Looks like I will soon have a logon at Fort Lewis College (and a sort of honorary adjuct professorship to make it legal). That will make communcating into a local call and I will no longer have to beat the clock.

I'm beginning to understand your approach to perception -- that is, not to misunderstand it so reliably. I especially appreciate your statement that neural nets are good at transformations in general, not just certain types. Have you looked into what is required by way of pre-organization to make the job of adapting to a specific environment easier? Perhaps there's a compromise between my position that different kinds of computing components are required at each level of perception, and yours that deals with the general case. I suspect, from the little I've seen, that most neural network models make a lot of assumptions about the kind of connectivity that is available before learning starts.

>The natural unit on which a neural net works is a region of similarity
>in some structure of its input and/or output.

Similarity in what terms? Defining similarity, it seems to me, is like defining invariance: you have to specify in what respect two things are alike, and in what respect they can differ. I'm a little confused by your reference to "structure." Does the network detect physical structure?

>Non-linear effects in the network may alter the magnitude of the
>relative change, and in a complex network one can get quite similar
>outputs from inputs that are in many ways substantially different, such
>as translationally invariant outputs from inputs in which most of the
>pixels have changed from black to white or vice-versa.

How do such similar outputs relate to perceptual phenomena? I can think of a nonlinear network that responds to the mean peak amplitude of retinal signals; this output would be invariant with respect to position, rotation, size, and so on -- everything but intensity. Yet it wouldn't be recognizing position, rotation, and so on. It would just be ignoring them. How do you get such a network to respond to, say, orientation while ignoring translation, amplitude, and size?

>What I am trying to get at is the potentially great increase of power
>provided by the nonlinearities. In your comments you point out many of
>these as if they were counter to my position. I want to bring out into
>the open the implications of the non-linearities

Sorry -- there I go again. I tend to think of nonlinearities in a continuous control system as a disadvantage, because they can make control more difficult (they introduce higher harmonics into regular variations and complicate stabilization). I tend to forget that at higher levels, we want nonlinearities to make logical operations possible, and (it seems) even more operations that I haven't grasped yet.

Can you tell me what perceptual phenomena you are trying to model?

Chris Malcolm (911107) --

>You forget that some of us are interested in building robots.
>Consequently we are interested in how biological systems do it,
>as clues to good design ideas.

Right. I think the design hints are there. In the Little Man V. 2, using the stretch and tendon reflexes almost "out of the box" gave me a stable model with extremely little difficulty. It seems that controlling acceleration, then velocity, then position, in a hierarchical manner, is the simplest way to go. Even with linear systems.

General question: if the ENVIRONMENT is nonlinear, does that also rule out a "flat" system? If so, the theorem people seem to be looking for is readily available. The environment is nonlinear in many ways (things bumping into each other, trigonometric and vector relationships, general dissipative phenomena). This matters when you deal with closed loops.

Got to go

Best to all

Bill P

==========

Date: Fri, 8 Nov 1991 14:30:55 EST
From: Joseph Michael Lubin <jmlubin@PHOENIX.PRINCETON.EDU>
Subject: thresholds and only thresholds; Nature wrong: retina backwards

[From Joe Lubin (911108.1200)]

Martin Taylor, Bill Powers, Chris Malcolm, Oded Maler (911107,911106) --

YUCK YUCK YUCK.

I didn't mean to start a discussion or a controversy.
I mostly wanted to add to Martin's first post

> In perceptron-like neural network structures, the power of multiple levels
> comes from the fact that the individual layers perform non-linear
> operations (they saturate).

that many neural networkers (the conventional variety who only consider perceptual-type architectures) consider the thresholding operation to be the fundamental nonlinearity and often discount the importance of the squashing-type operations.

I totally agree with Martin's second post

> That's a very restrictive notion of the functions of neural nets. The
> need for a decision is imposed from outside, in the experimenter's choice
> to use the net as a classification device. Nns are quite good at that,
> but I think that to treat them as primarily classification devices is
> to ignore most of what they do best. Typically, there is no need for
> a decision until there is a need for behaviour. The nets usually do
> not control behavioural choices. What nets do well is to bind the
> perceptual space into topological neighbourhoods, in contrast to rule-based
> systems which treat the perceptual space as being built up from a myriad
> of unrelated entities.
>
> Actually, I don't care which kind of non-linearity you introduce. But you
> must have some to do more than some kind of generalized rotation of the
> perceptual space, such as a Fourier transform.

and should have stated so but I felt that I had nothing to add.
I should have realized that I was speaking to people who
generally try to address the full scope of sensory-motor
processing and are not interested merely in models of
categorization.

And now to stir up another can of worms ...

Chris Malcolm (911107) --

> You forget that some of us are interested in building robots.
> Consequently we are interested in how biological systems do it,
> as clues to good design ideas. But there are some things (such
> as wheels, or the von Neumann computer) which biology eschews, but
> which are nonetheless useful in artificial animals. Very often nature
> has done it the most economical way, but sometimes screws it up,
> such as the inverted retina in the vertebrate eye.

The photoreceptors need nourishment and pigment recyling
functions done by the pigmented epithelium behind them. And this
epithelium, in turn needs to be connected to other structures,
(blood circulation, etc.) so it can only be at the back of the
retina. If this layer of cells were in front of the
photoreceptors, most photons would be absorbed before being
registered by the transducing elements, the photoreceptors. The
epithelium is also responsible for anchoring the photoreceptors
in a very precise spatial array.

Since the retinal network is quite translucent, it is not much of
a problem having it in front of the receptors. At the fovea
(high resolution area) the retinal network is mostly pushed to
the side to remove the mild distortions it would cause.

[Interestingly, retinal blood vessels (which do not feed the
receptors) lie within the retinal array and absorb photons.
These produce no psychologically noticeable visual effect on the
transduced image even though they do obscure it. This is due to
a visual mechanism which "fills in" occluded regions (like the
blind spot where the ganglion axons leave the eye for the optic
nerve). It is possible that such a "filling in" mechanism may
have evolved in response to the variety of occlusions that exist
because the receptors are at the back.]

Other thoughts:
(i) the farther back in the eye the receptors lie the larger is
the surface that the receptors must tile, so given fixed-diameter
receptors, you get a higher resolution lattice.
(ii) receptors, even if out-front, should have their
pigment-absorbing outer segment pointing to the back of the eye
so that their can be a funnel-shaped waveguide inner segment to
halp capture photons. This has signal-to-noise, and sensitivity
implications.

So having the receptors out in front is not necessary and

probably not a good idea for high resolution vision.

Date: Fri, 8 Nov 1991 23:05:00 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: levls

[From Bill Powers (911109.1000)]

Chris Malcolm (911107) --

I've had some thoughts about why we have a hierarchical system instead of a one-level system. I think I alluded to this subject previously, but maybe I didn't elaborate on it enough. You're asking how you can get grant mony for looking into hierarchical robotic systems. I think you can make a case that putting money into any other kind is a waste of resources in the long run.

The main reason for the hierarchy is modularity. Lower-level systems are complete control systems in themselves, capable of bringing variables to reference conditions and keeping them there despite disturbances tending to alter the variables. The tendon reflex controls sensed tension in the tendon, which is to say the force applied to the attachment. This force and the angular acceleration at a joint are now precisely controllable even if the muscle's response to signals changes, and even if there are external disturbances tending to alter the force or the acceleration. The phasic stretch reflex adds damping, so the limb containing the muscle can be dynamically stable, taking care of inertial forces by itself and allowing position control to be carried out by a simple proportional angle-sensing system. The phasic stretch reflex is a velocity control system; it makes sensed velocity match commanded velocity by altering the reference signal for applied force.

If this same arrangement were to be accomplished by a single-level system, then the position-control system would have to contain the damping factors and would have to compensate for external loads, changes in mass, and fatiguing of the muscles.

That could be done. But what if you now want to build a system in which the arms are used together, suspending a load between them? If you're allowed only one level, you have to put the old system aside and build a new one having different connections and different stabilization requirements. In effect, the new system has to duplicate all the features of the old ones for controlling each arm, plus whatever is required by the need to operate both arms with a common load.

When you consider even slightly higher levels of organization, the duplication gets worse and worse. You need to handle muscles differently when you drive or walk; when you are walking up stairs or up an escalator; when you are walking, jogging, trotting, or running to win; when you are walking or running in a straight line or around a curve. And if you go a little higher yet, you have to incorporate complete duplicates of all those duplicated systems when you run away from danger or toward someone in order to give help, when you walk to improve your fitness or to work off anger.

If you're just building a robot to do one job under unchanging circumstances, there's no choice between flat and hierarchical models, except perhaps for easing the designer's job. But if you want the robot to be useful in varied circumstances and to accomplish different types of goals, modularity is the only sane way to go.

The human system is modular to an extreme. Every critical variable is under local control -- very little of importance happens open-loop. It's as though organisms learned early -- at the soup stage -- not to trust nature. When you send a signal to a muscle telling it to create a force, there's a sensor that comes back with the report, "generating this much force, boss." When you tell a joint to assume angle X, joint receptors send back the message "Feeling this angle, boss." When you tell the oculomotor systems to look at that object over to one side, the retina comes back saying, " Whatever it was you wanted to see, boss, here's what's in front." When you send a signal to your adrenal cortex telling it to jazz up the system some, you get back a message saying "I'm this excited, man." In fact all these commands are reference signals specifying what is to be sensed, not what is to be done. The local control loop automatically adjusts its output, doing whatever is needed to make the sensed result match the requested result.

I've used the principle of non-duplication of function throughout my model. The imagination connection is where it is because I didn't want to have to have a whole separate set of perceptual functions just to recognize internally-generated memory signals. The imagination connection runs the output of memory into exactly the same perceptual functions that are used for handling sensory inputs. I didn't want separate memory devices for perceiving and for doing, so I put the memory in a position where it could provide either remembered/imagined perceptual signals, or reference signals for lower systems.

The properties of feedback systems see to it that you don't need one organization to organize incoming information, and a separate one that does the inverse thing to create structured output commands. As I showed in my algebra for Oded, the perceptual organization plus the closed-loop process is enough to make the outputs have the required structured effects. You don't need an output system that elaborates commands into the appropriately-weighted outputs. The input organization does the main part of the job. All the output systems have to do is map error signals into changes in lower-order reference signals that will tend more to reduce the higher-level error than to increase it. The output systems can even contain many wrong connections, as long as they're outnumbered by right connections.

In the model as a whole, the division into levels and the subdivision of levels into independent control systems means that there is modularity at all levels: each level consists of general-purpose control systems of a given type, available for an infinite variety of uses by whatever systems you want to add at the next level. This makes sense in design terms, and in terms of development, evolutionary or within one lifetime. If you want to build a self-organizing robot, make the system modular. That is, hierarchical.

Best

Bill P.

==========

Date: Sat, 9 Nov 1991 20:10:00 CST
From: TJ0WAH1@NIU.BITNET

Subject: epistemology

[from Wayne Hershberger]

(Bill Powers (911104.0700)
>I agree with you. Reality is precisely what we experience. But
>to add "and nothing more" is to assert what we can't know
>directly, and to deny, for no good reason, the implications of
>physics, chemistry, and neurology, all of which claim to
>represent a world of immense detail that is inaccessible to our
>senses. You have never answered my comments about physics and so
>on.
     It would be a comfort were you to agree with me, but I do not agree with the statement, "Reality is precisely what we experience." Of course, I agree even less with the obverse idea.

That is, as an empiricist, I do not endorse the idea that reality is precisely what CAN NOT be accessed empirically. Hence, I can not imagine ever saying that physical, chemical, or neurological phenomena are epistemically inaccessible. In fact, I have been championing the antithesis. I guess I am not making myself clear.
     Let's go to square one. I claim that the epistemological challenge is not to explain how the truly inaccessible can be accessed (a logical impossibility), but rather to explain how the truly accessible can have appeared to be inaccessible (a logical possibility). Do you agree?

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu
==========
Date: Sun, 10 Nov 1991 09:47:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: retina,,little stick man

[From Wayne Hershberger]

Joe Lubin (911108.1200)

I enjoyed reading your defense of nature's design for the chambered eye. Let me offer you another round of ammunition.
You say,
>Since the retinal network is quite translucent, it is not much
>of a problem having it in front of the receptors. At the fovea
>(high resolution area) the retinal network is mostly pushed to
>the side to remove the mild distortions it would cause.

The foveal pit you're describing here would also function as an optical element (a negative lens) collimating the pencils of light so that the refracted rays are better aligned with the long axes of the rod-shaped foveal photoreceptors.

Bill Powers (911109.1000)
>The phasic stretch reflex adds damping, so the limb containing

>the muscle can be dynamically stable, taking care of inertial
>forces by itself and allowing position control to be carried out
>by a simple proportional angle-sensing system.

Does this imply that the new version of the "little stick man"
program includes an angle-sensing system as well as a muscle-
length sensing system?

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu

==========

Date: Sun, 10 Nov 1991 10:56:45 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: Epistemology

[From Bill Powers (911110.0900)]

Wayne Hershberger (911109) --

Epistemology:

I'm game if you are. I keep having this feeling that if I could just find the right words, the whole thing would come clear. This is undoubtedly a pernicious fallacy, but it keeps me going.

>Let's go to square one. I claim that the epistemological
>challenge is not to explain how the truly inaccessible can be accessed
>(a logical impossibility), but rather to explain how the truly
>accessible can have appeared to be inaccessible (a logical possibility).
>Do you agree?

Square one it is. I agree that the truly inaccessible can't be accessed. That's a definition, not a proposition. The truly inaccessible is that to which we have no access at all, either direct or indirect.

The second part of the statement is not a definition: it asserts that the accessible sometimes appears inaccessible. I think this is an attempt to create a two-valued situation (either something is accessible or it is not) out of one that has more than two possible values.

I see accessibility of the workings of nature to be a matter of degree, with the maximum degree falling short of 100 percent. To explain my view, I will resort to a thought-experiment.

Suppose we have before us a black box which we have no means of opening. Let's call this box, to humor me, the Reality Box. On its surface are numerous buttons and lights. The buttons and the lights are undoubtedly real, because we experience them directly and unequivocally: they are totally accessible, and cannot be mistaken as being inaccessible to our observation. So the box, the buttons, and the lights are not an epistemological problem.

When we press various buttons, we find that certain lights and combinations of lights turn on or off. With sufficient experimentation and record-keeping, we can discover consistent effects of the buttons on

the lights. As our experience grows, we can discover that some buttons alter the effects of other buttons on lights, or make certain lights come to depend for their state on new combinations of buttons. We can discover that only certain sequences of button-pressing will have predictable effects on one or more lights. We can find that certain lights have mutually-exclusive states; if a member of one set is on, another set is always off, and vice versa. We could uncover logical relationships, relationships corresponding to arithmetic operations, and so on. We could even develop heuristics; some ways of turning lights on work best under one set of circumstances (combinations of lit lights), usually, than other ways.

Thus we arrive eventually at a sophisticated empirical understanding of the Reality Box. At no point have we asked what is inside the box. We have simply observed, recorded, and tried to recognize consistencies.

Perhaps I should motivate this investigation by saying that for reasons we only vaguely understand, certain of the lights on the Reality Box have extraordinary value to us; indeed, their states of illumination seem to us to be a matter of life and death, or at least make the difference between enjoyment and disappointment. So we have an interest in pushing the buttons to keep the most important lights in the states that seem the most desirable, especially as they will not stay in those states without the button-pushing. To be crass and less mysterious, we could say that each time we succeed in maintaining the critical lights in the critical state for one minute, we receive $5 -- that's $300 per hour for this job if we can learn to do it perfectly.

I think I have now described the state of human understanding of nature in the pre-Galilean era.

In fact, we find that we are a long way from making $300 per hour -- the actual payoff isn't nil, but it's just barely a minimum wage. It would be greatly to our advantage if we were allowed to cheat: to open the box and trace out the circuits (or talk to the little men, or analyze the chemicals, or take whatever action is appropriate to what we find in the box). If we knew WHY the buttons affect the lights as they do, we could abandon the trial-and-error empirical approach and simply deduce the actions that would have the effect we want.

I am now describing the advent of the physical sciences.

We are not, however, allowed to cheat. Nobody knows how to open the box.

Nevertheless, once we get the idea of explaining the dependence of lights on buttons rather than just observing it, we might well decide that even a good guess about what is in the box might be more valuable than random experimentation. So we begin to construct a model of the internal workings of the Reality Box, trying to outguess its designer.

This project turns out to be extraordinarily successful. By imagining circuits and functional devices inside the box, we succeed immediately in explaining why some buttons cause some lights to change their brightnesses. Numerous revisions of the model are required, however, because just when we think we have the right connections, an anomaly turns up and we have to modify the design of the hidden devices or the connectivity between them. But by demanding that the model ALWAYS work, no matter what combinations of buttons we press, we eventually get this model to the point where it never fails in any way we can notice.

We now begin to believe in the reality of the model. What APPEARS to be happening is that buttons activate lights, but what is REALLY happening, we say, is that the buttons are feeding their effects into a hidden complex device that in turn operates the lights. Gradually the status of the insides of the Reality Box changes. Those insides no longer seem hidden to us. In fact, even though they are complex, they are far simpler than our records of empirical findings are. They also permit us to predict the effects of button-pushings, even combinations never tried before, with exceeding accuracy, whereas our empirical predictions, based only on unexplained frequencies of occurrance, are wrong nearly as often as they are right, and are essentially useless in unfamiliar circumstances.

I have now described the rise and maturation of the physical sciences, and their essential difference from the purely empirical sciences.

The penultimate stage in this development arises when someone notices a fact that by now is considered a very strange fact. Those who are engaged in the exploration of the Reality Box by now feel that its devices and connections are IN THE BOX. The lights tell them what is happening inside the box; the buttons let them influence what is happening inside the box. What this someone says that is thought so strange is merely a reminder that in the beginning, nobody knew what was in the box because only the lights could actually be observed, and no effect of pressing the buttons could be seen except in the lights. This is, in fact, still the case. So the model of what is in the box must exist in the minds of those who are observing the box. It is not in the box. In fact, it is perfectly possible that what is in the box is entirely different from what is in the mental model, but is equivalent to what is in the mental model under all the button-pushing operations so far tried. Even what seems to be a simple direct connection through a hidden wire might actually involve a hidden modulator that converts the button-press to a radio frequency and broadcasts to a receiver whose output lights the light. That would not be a parsimonious design, of course, but it might be the one that exists.

What this upsetting stranger is doing is reminding everyone that all they can actually observe are the lights, and the only effect they can know they are having is to press the buttons. All the rest is imaginary. Therefore we should throw away all those figments of the imagination, and admit that all we know is how the buttons affect the lights, and to remain pure of heart we should talk about nothing else.

I have now described the advent of stimulus-response theory, behaviorism, biology, empirical psychology, and so on.

The final stage entails the epistemology of the Reality Box that I propose.

In fact, the model works much better than it should. Moreover, there is evidence in the relationship between buttons and lights that tells us something consistent is going on independently in the Reality Box. The lights that we can affect with our buttons sometimes turn on and off when we aren't pushing anything. Very often, we have to change which buttons we push in order to reproduce the same state of the lights, and there seems to be no way to predict when, by how much, or in what direction we will have to make these changes. Something else is interfering with the effects that the buttons have. This something else can be inferred, to some extent, because it may occur regularly, or in some regular pattern, as we can tell by watching what different buttons we have to press to reproduce the same effect, and when we have to do this.

So we are led, in the end, to recognize three major facts. First, our mental models of what is in the Reality Box have an unknown relationship to what is actually there. Second, the regularities implied by the model actually do occur, even through we can't know that they occur for the reasons we propose. And third, there is something in the Reality Box that can act independently of us. So we can say that in some regards, what is in the box is accessible to us, but we must also admit that our interpretation of its inner workings is not necessarily the only one that would be as good at explaining what happens.

The lights, of course, are our perceptions, and the buttons are our actions. The Reality Box itself is invisible; we experience only the input and output devices mounted on its surface. We conjecture that the buttons do something that we don't observe. We conjecture that the lights indicate something that is also not observed -- if only the presence of a wire from the button to the light, and an invisible power supply.

So what does "accessible" mean? Does it mean that we observe Reality exactly as it is, or that there is neither agency nor order other than what is evident to us in our sensory experiences? I feel that such questions are not matters to be deduced logically so that we can know once and for all the truth about experience. I think that they are matters to be settled as we settle all factual questions: by the examination of evidence, and by settling for the inference from the evidence that seems most supportable by all the rest of our experiences and knowledge. Pure philosophy can't provide that sort of conclusion: it demands an end-point, a certainty. That, I think, is definitely inaccessible.

Best re-gardes,

Bill P.

==========

Date: Sun, 10 Nov 1991 16:49:21 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Test; config control; linear transform

[Martin Taylor 911110 16:45]
(Bill Powers 911107.1000)

A quick response on technical terms:

>
>"Linear transformation" is in the eye of the beholder. If you start with
>a point (say a corner of a square) on the retina at coordinates x,y, and
>want to rotate it (with the square) by some angle, the new coordinates
>x', y' are
>
>
>
>
>
>
>Sine and cosine functions are certainly not linear; they aren't even
>single-valued.

A matrix multiplication like this is a linear operation. The fact that the constants in the multiplication could be different on other occasions according to a non-linear function of some other variable is irrelevant. All translations and rotations and scalings in a Euclidean space are linear operations. They are all describable by a matrix multiplication. But

if x and y appeared in the matrix, then the operation probably would not be linear (I say "probably" because in some instance it is conceivable that the quadratic terms in x and y might vanish).

I think we should follow Bill's earlier lead and say "Of course there are important nonlinearities; forget the linear analyses," and see where we get to.

Martin Taylor

==========

Date: Mon, 11 Nov 1991 07:49:02 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: Optics; nonlinearities

[From Bill Powers (911111.0800)]

Wayne Hershberger (911110) --

Re: retinal optics

>The foveal pit you're describing here would also function as an
>optical element (a negative lens) collimating the pencils of
>light so that the refracted rays are better aligned with the long
>axes of the rod-shaped foveal photoreceptors.

A Galilean telescope! Now all you need is another eye behind it, etc.

Re: angle sensors in little man model

>Does this imply that the new version of the "little stick man"
>program includes an angle-sensing system as well as a muscle-
>length sensing system?

No, not this time. In this model I assume a constant lever arm (pulley effect) for the muscles. The muscle length sensors (proportional component) therefore produce a signal proportional to joint angle, so adding joint sensors wouldn't provide any new information. After the C version of a merged Little Man and Dynamics model is running, Version 2, I'll start work on Version 3, which will include some nonlinearities, mainly the nonlinear relation between muscle length and joint angle (due to changing lever arm). Greg Williams has found some nice literature in which these and many other useful relations are given, so they can be put into table lookups for fast conversion between angle and length. Then it might make sense to use joint angle info for a linear representation of angle -- although the model seems to run just as well with the nonlinear kinesthetic perception of joint angle (I've tried it). The visual control system dominates -- little nonlinearities just make the arm move a little bit wrong, and the visual systems instantly alter the kinesthetic reference signals to correct the error.

Martin Taylor (911110) --

Re: rotation matrix

>A matrix multiplication like this is a linear operation. The fact that
>the constants in the multiplication could be different on other
>occasions according to a non-linear function of some other variable is
>irrelevant.

I was thinking of the problem of rotating and translating a map of the retina to match an image to a reference image. The x',y' coordinates

(after rotation) would be a function of the rotation angle. So the "constants" in the rotation matrix are actually functions of the angle: sines and cosines. Given some way to compare the rotated map in x and y with a reference map and generate a signal roughly proportional to the angular difference in orientation, the error signal would be input to the perceptual function to alter the constants in the rotation matrix according to sine and cosine functions. The amount of rotation-signal required to achieve the match would then be the perception of orientation.

At the same time, a second system could alter constants added to the elements of the rotation matrix, to achieve a translation-match.

I'm just doodling with designs here. Don't take this too seriously. There are probably less elaborate ways to accomplish what I'm thinking of.

Linearity:

I use linear models because they're easy to analyze and they compute fast. I've tested all sorts of nonlinearities and have verified that if they don't get too extreme, you get essentially the same (or easily predictably different) behavior. Just consider the comparator. It can be VERY nonlinear, because normally control processes keep the error signal very small -- you're always looking at a small span around zero error, so it doesn't much matter what the shape of the function is for very large errors (except at startup). Perceptual nonlinearities just redefine the controlled variable (and affect loop gain). Output nonlinearities act like changes in the output sensitivity with signal strength; as long as the incremental loop gain remains high enough, there's no important effect of the nonlinearities. In the arm model with dynamics, I've tried putting in the sixth-power force/length relationship that is the actual "spring constant" of the muscle. It just seems to make the arm a little more stable. You can't see any interesting difference in behavior from the linear muscle model. The main effect would be to increase the loop gain for large disturbances -- where sensory sensitivity is falling off.

>I think we should follow Bill's earlier lead and say "Of course there
>are important nonlinearities; forget the linear analyses," and see where
>we get to.

Agreed. I hardly ever do any analytic arithmetic any more. Simulations are easy to work with and let you put in all sorts of weird relationships that you could never handle by solving equations. If you suspect that nonlinearity might be important, just put some in and see. It takes ten minutes instead of a major computational project.

Best to all

Bill P.

==========

Date: Mon, 11 Nov 1991 11:44:23 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>

Subject: mmt: information (long)

[From: Bruce Nevin 911110 1747]
(Martin Taylor, various posts) --

I spent some time this weekend, Martin, reviewing some of our past correspondence, trying to sort out our divergent use of terms.

You identify "information" with the reduction of uncertainty.

> The information in a signal can never be determined by examining the signal itself. the information resides only in how the signal is usable by a receiver to change the probabilities the receiver assigns to possible states of the world. . . . The real information conveyed by the signal is determined by how concisely it . . . can be predicted. . . . That depends on what the receiver knows ABOUT the content of the message--in other words, how the message fits with the receiver's pre-signal ideas. . . . (7/11/91 17:24)

By "information" I refer to entities and relationships among them. Of course, the entities are perceptions, and the relationships such as dependencies are perceptions too. I will go through this here, and return to it later where hopefully it will make more sense if it does not do so now. A speaker presents certain linguistic entities and relations in an utterance. In the perceptual hierarchy of the speaker, these correlate loosely with nonlinguistic entities and relations among them. The linguistic entities, gross restrictions on their combinability (operator and argument classes, reductions), and the loose correlation with nonlinguistic perceptions, are socially shared because of the way in which they are socially acquired and socially maintained. In particular, they are shared by the hearer. Because of this, the particular linguistic entities and relations in the speaker's utterance constitutes an occasion for the hearer to consider nonlinguistic perceptions and relations among them that she might not otherwise have considered.

It is from this perspective that I said (quoted by you, ibid):

> Language is *not* a code. Formal relationships among parts of an utterance themselves constitute the information "in" or "carried by" that utterance. (Vessel metaphor or portage metaphor, you pick.)

To which you replied:

> That [is the] position [t]hat I am arguing against, above. The formal relationships can constitute only the "third-party" information--the information detectable by someone who knows nothing of what the partners are talking about. Information is NEVER computable without knowledge of the other possibilities and their probabilities. It is ALWAYS the reduction in uncertainty that the message permits. And in any case, the formal relationships constitute the redundancy of the message, not its information content.

Later (7/16/91 11:15), you wrote:

> . . . I believe the main presuppositional disagreement is on the definition of "information." . . . Information is always the reduction of uncertainty. Uncertainty is the lack of knowledge

> about some universe of possibilities, and it can be measured by the standard Shannon measure (sum p log p) over all the possible events, some subset of which may happen. After an event, the uncertainty is reduced, and the amount of reduction is the information conveyed by the event.

In all of this, "information" is identified with the reduction of uncertainty. This I take to be a numeric quantity that can be "computed" (your term above), determined by counting the number of choices available to particpants in the process of informing.

[Choices of what? You never say, but it appears to be choices of values for a particular X. What is X? A variable? A term in a proposition? These are metalanguage words referring to linguistic entities. That way lies a theory of the integration of the current discourse with (memories derived from) prior discourses, something I have proposed and you have disavowed. Choices for action? That way lies S-R theory. But let that issue lie for now. It would be a distraction at this point.]

Instead, note that a measure of something is not the thing measured. The reading on the gas guage is not the gasoline in the tank. The guage might be hooked up for tire pressure instead, or to an air-speed sensor. To illustrate this, imagine two communicative scenarios. In one, a man in a restaurant speaks to the waitress, and what he says reduces her uncertainty with respect to items in the menu by an amount n. The information in his utterance is, by your account, identified with this quantity n by which her uncertainty is reduced. In the second scenario, the man overhears his child talking with Santa Claus in a department store. What the child says reduces the man's uncertainty with respect to items that he had considered buying as Christmas gifts for the child by an amount m. The information for the father in the child's utterance to Santa Claus is, by your account, identified with this quantity m by which the father's uncertainty is reduced. But as it happens, the two quantities turn out to be the same number: n=m. By your account, then, must we say that the information is the same in the two cases? The reduction in uncertainty is the same computed quantity.

But no, you may say, of course not! The reduction of uncertainty is always with respect to a particular universe of possibilities--edible items on a menu in one case, gifts for a child in the other. You will admit then, perhaps, that the information in an utterance is not to be identified with the reduction in uncertainty as an abstract quantity "measured by the standard Shannon measure (sum p log p) over all the possible events, some subset of which may happen"? That the information has to do with something about the particular "events" subject to that uncertainty, and not with reduction in the uncertainty itself? Can you see, then, how I might say that information concerns entities and relations?

[Resuming: Bruce Nevin 911111 0732]

Our Sun server is down, so I can't produce the books that were due last Friday (deadline slipped from Wednesday). So I have to wait. Good thing I have something to occupy my time! . . .

This information-theoretic identification of information with a measure of uncertainty obfuscates the issues, I think. For example, still later in the same post (7/16/91 11:15), you wrote:

> . . . The overall uncertainty of the universe of possibilities can

> be partitioned [among many other ways as follows]: (1) the residual uncertainty after an event has happened, caused by difficulties in distinguishing events, (2) the information carried by the event, and (3) information about what subset of potentially possible events actually are likely to occur (e.g. in a universe of word strings, some strings are highly unlikely, whereas others are relatively highly likely). I have called (2) the information, and (3) the redundancy . . . You call (3) the information of the signal.
>
> . . . in my terminology [these] are (1) noise, (2) information, and (3) redundancy. The redundancy helps overcome the effects of noise, permitting the maximum transmission of information.

Here, you are talking about "(2) the information carried by the event" as though it meant information in the familiar sense ("bacon and two eggs over easy" or "I want a toy fire truck") rather than the numeric reduction in uncertainty n which prior context indicates you do indeed intend. What happens when we substitute "reduction of uncertainty" for "information"? Retaining your numeration, but in reverse order, this appears to me to be as follows:

> The reduction in the overall uncertainty of the universe of possibilities can be partitioned as follows:
>
> 3. Reduction in uncertainty as to which of the set of possible events {E} are likely actually to occur.
>
> 2. The reduction in this uncertainty when a particular event e <- {E} in fact occurs (where <- means "is a member of").
>
> 1. The residual uncertainty after e has occurred, due to difficulty identifying precisely which event in {E} has occurred.

I would call these something like (3) redundancy, (2) reduction in uncertainty as a consequence of information transmission, and (1) subsequent uncertainty as a consequence of ambiguity. To (1) I would add uncertainties due to inherent ill-definition of any {E} in nature--that is, due to the approximative character and "subsumption with neglect" of any category perception.

Note that "information" appears nowhere in this formulation. Nor are any linguistic utterances a necessary part of it, and consequently it most certainly does not specify my particular concern, which Harris has termed linguistic information.

[As an aside: I have said elsewhere that your (3) is related to redundancies that we call grammar (by which we determine operator-argument dependency classes and reductions), and not the information in particular discourses, which is in a sense built upon these redundancies--see Harris, _Language and Information_.]

In your recent statement (911030 14:10), you appear to be using "information" in a different sense:

> . . . A finds that the world is not as he wants it to be (perception does not match reference). There is something A can do to improve the situation ([A initiates] a goal and a plan to achieve it). This . . . might better be done with the

> participation of B, and A's model of B does not predict that B will do the necessary things as matters stand, so A must do something to affect B's future behaviour (a communicative goal emerges as a subgoal of the main goal). To achieve the communicative goal, A determines what information must be received by B. . . . that information is what I call "the top-level primal message".
>
> [Some caveats about ongoing calibration of an avowedly incomplete and inaccurate model.] . . . The "information" A wants B to get is an initial multidimensional representation of what A thinks will get B to act as A wishes, and it will be continuously modified as A observes B's reactions.

Now, try as I might, I cannot substitute "reduction in uncertainty" for "information" in the above passage (repeated below with a few amendments for concision en passant intended to be entirely friendly):

> A finds that a perception does not match a reference signal. A initiates a goal (a situation in which the discrepant perception would better match the reference signal), and a plan to achieve the goal. A recognizes that this might better be accomplished with the participation of B. In memory and imagination, A has a model of B. In A's model of B, B might participate as needed if B is aware of certain entities and relationships. A initiates a communicative goal (a subgoal of the main goal) by which B may become aware of these things.

So far so good. Now the sticky part:

> To achieve the communicative goal, A determines what reduction in uncertainty r must be received by B, where:
>
> r = n - m
>
> n = (sum p log p) over all the possible events before A's communicative actions
>
> m = (sum p log p) over all the possible events after A's communicative actions
>
> That reduction in uncertainty is what I call "the top-level primal message". The reduction in uncertainty A wants B to get is an initial multidimensional representation of what A thinks will get B to act as A wishes, and it will be continuously modified as A observes B's reactions.

But how can a reduction in uncertainty, the numerical value n, be a "multidimensional representation of what A thinks will get B to act as A wishes"? You see, the notion of uncertainty or reduction in uncertainty simply does not work as a concept of information. The fact that you put "information" in quotes suggests that you seem also to know that you need a different and more colloquial sense of the word, and that your information-theoretic notion of socalled information simply will not do. And that is because reduction in uncertainty is really only a measure of the *result* of successful information transmission, an indirect, pale, "skeletal" reflection of the information transmitted. It is a meter reading and not the gasoline in the tank, a number on a scoreboard and not the play of the game. Reduction in uncertainty cannot even be taken as a goal, in PCT terms, it is rather a particular arrangement of entities and relations (perceptions thereof) that is the goal.

Ambiguity (uncertainty) is a disturbance, and any control system will work to reduce disturbances, so in a sense reduction of uncertainty is an instance of the general meta-goal of control itself, which any control system has by virtue of being a control system, but it is not a goal in the technical PCT sense.

Now let me say again: I use the term "information" to refer to (perceptions of) entities and relationships among them. A speaker presents certain linguistic entities and relations in an utterance. In the perceptual hierarchy of the speaker, these correlate loosely with nonlinguistic entities and relations among them. The linguistic entities, restrictions on their combinability (operator and argument classes, reductions) that define their relations, and the loose correlation of these with nonlinguistic perceptions, are socially shared because of the way in which they are socially acquired and maintained. Because of this, the speaker's utterance constitutes an occasion for the hearer to consider (perceptions of) nonlinguistic entities and relations among them that she might not otherwise have considered. I hope that makes more sense this time around.

Language is more clearly organized than the world appears to be, or even than our nonlinguistic perceptions seem to be. What is the relationship between the socially constituted self-organizing system of language on the one hand and the perceptual hierarchy without language on the other? (Whatever that latter may mean--various of us have noted previously the enormous difficulty isolating language about perceptions from the perceptions themselves. I think for very good reason, as I have previously indicated.) Here are some options:

1. Language has its own structure, which maps loosely onto entities and relations in the perceptual hierarchy. There is hierarchical control of the perception and production of language, and there is hierarchical control of nonlinguistic perceptions, and there is control of the mapping from one to the other. There is no reason to suppose that this is a simple mapping of perceptions onto individual words in the lexicon, there is too much polysemy. (Consider for example the great diversity of perceptual associations with even a simple word like "up".)

2. The structure of language is precisely the perceptual hierarchy from the category level up. This, I take it, is at one end of a scale with respect to how closely integrated linguistic and nonlinguistic perceptions may be, with (1) at the other end.

3. There is no category level. We may talk about absolutely any perceptions, given attention and the development of agreed-upon vocabulary. When we talk about and compare intensities, sensations, configurations, transitions, events, and relationships, we categorize them in order to do so. Likewise when we talk about sequences, programs, principles, and system concepts. What appears to be the category level is simply the mapping from the perceptual hierarchy per se to the socially-inherited system of objects and relations that is language. On reflection, this appears to be an elaboration of (1).

There can be no social action without the perception of repetition, of equivalence, of a particular perception constituting an instance of a class or category. In short, category perceptions are the essence of social reality, as they are of course at the heart of language structure as well.

[Martin Taylor 911105 17:45]

>>The only thing we have for representing the information in natural
>>language is natural language itself.

>I do not see why a theory of language needs any more than a description
>of the circumstances in which the language is used and the circumstances
>in which it is used.

First, your description of circumstances is a description in--um--let me see . . . language? Secondly, language has a complex, socially constituted structure, and if you deny this or ignore this I suggest you get acquainted with it. I have posted references for descriptions of the structure of language (as opposed to descriptions of elaborate formal machinery for talking about language perversely confused with it). Suppose that you had a "description" of circumstances for language use couched alinguistically in memory and imagination perceptions, and perceptions of your current ongoing circumstances, you would need . . . no, I simply cannot make this anything like coherent.

>   There seems to me to be zero probability that a
>formal description can encompass the uses or the observed structure of
>language. I know there exist many theories of language that treat
>"language" as if it were an object having a finite description. That is
>not a view with which I can sympathize. To me, language is a means of
>affecting the knowledge or behaviour of another party. To some extent
>the user of the language takes advantage of conventionalized forms, but
>to a greater extent the communicative process is one of situation-specific
>negotiation. I cannot believe that natural language (and still less that
>formal language) can describe all these possible negotiations, and that
>is what would be required "to represent the information in natural
>language".

Language is a socially-extant system that we learn and that we use for transmission of information (perception of entities and relationships). We also use other means for transmission of information (Iago dropping Desdemona's handkerchief) and for communication (Cassio's tone of voice, gestures, posture as he says "I've lost my reputation!" in the same play). All of these things can be described in language, OF COURSE. We can talk about absolutely anything. Talking about it does not constitute (or substitute for) doing it, of course, even with respect to language, as we see when explaining a Yiddish joke (or a Hopi joke, or a joke in English, for that matter, if it needs explaining) fails to convey the effect.

The fact that language contains its own metalanguage entails a number of subtle and I think quite profound points. Some of them are of concern to linguists: there is no privileged position outside of language for describing language, any redundancy in the means of talking about language obscures what it is supposedly talking about, which is after all comprised of socially-established redundancies, etc. For those of us who are involved with PCT, it points to our difficulty talking about category perceptions and higher-level perceptions that depend upon category perceptions. There is no priviliged position outside of our socially-constituted system of category perceptions from which to describe these perceptions. Category perceptions (and sequence, program, principle, etc. perceptions) are of the SAME KIND as the objects and relations of language, and I believe they are acquired in the same way, socially, not by induction from cases. They are acquired at the same time as language, and by the same means, in one unified

process.

This entails the prediction that category perceptions may contradict and override lower-level perceptions. We may perceive things that are not evident because their category perception calls for them, and we may ignore perceptions that are present because they do not fit our category perceptions. I do not think I need to belabor this (it seems to me a truism), but I will go into one interesting case.

Kay and Kempton (_American Anthropologist_ 86:66, 1984) surveyed research on the Whorf-Sapir Hypothesis, particularly as related to color perception (research pioneered by Kay and Berliner in the late '60s). Last month, one of the authors posted the following interesting account to the linguist email list:

>Date: Wed, 23 Oct 91 17:16:06 EDT
>From: Willett Kempton <willett@Princeton.EDU>
>Subject: Whorf and color
>
>I'm a coauthor of the Kay and Kempton study discussed in several
>earlier messages. (I don't follow this newsgroup regularly, but a
>colleague passed on the thread.) As pointed out earlier, from the
>tangled cluster of hypotheses referred to as the Sapir-Whorf
>hypothesis, we tested only one question: Do the lexical categories of
>a language affect non-linguistic perceptions of its speakers to a
>non-trivial extent? (P. Kay & W. Kempton, "What is the Sapir-Whorf
>Hypothesis?", American Anthropologist, vol 86, No. 1, March 1984.)
>
>Considering the complexities of proir research efforts, our primary
>experiment was simple: Present three color chips (call them A, B, C)
>to speakers of two languages, such that colors A and B are slightly more
>different in terms of (universal) human visual discriminability, whereas
>B and C have a linguistic boundary separating them in one language
>(English) but not the other (Tarahumara, a Uto-Aztecan language).
>As noted earlier, the English speakers chose C as most different,
>whereas the Tarahumara chose A or split evenly (there were actually
>eight chips and four sets of relevant triads).
>
>I'll add a couple of points of interest that were either buried in
>that article, or have not appeared in print. First, as the speaker
>of a language subject to this perceptual effect, I would like to
>report that it is dramatic, even shocking. I administered the tests
>to informants in Chihuahua. I was so bewildered by their responses
>that I had trouble continuing the first few tests, and I had no idea
>whether or not they were answering randomly. In subsequent analysis
>it was clear that they were answering exactly as would be predicted by
>human visual discriminability, but quite unlike the English informants.
>
>An informal, and unreported, check of our results was more subjective:
>I showed some of the crucial triads to other English speakers, including
>some who had major commitments in print to not finding Whorfian effects
>for color (several of the latter type of informants were available on
>the Berkeley campus, where Kay and I were). All reported seeing the
>same effects. We tried various games with each other and ourselves like
>"We know English calls these two green and that one blue, but just
>looking it them, which one LOOKS most different?" No way, the blue one
>was REALLY a LOT more different. Again, the Tarahumara, lacking a
>lexical boundary among these colors, picked "correctly" with ease and in
>overwhelming numbers. The article includes the Munsell chip numbers, so
>anyone can look them up and try this on themselves.

>
>Some of the triads which crossed hue and brightness were truly
>unbelievable, as it was perceptually OBVIOUS to us Engligh speakers
>which one was the most different, yet all the visual discriminability
>data were against us. (The article did not mention the hue/brightness
>crossovers for the sake of simplifying the argument in print.)
>
>Our second experiment, like the original visual discrimination experiments,
>showed only two chips at a time. We additionally made it difficult to
>use the lexical categories. And we got visual discrimination-based
>results, even from English speakers. So there are ways to overcome
>our linguistic blinders. (Which we knew already, or the original
>visual discriminability work could not have been done in the first place.)
>I don't feel that the differences across these tasks was adequately
>explored, and represent a golden opportunity for a research project
>or thesis.
>
>I didn't expect to find this. The experiment was a minor piggy-back on
>another project. I believed the literature and the distinguished
>scientists who told me in advance that we wouldn't find anything
>interesting. The experiment was going to be dropped from the field
>research, saved by a converstion at a wine party with a "naive" sociologist
>(Paul Attewell) who had read Whorf but not the later refutations.
>
>A simple experiment, clear data, and seeing the Whorfian effect with our
>own eyes: It was a powerful conversion experience unlike anything I've
>experienced in my scientific career. Perhaps this all just goes to
>affirm Seguin's earlier quote, as applying to us as both natives and
>as theorists:
>
>"We have met the natives whose language filters the world--and they
>are us."
>
>- Willett Kempton
> willett@princeton.edu

Kay & Kempton show that configuration perception, after naming,
overrides sensation perception, prior to naming. Techniques to avoid
assigning name-categories to color chips being compared are described in
the article. I can summarize what I remember if you wish.

Bill Powers (911106.0800) --

>I'm having second thoughts about "language is perception", too. It seems
>to me, for example, that Harris's partial orderings can't be recognized
>without using a mode of perception that isn't expressed in words.

Harris describes the bootstrapping by which this can be developed both
onto-and phylogenetically. Do infants (by definition, children before
speech) have category perceptions? Yes. But do they not thereby also
have aspects of language? Passive command of language as is well known
precedes active production by a long way.

>I'm
>having doubts that language has any metalanguage at all: it's just
>language -- as Martin says, a process instead of a thing, or as I would
>say, signals handled by a process that is not represented in the signals.

To adhere to this position you will have to claim that control of language
does not require sequence and program perceptions referring to entities

and relationships of language.

>That which does language is not the same thing as the language that it
>does. Even to recognize the difference between two different word
>orderings requires sequence perception, which can't be expressed in words
>but can only be done.

To recognize the difference between two word orderings you must first recognize the words involved as words and refer to them as words being in this order or in that. That recognition and that reference as to order are metalinguistic perceptions. They are perceptions that are required for the hierarchical control of language. Are you saying that perceptions required for hierarchical control of language are not themselves part of language? Where are you placing the boundary between language and non-language, my friend?

>In other words, I think that a "metalanguage" is
>being asked to do a job that only a neural perceptual function can do.

Language and metalanguage are neural perceptual functions. Some neural perceptual functions are socially acquired and maintained, and these are among them.

This has been a very long post (over 500 lines). I am sorry if it is a burden. If it does not get through intact, let me know and I will break it up and re-send it. My bronchitis is on the mend. Sarah is at home resting, pheumonia at bay, but will be frail for 8-10 weeks, and that will naturally make life more complicated.

Bruce

==========

Date: Mon, 11 Nov 1991 12:14:43 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: Oded Maler: correspondence of perceptions to events

[From: Bruce Nevin (911111 1212)]

(Oded Maler 911104) --

>(Bill Powers 911104.0700, last paragraph):
>
>>The most important aspect of this analysis is the idea that these systems
>>are controlling for a subjective condition, not an objective one. They
>>are only trying to control their own perceptual signals.

>I find it reasonable, although the problem of *correspondence* between
>some perceived combinations of sensory signals and what happens to the
>organism in the "real" world, is sometimes under-emphasized. Even if
>the organism behaves according to simple perecption-based rules, in order
>to show that this behavior achieves something in some external physical
>environment, you can get involved in very complex models which are very
>hard (if not impossible) to analyze.

I don't understand this. There is no need for anyone to model the correspondence between perceptions and events. The feedback loop through the environment takes care of this, so long as the organism does whatever it takes to maintain the preceptions that matter to it at the reference values that it cares about. This quite apart from the fact that there is no way to know about events in the "real" world other than other perceptions, viz. of physicists' meter readings, so that the

correspondence to be modelled would be between the organism's perceptions and the investigator's quite disparate other perceptions--no wonder such "models" are complex and hard to analyze! Why should we engage in such pursuits?

Bruce Nevin
bn@bbn.com

==========

Date: Mon, 11 Nov 1991 16:45:52 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: on to 5 or 3 to 4

[from Joel Judd]

I've been waiting for a resolution to the level 3/4 debate, mainly to settle on an order in the dissertation. Since a reflective equilibrium (with apologies to Hugh Petrie) seems to be all I can hope for, I'd like to ask some questions that may have been answered in all the language polemic, but which I missed.

Is there phonological evidence, analgous to visual, that would argue for configuration arising from transitions? Is there something [inherent] in the phonemic sound train that would "require" a perceived transition prior to a "static" configuration? The development of these levels of perception happens so early that the phonological development I've looked at doesn't help at all. There was only one type of research that seemed to begin to approach this question. It was cited by Studdert-Kennedy (1986) in _Invariance and Variability in Speech Processes_, which I think I checked out for the Abbs article on disturbing mouth musculature during speech (Bruce brought the topic up several weeks ago)--I can never remember why I check books out by the time I get them. Anyway, Studdert-Kennedy mentions a study by Kuhl and Meltzhoff (1982; The bimodal perception of speech in infancy. _Science_ 218, 1138-1141). This was one of the "face preference"-type studies where infants (4-5 months) look longer at faces articulating the vowel ([i] or [a]) they're hearing than a face saying the other vowel simultaneously. They added pure tones matched for length and amplitude and found no effect. I haven't found a similar experiment for phonemes and tones. Overall, though, how would you determine one or the other perceptual ordering in infants anyway?

This leads to another question: Is it forcing the issue to insist (at this point) on one or the other level ordering anyway? A number of developmental linguists already emphasize the word level, even at the initial stages of language learning where no adult "words" are even discernable (eg. Charles Ferguson). Said Ferguson goes so far as to say the basic unit of phonological production is the "word shape," exemplified by utterances such as [kyu] for 'thank you.' How the child comes to use even somthing like [kyu] is still unanswered, but perhaps the initial sound development is so early that for many purposes forcing the issue now is not necessary (?)

Locke (1983; _Phonological Acquisition and Change_) offers a fairly simple explanation of early phonology (pp.83-4):
  It appears that as a child reaches that point in his social and cognitive life in which it is desirable and possible to designate things with strings of sound [then he can utilize sounds he already produces as a result of babbling]...in projecting his available articulations the infant will produce a number of "hits." Many of his [d]s and [b]s will land on lexically standard /d/s and /b/s. A number will land elsewhere...because the child (a) may not notice that those sounds differ from [d] and [b]; (b)

may not know that the difference matters, in some sense; or (c) could not make the necessary articulatory [even if he recognized a difference].

The interesting question, of course, is How does the child know what to 'shoot for' in order that he might make a 'hit'? In saying that we might not have the answers for a while, he gives an interesting quotation from Peters (1974), which I give in part:

> If a child is aiming at a constant target then why do we observe that her pronunciation of a specific word may improve right after an adult has pronounced the word--and then revert to her current level of accuracy [p.97]

The same thing happens in [pronunciation] language drills with older L2 learners.

To avoid ending with a question I'll mention the Abbs (1986) paper also in the _Invariance..._ book. It is the same author who is with Gracco in the paper Bruce mentioned, where Ss mouths are disturbed as they speak. This article seems a little more general and Abbs mentions some other related literature others might be interested in, particularly arm movement interruption. Some of the concluding comments are worth reproducing:

> What is apparent from the study of speechas a motor behavior is that the phonological goals represent the *intent* of the system to generate an acoustic output that is perceptually acceptable. This concept of motor-speech intent is not different, in principle, from the abstract motor objectives underlying handwriting, sign language, or perhaps even nonlinguistic motor behaviors such as the successful parallel parking of an automobile or reaching for an object in three-dimensional space. In none of these cases does the goal prescribe the specific muscle actions or isolated movements involved in its successful achievement. On this basis, it is of limited value to describe isolated movements or muscle actions of teh speech-production system without parallel observation of synergistic actions that may covary in achieving overall system goals. (pp218-9)

==========

Date: Sat, 9 Nov 1991 15:12:34 +0100
From: Oded Maler <Oded.Maler@IRISA.FR>
Subject: Re: Test; config control; linear transform

[From Oded Maler 911109]

Just a minor terminological note.

I wrote

>>This [range of control] might be meaningful when we move up to multi-
>>dimensional systems, because some trajectories for correcting a
>>"virtual" signal may go outside the controllability region of the
>>elementary inputs.

And Bill replied with

>A position control system will head directly toward the reference
>position without regard to intervening obstacles. The trajectory is
>whatever results from the way muscles are used on the way to the goal --
>it isn't controlled at the configuration level. Normally, however, the
>actual configuration equals the reference configuration, so trajectory
>planning is really planning how to change the position reference signals.
>I think that "planning" is the wrong word to use at these low levels,
>however (an aside to Martin -- we can get into that later). Outputs are
>not planned -- not if you want to deal with disturbances.

>I think the problem about trajectories is related to the reason for
>development of higher levels.

But I didn't mean "trajectory" in the sense of pointing behavior, but rather in describing the trajectory between 2 points in the sensation space, which is passed when you try to control a configuration.

--Oded

==========

Date: Mon, 11 Nov 1991 15:48:53 PST
From: marken@AERO.ORG
Subject: The Blind Men and the Elephant

[From Rick Marken (911111)]

The blind men and the elephant: Three different views of the phenomenon of control.

The output of a control system depends on the reference input and the disturbance to the controlled variable as follows:

o = Ke*r/(1+Ke) - Ke*g(d)/(1+Ke)

where o is the output variable, Ke is the output amplification factor, d is the disturbance and g() is the function which transforms the disturbance into output. With sufficient amplification (Ke >> 1) this equation simplifies to

o = r - g(d)

I believe that this property of control (output directly proportional to reference input and negatively related to disturbance) is the basis for the three major ways of looking at "behavior" in scientific psychology.

1) stimulus - response

Behavior seems to be a response to stimulation when r is relatively constant and all disturbances to the controlled variable except d are held constant (controlled). Then what you are seeing is
o = -g(d)

and it looks like stimuli (d) causes reponses (o). The nature of the relationship between o and d (g) will appear to be the "psychological law" relating stimulus and response. This relationship will be "noisy" due to variations in r (which cannot be controlled) and variations in other disturbances to the controlled variable.

2) reinforcement

Behavior seems to be selected by its consequences. This happens when r is constant AND d is constant (call it D). All that is changed is the relationship ,g, between o and d. For example, d could be a certain amount of food deprivation. The organism has a fixed reference for the amount of food (say 10 units). The deprivation causes the input food to be perceived at 0 units. g determines how much output must be generated to compensate for this deprivation to bring the food input to 10 units.

g determines how much output must be generated to make up for x amount of deprivation. If g is demanding, then the animal responds more per unit deprivation than it would if g were lenient. This is the schedules of reinforcement and, probably, matching findings.

o = g(D)

3) planned output -- cognitive behavior

Behavior seems like planned output if you remove the disturbance (set g(d) to zero) and watch output "generated" by the subject. In this case you are seeing

o = r

An example of this might be the protocol analysis of people solving puzzles (there are disturbances here but much of the subject's activity seems "spontaneous" -- this is the behavior of the reference signal).

This little exercise is the beginning of a proposed essay (with the title given above) in which the phenomenon of control is shown to be analogous to the famous elephant described by the blind men. Each man sees the elephant (control) from a slightly different perspective. And each concludes that it is a different thing -- s-r, reinforcement or cognitive behavior.

This is obviously just a start. It needs much more detail. But any comments, examples or flames would be greatly appreciated. I'll be out of town 'til the end of the week so I thought I'd leave you with something to chew on.

Best regards

Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Tue, 12 Nov 1991 09:02:09 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: 5, 4, 3

[From: Bruce Nevin (911112 0755)]

(Joel Judd Mon, 11 Nov 1991 16:45:52 -0600) --

>Is there phonological evidence, analgous to visual, that would argue for
>configuration arising from transitions? Is there something [inherent] in
>the phonemic sound train that would "require" a perceived transition prior
>to a "static" configuration? The development of these levels of perception

Martin suggested that both orderings might exist in different parts of the system. As I recall, foveal vision seemed to be one way, and other

vision the other way. (Martin?)

Again, Martin is more expert than I in the details of phonetic analysis. However, I will venture some speculation out of my reading and thinking about this.

It appears that both acoustic and articulatory goals are involved in speech, with acoustic goals generally predominating. It further appears that articulatory goals are more important for consonant articulation than for vowel pronunciation.

Acoustically, consonants are distinguished by transients in adjacent vowels, by bursts, and by other transient phenomena. Articulatorily, they are characterized by steady-state "target" positions of articulators that are relatively easy to specify. They are also relatively easy to control for, since there are tactile as well as kinesthetic cues as to shape, orientation, and location of articulators.

Articulatorily, vowels are characterized by "coarticulation effects" of transition between adjacent consonants. Articulatory "target positions" are difficult to specify (they turn out not to map onto the traditional "vowel triangle" at all well, see Lieberman and Blumstein). Articulatory positions of vowels are difficult to control for, as there are only kinesthetic cues, and coarticulation effects can be overwhelmingly strong. Acoustically, vowels are characterized by steady-state "target" relationships of formants in contrast with those of other vowels in equivalent phonemic environments.

From these observations comes a guess that consonant pronunciation is governed by articulatory goals, while vowel pronunciation is governed by acoustic goals. A language with vowel harmony (e.g. Turkish) suggests a control system for vowels operating across words regardless of consonants; conversely for a language with consonant harmony (e.g. Kurdish, I think, some of the Salishan languages). Similarly for assimilatory changes in historical linguistics, slips of the tongue, etc. Such a view dovetails nicely with autosegmental phonological theory: each tier of autosegments (not just the C-V tier) would have its own control systems for the contrasts maintained on that tier.

There is lots of evidence to be mined from the literature amenable to reinterpretation in PCT terms. I recommend particularly David Stampe's "natural phonology". Also Vicky Fromkin's work on slips of the tongue some years ago and the experimental work I cited a while back. Refs if you want them, but I expect you know the literature better than I do.

Your question seems to presuppose that the developmental order of levels of control follows the order of internal organization of the control hierarchy. I would suggest that some bootstrapping is possible, perhaps even necessary, and certainly plausible given the evidence for social matrices (Bruner's "formats") in which children learn social action. Thus, a child may acquire a higher level of control with structured contextual support from parents and other adults before being able to control lower levels well at all, making do with approximate and substitute outputs and very low gain, so that the socially propped-up higher level of control both directs and instigates (i.e. provides goals for) acquisition and refinement of lower-level control that the former is logically "built upon".

Similar but I think different is Ferguson's observation about word-level control. Language seems to have a hierarchical structure: features,

(auto-)segments, syllable constituents, syllables, metrical units, phonological words, breath groups; morphemes, syntactic words, hierarchies of morpheme- and word-classes, syntactic constituents, etc. I think that control of different aspects of language is parallel, not hierarchical. This is obviously the case around the break between phonological units and syntactic/semantic units marked by semicolon in the above list, less obviously but I think necessarily the case elsewhere in this organizational or structural (but not control) hierarchy.

However, relationships in this structural hierarchy can guide our thinking and research about attention and the relative tightness of control (gain) in different parts of the control of language. An illustration:

>Locke (1983; _Phonological Acquisition and Change_) offers a fairly simple
. . .
>The interesting question, of course, is How does the child know what to
>'shoot for' in order that he might make a 'hit'? In saying that we might
>not have the answers for a while, he gives an interesting quotation from
>Peters (1974), which I give in part:
> If a child is aiming at a constant target then why do we observe that
>her pronunciation of a specific word may improve right after an adult has
>pronounced the word--and then revert to her current level of accuracy
>[p.97]

I think we must integrate this with Bruner's work on socially constructed "formats" supporting the child's acquisition of language, and with the suggestions just above about parallel control. When the child attends closely to pronunciation in a context that provides social support, gain is higher on the control systems for pronouncing the utterance singled out by adults for attention. Immediately after, however, the child's attention reverts to larger questions of accomplishing communicative goals ("doing things with words"), to which the detailed control of pronunciation at an adult level is greatly subordinate.

Hope this is helpful.

Bruce Nevin
bn@bbn.com

==========

Date: Tue, 12 Nov 1991 14:51:40 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: on to 5 or 3 to 4

[Martin Taylor 911112 14:15]
(Bruce Nevin 911110 1747)

Thanks for the long post about the word "information". I agree, it is polysemic. Your posting helps in the clarification, but I fear that I will be muddying the waters more in the future, because I believe that the Layered Protocol approach helps to relate the meanings of "information" as amount and "information" as content.

I don't remember the context of the prior discussion, but I think the disagreement was in your identification of the information content of a sentence with its structural relations. In that case you are playing in the "amount" field, and my comments about information being the reduction of uncertainty apply. The structure of a message is its redundancy, which does not go to reducing the uncertainty about the

content achieved by the recipient on receipt of the message. Hence the structure of the message is not normally related to the information conveyed by the message. I agree that the amount of information conveyed by the message is not the same as the content, which is in colloquial language (which I have used often enough) the "information" in the message.

The above should be slightly modified, in that the choice of structure used by the originator of the message may have been uncertain to the recipient, and in some cases that uncertainty can be used to convey the main content, the surface content being of less informational importance (i.e. almost anything could have been substituted, provided the same structure was used; in a convivial social group, the use of a very formal register may signify something, and the verbal content may just be a carrier for the message contained in the structure--but this is not the normal case).

I confess to a loose use of many words, when I am not concentrating on the usages of those particular words. "Information" is one such.

Martin

==========

Date: Tue, 12 Nov 1991 15:12:09 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: 5, 4, 3

[Martin Taylor 911112 15:00]
(Bruce Nevin 911112 0755 in response to Joel Judd)

I endorse Bruce's comments about phonology. I was going to answer Joel directly, but I couldn't find the reference; for what it's worth, here's what I remember about acoustic transition detection.

Some experiments around 20 years ago studied the ability to discriminate acoustic events (easily discriminated from each other, such as high-pip low-pip or loud-click soft-click) that differed in the order of presentation (loud-soft vs soft-loud, for instance).

The major finding that I remember is that when the two events are close in time (<20 msec?) or far apart in time (>70 msec?) the two orders are easy to distinguish, but at intermediate time separations they are not. What this says to me is that a fast transition is a single phenomenon that can be incorporated as part of a configuration, whereas the slow transition is a sequence of configurations. In between, the two events are too far apart to be treated as one phenomenon, and too close to be elements of a sequence. (I use the word "phenomenon" because I don't know which CT level I should talk about here.)

As for the fovea, all I can ever remember saying was that the visual periphery seems to act as if its basic task was the detection of movement, which it does exquisitely well though it is terrible at shape, whereas the fovea is good at static shape and somewhat poorer at movement.

Martin Taylor

PS. Bruce: Is there any information about the relative precision of target location for vowels and consonants in different languages? In my casual observation of Danish, I have the impression that they are much less concerned than we are with precision of consonants but much more concerned than we are with precision of vowels. Have there been any

phonological studies in this area?

Martin

==========

Date: Wed, 13 Nov 1991 00:39:11 CDT
Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: Re: Blind men and elephant

From Tom Bourbon [911113 0:27]
Rick Marken [911111] Your interpretation of three classical theories of behavior as special cases of the functioning of a PCT control system squares with the interpretations of SR and plan-driven models Bill Powers and I develop in our manuscript ("Models and their worlds") that is nearing re-submission. The reason those earlier models developed was that, when viewed under certain constraining conditions, a control system can appear to be an entirely different kind of system.

There is more: a person can decide to act like one of the other systems -- attempting to make actions track or follow the pattern of some external event (that thereby assumes the role of "stimulus" to the person's "response"), or to make actions follow some previously decided-upon "plan" (schema, cognitive program, etc), in which case the person will stumble and fumble in a manner similar to that of a genuine program-driven system. And the PCT model can emulate the actions of SR and plan-driven systems as well as can we. That is the topic of the sequel to "Models and their worlds"-- which is in its early stages of preparation.

Wayne Hershberger has also touched on some of these ideas, in his letters published in _American Psychologist_. We seem to be converging on another gambit, in our attempts to speak to traditional behavioral scientists.

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402

==========

Date: Wed, 13 Nov 1991 06:33:34 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: Information; continuity; transitions; trajectories; elephants

.LM0/.RM79/[From Bill Powers (911112.0700)]

So much to comment on -- probably won't finish in time to send this today.

Bruce Nevin (911112) --

My sympathy to Sarah, with an appropriate fraction of it to you. Being sick messes up life, doesn't it? I know what you're going through, somewhat, even though I didn't have to keep a job going while taking care of Mary.

I'd like to horn in a little on the debate concerning information -- next June I have to give a talk on this subject to a Systems Science meeting in Denver, so a little warming up will help. I'll also get into other subjects in your post.

Originally, "information" was defined by Shannon and Weaver in an attempt to reach a definition of channel capacity. It had nothing to do with semantics: it was purely an engineering measure. It still is an engineering measure, like signal-to-noise ratio but most appropriate for digital communication. It says absolutely nothing about what is in a message. All it gives us is another measure of bandwidth. The connection with energy and entropy is metaphorical; energy can flow in the direction opposite to "information." The battery that supplies the telegraph line could be at the receiving end. Neural impulses don't carry en
ergy in the direction of propagation; they dissipate energy. So be careful about what you say about information "reducing entropy!"

The "reduction of uncertainty" definition has been extended to the content of messages. This definition of uncertainty depends on thinking in terms of a finite or enumerable set of possible meanings, shared between transmitter and receiver, about which one can be uncertain. This is not what the Shannon-Weaver concept was about.

Note that the question of channel capacity, which entails a conderation of random noise that destroys "bits", implies *variability*, which is not the new meaning given to "uncertainty." The new meaning (as of the 1950s) assumes that there is a finite universe of m possible choices, and that the message pins down exactly n of the dimensions, leaving an ambiguity -- a choice among alternatives -- exactly equal to the number m-n of choices not pinned down. So there is really no noise involved in this kind of uncertainty. I say "Buy the Honda," and you reply, uncertainly, "The green one or the red one?" It never occurs to you that I might have said "Buy the Hyundai." That is not the kind of uncertainty now meant
. It's assumed, in fact, that the communication channel is perfectly noise-free, so that if the same message is sent 100 times, the same uncertainty will result 100 times.

Quantitative measures on a continuum, such as are represented by the position of a meter needle on a scale or the time-interval between two neural impulses, can't be divided into convenient discrete packages. The basic "uncertainty" in such measures isn't visible until you get to the quantum level. With continuous measures, the more appropriate concept is signal-to-noise ratio. This concerns the repeatability of messages, not a choice among alternative meanings.

To speak of choices among alternatives means that one is thinking in terms of discrete categories from the start. This is a function of language. Language, by assigning labels to categories, reduces the universe to a set of discrete elements. It thus creates a special set of perceptual entities that have properties unlike the properties of other perceptual entities. The universe of which one speaks is therefore not the whole universe that is experienced. It is only a sampling that leaves out the continuum of change between the samples. The choices presented to us in direct experience differ only infinitesimally from each other; in fact they lie on a continuum. The artefact that is language creates an artificial world in which symbols with discrete values substitute for perceptions with continuously-variable values.

Suppose I ask you, "Will you bring me the rest of the standiformers?" You, not knowing the verbal category "standiformer," reply "Bring you WHAT?" I pick up my 13/16-centimeter standiformer, and say "The rest of THESE." Oh, you say, and you go get the rest of the things in the toolbox that "look like" the standiformer. That is a nonverbal category.

What is the basis on which you picked out the standformers from the

screwdrivers and wrenches? It can only be that in some sense each object in the toolbox looked more or less "like" the standiformer in my hand (sorry about misspelling standiformer in the first sentence of this paragraph -- or was that the only correct spelling?). This capacity for perceiving "likeness" is part of the capacity for forming categories of perception. Some objects looked "very like" the standiformer; some looked only a little like it; some didn't look like it at all. If we're to understand categories, we have to try to guess what sort of perceptual processing is needed to form this variable -- continuously variable -- sense of categoriness.

We do this sort of thing all the time. We pick up a screwdriver even though screwdrivers vary all over the place in size and configuration (think of one of those gadgets bent at right angles at both ends, with one blade set ar right angles to the direction of the other). We treat items that are actually different as if they were equivalent. Language isn't necessary for doing this. But this ability is necessary for language.

So the continuum that supplies the meanings of messages is not organized in such a way that we can calculate the information in the messages on the basis of a finite set of possible meanings. The set of meanings is infinite. It seems finite only when we use the words themselves, which are discrete, in place of their meanings. But the words themselves, treated merely as symbols, have no more meaning that "x" does in algebra.

Segue to another topic:

Ordinary language is about things we can discretize without inconvenience. But not everything can be treated that way. We have to handle continuous variables, too. In order to do this with language, we have to approximate. A meter reading shows us a needle in one particular place on a scale, within the limits of resolution of the eye. That appearance of the needle is the perceived meter reading. But to speak of this reading or use it in formal manipulations, we have to try to represent it in language somehow. We write down "3.14 volts." Looking a little closer, we write "3.142 plus or minus 0.001 volts". Referring to theory, we see that this voltage *should* be pi volts. If it isn't, we've misread the meter a little, or the meter is a little off calibration.

So we can conceive of quantities, and represent them with symbols, that have an exact place on the real number scale. We can conceive of variable quantities -- y = sin(omega*t) -- that live in a continuous universe and do not jump from one state to another. This sort of conception is just as much a part of our program-level perceptions as is a digital operation like "if (n > 2) goto 20" or "If you aren't out of there in ten seconds I'm going to break the door down." The outfielder backpedaling to keep the "ball" doing something called "rising" at a rate labelled "constant" is performing a continuous quantitative calculation that is adjusting his spatial relationship to the ball, as computed from continuous retinal and kinesthetic information. The signals involved in this calculation do not jump from one frequency to another: they vary as smoothly as the interval between successive nerve impulses can vary.

The machinery behind activities like these can compute in many ways. One of the ways it can compute we experience as language. Another is logic. Another is continuous analogue computation. I can't make a machine that can speak if you just hand me a dictionary or a grammar. You have to hand me a soldering iron. This is why I no longer think that language is bootstrapped by a metalanguage. Language and metalanguage are just two of

the things that this machinery can do. It's easy to confuse the program with the computer. But to understand what we observe, we have to try to understand the computer, not just one example of what it can do.

Joel Judd (911111) --

Another interpretation of the Whorfian hypothesis is that BOTH perception AND language reorganize so as to become consistent with each other.

No matter how you speak of a chest of drawers, both language and perception have to recognize that the drawers can be pulled out. A language that implied that the drawers were an intrinsic part of the chest would have to be changed. A perception of the whole chest as a unit would have to reorganize as soon as you saw a drawer being pulled out. What you are seeing is not a chestofdrawers.

Third and fourth level:

We have to distinguish the internal operation of perceptual systems from the signals they produce that represent reality to us. Even though microsaccades are necessary to maintain vision, the continuing jumps aren't represented in perception.

So the transition level, which does produce signals we recognize as "motion" or "change", goes back above the configuration level. A steady change-signal arises from the second hand of a clock, but not from the minute hand or the hour hand. The second hand can still be seen (perhaps because of microsaccades) if the clock stops, but the transition signal disappears. When the power has been off, a digital clock says "12:00" with a superimposed blinking on and off: the on-transitions and the off-transitions are something added to the perception of "12:00", which can also appear steadily twice a day when the clock is running. The clock does not have to blink in order for us to read it; but it can't blink if there is never any number showing.

The principle is: if you must vary X in order to affect Y, and if Y cannot exist if X does not exist, then Y is of a higher level than X.

So reorganization can bring us back where we were, too.

Oded Maler (911111) --

>... I didn't mean "trajectory" in the sense of pointing behavior, but
>rather in describing the trajectory between 2 points in the sensation
>space, which is passed when you try to control a configuration.

In passing between two points in sensation-space, the perception follows a trajectory. It must, because this is a continuous space and sources of stimulation can change only continuously. The question is whether that trajectory is the object of control, or whether it is simply due to the fact that in order to get from x,y to x',y' some set of intervening points must be traversed.

Suppose that the change occurs only in x. Initially, the perception x is near the reference level x*, set by a higher system. Now the value of x* steps suddenly to a new value, x*'. The error signal, x*' - x, jumps to a large value, sending a large signal to lower systems that activate muscles. The muscles tense, affecting whatever sensation is under control, and x begins to change toward x*'. If the system is well-damped, the approach will be exponential: as the error decreases the velocity

will decrease, with an asymptotic approach to the new condition x = x*'.

If this happens simultaneously in two axes, y also will approach y*' exponentially. If the exponents are the same in the two systems, the trajectory will be a straight line from x,y to x',y', traversed at a decreasing speed. If the loop gains and damping of the two control systems are different, or nonlinearities are different, the approach will be along a curve. The shape of the trajectory, however, is not under control: it results from the details of gain and damping in the two control systems.

Now consider what happens if the reference position x*,y* changes slowly in comparison with the speed of control in the slowest of the two systems. As x* and y* begin to change, x and y change too. There is still a curved trajectory, but now the positions lag only slightly behind the reference positions: the curvature is negligible. Viewed from afar, the points x and y move right along with the changes in x* and y*.

Now by manipulating x* and y*, a higher system can cause the point to describe any desired trajectory in sensation-space. The control of trajectory is occuring in the higher system. If x* varies as a sine function of time and y* varies as a cosine function, the point in sensation space will move at constant speed along an arc of a circle.

The same consideration holds at the next level: is the trajectory produced at the sensation level by changes in x* and y* purposeful, or is it merely a byproduct of the gain, damping and nonlinearities of the higher system? In the case of configuration control, I would guess that it is a byproduct. The reason is that a static configuration reference signal specifies a static condition, a fixed arrangement of sensations, like a posture. An unchanging set of sensations goes with an unchanging configuration. So the manner in which a configuration gets to its reference state is determined by the design of the control systems, and is not purposeful.

Only when we get to the transition level, where a constant reference condition implies continuing change, can we call a trajectory purposive. Now the control system senses a dynamic function of lower-level signals. If the output of this dynamic function is to match a reference signal, the lower-level perceptions must be changing as the inverse of the dynamic function.

One of the startling features of control systems is that they can achieve such inverses easily, without computing them. If the input function is computing the time-integral of a lower-level variable, the error signal need only be amplified and used to specify the amplitude of the lower-order variable. That variable will automatically change as the first derivative of the rate reference signal. This is the basis for the way operational amplifiers are used in analogue computing. There is a strange bootstrapping effect, in that the feedback provides just the variation in input that makes the time function of it match the reference signal. There is NO inverse time-function in the output part of the control system. The inverse is created completely by the feedback effects. It would be well worth anyone's time to study some texts on analogue computing and understand how this phenomenon works. It is essential in understanding how PCT works, of you want to go beyond qualitative understanding.

So I believe that trajectories first come under control at the reinstated fourth level, the transition level, where dynamic functions of inputs are

perceived and made to match reference signals.

Rick Marken (911110) --

The blind men and the elephant:

Another beautifully simple parable from Marken. Publish if they'll let you. This will bring together all our complicated arguments about the relationship of control theory to others into one simple framework that nobody in his or her right mind could reject. Now all you have to do is find an editor and a set of reviewers in their right minds. I wouldn't recommend expanding it too much. Simple and concise.

Best to all

Bill P.

==========

Date: Wed, 13 Nov 1991 07:41:56 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>
Subject: Levels

[From Bill Powers (911113.0745)]

Bruce Nevin (911112) --

One thing to keep in mind about the relationship between phonology and articulation is that articulation is NOT at a lower level than phonology. It must involve the same hierarchical levels as control of phonology. That is because articulation is not an output process but a process of controlling (kinesthetic) perceptions. Muscles are used to create a sense of effort sensations, articulator configurations, and transition from one configuration to another. They are also used to control (auditory) sensations, configurations, and transitions.

It is likely that some phonic configurations are composed of sound and kinesthetic sensations: saying "fffff" or "zzzzz" has both a feel and a sound to it. Saying "zzzzzip!" ends with the sound of the lips pressing abruptly together, which is partly a feel. When you say and act out "Pro - nounce eet cleeer - ly" sound and feeling form one set of configurations in transition.

Martin Taylor (911112b) --

>The major finding that I remember is that when the two events are close
>in time (<20 msec?) or far apart in time (>70 msec?) the two orders are
>easy to distinguish, but at intermediate time separations they are not.
>What this says to me is that a fast transition is a single phenomenon
>that can be incorporated as part of a configuration, whereas the slow
>transition is a sequence of configurations.

Rick Marken should joint this conversation and describe again his experiment with recognition of perceptions of different levels.

In dealing with temporal phenomena we have to tread carefully, because (I think) there are different *aspects* of temporal phenomena to which we attend at different levels.

Presenting two different sounds very close together in time creates one cycle of an oscillation. If you hear those two sounds in *continuous*

rapid alternation, at something like 5 or 10 alternations per second, you hear two things: both sounds simultaneously (a configuration) and a very low-pitched sound (reperception of the envelope as a low-frequency tone -- a sensation). A shareware program called Pianoman takes advantage of this two-level perception to produce four-voice music out of a PC loudspeaker. It sounds like bubbly piano music, with chords clearly audible, but with a disconcerting undertone at the frequency of switching.

At a somewhat longer interval between sounds, you hear an event: that is, you can distinguish one pattern of changes from a different one, but you can't tell which element is out of order when the pattern fails to repeat exactly. Something like this happens in fast typing -- you pause, knowing that your fingers have just produced a misspelling, but you have to look back at the word or remember the feel of the keystrokes to see which letter was mishit.

At still longer intervals, you have time to categorize the separate events and perceive the ordering directly, in detail. Now if I say "A, B, C, E, D ..." the incorrect element of the sequence leaps out the instant it occurs. This is sequence perception.

I suppose there's even a higher-level example: "Go through the door and then open the door" creates a *logical" problem with the meaning of the sequence, even though you heard the sequence as spoken. Logic tells us that if you're to to go through the door, then the door must be open already.

And of course, way back down at the bottom, it seems that there must be changes of intensity in order for at least some sensations to persist in time (or maybe it's changes of sensation for configurations to persist in time). That kind of change is not perceptually represented; it's part of the mechanics of perception. Another example of this is edge enhancement in vision -- we don't see objects outlined, so edge enhancement is part of the internal organization of a perceptual function, not appearing in ordinary perception where it can be controlled. Edge enhancement may simply compensate for a rounding of edges that occurs at a higher level. I wouldn't be surprised to find a phenomenon similar to Mach bands in hearing changing sound-intensities. Mach bands are the penalty you pay for perceiving sharp edges as sharp. Ah, epistemology.

Best

Bill P.

==========

Date: Wed, 13 Nov 1991 09:09:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: epistemology

[From Wayne Hershberger]

Bill Powers (911110.0900)
>Square one it is. I agree that the truly inaccessible can't be
>accessed. That's a definition, not a proposition. The truly
>inaccessible is that to which we have no access at all, either
>direct or indirect.

>The second part of the statement is not a definition: it asserts
>that the accessible sometimes appears inaccessible. I think this
>is an attempt to create a two-valued situation (either something

>is accessible or it is not) out of one that has more than two
>possible values.

>I see accessibility of the workings of nature to be a matter of
>degree, with the maximum degree falling short of 100 percent.

While agreeing that the truly inaccessible is inaccessible, you assert that much of that which is accessible is not all that accessible. Here you are changing the subject (i.e., I'm assuming that you are not simply contradicting yourself). That is, you are using the word accessible to refer not to the possibility of epistemic access but to the amount of X that is accessible or to the difficulty of achieving access to X, which could vary, of course, with the directness or complexity of the epistemic process. This is as misleading as referring to the length of a pregnancy as a degree of pregnancy (e.g., the unwed mother who claims she is "just a little pregnant"). More specifically, if 90% of X is truly accessible and 10% is truly inaccessible, it does not follow that X (i.e., every bit of it) is 90% accessible.

>Suppose we have before us a black box which we have no means of
>opening. Let's call this box, to humor me, the Reality Box. On
>its surface are numerous buttons and lights. The buttons and the
>lights are undoubtedly real, because we experience them directly
>and unequivocally: they are totally accessible, and cannot be
>mistaken as being inaccessible to our observation. So the box,
>the buttons, and the lights are not an epistemological problem.

On the contrary, the epistemological problem is usually stated in just such terms; for instance, it might be said that although the buttons "appear" to be solid and stationary they "really" comprise a swarm of whirling dervishes known as atoms. Or alternatively, it might be said, as you are wont to do, that they really comprise a collection of neural signals. In fact, the buttons, lights, etc. are the phenomenal objects that the empirical process we call perception provides us, and the epistemological question concerns whether or not these phenomenal objects are as objective as the label "object" implies. The box, the buttons, and the lights pose the epistemological problem! Your elaborate example begs the question. Your recent comments about objectivity were, as I recall, more to the point.

To suppose that the phenomenal objects are accessible appearances comprising indirect representations of an inaccessible reality is to embrace a radical skepticism, because there is no way to assess the fidelity of the representation without accessing the inaccessible (i.e., in order to test the correspondence between the reality and the appearance, one needs access to both and that, by definition, is not possible). This question of correspondence between what is accessible and what is inaccessible (i.e., between what is internal to and what is external to the limits of experience) is readily confused with the correspondence between what is internal to and what is external to the nervous system. But whereas the former type of correspondence is impossible (by definition), the latter type of correspondence is easily determined--neurophysiologists, like David Hubel, do it all the time.

>What this upsetting stranger is doing is reminding everyone that
>all they can actually observe are the lights, and the only

>effect they can know they are having is to press the buttons.
>All the rest is imaginary. Therefore we should throw away all
>those figments of the imagination, and admit that all we know is
>how the buttons affect the lights, and to remain pure of heart
>we should talk about nothing else.

If I get your meaning, that stranger is no stranger. His name was Bishop George Berkeley. Later, Johannes Muller's echoed Berkeley's Subjective Idealism in Muller's doctrine of specific nerve energies; as Muller put it (in article 5), the sensorium is aware not of the external object but of the state of the nerves only. Having said that, Muller then seemed to recognize belatedly that Berkeley's thesis makes no sense expressed in physiological terms because in article 8 he says that the sensorium is aware not merely of the state of the nerves but of the external causes as well. (There's nothing like having your cake and eating it too.) Muller was confusing the two correspondence questions described above. Berkeley's philosophy concerns only the former type, as Kant observed; that is, the only objects we experience are phenomenal not noumenal things.

I share you concern with the essential nature of nature. That is, when an experiment asks a question of nature, "someone" answers. But I see no reason to exclude this final arbiter of empirical truth from the phenomenal domain. Banishing this arbiter to an inaccessible realm from which it creates accessible appearances (like the Wizard of Ozz) makes about as much sense to me as claiming that today is but a representation of yesterday's REAL TOMORROW. What does it buy one, but a big headache? Who needs it? What's wrong with immanent truth, as reflectd in phonomenal coherence?

>Best re-gardes,

I am not interested in a fencing match. I AM interested in your staking a claim to what I see to be the epistemological high ground (in my view, solipsism is not the high ground). My motivation is selfish. Because you are the principal champion of psychomodular control theory I have a vested interest in your being in the best position to defend yourself, both your psychomodular theory and your epistemology. And since I do not see your psychomodular control theory as implying a solipsistic epistemology I see no reason for you to defend that indefensible epistemological position. My inability to persuade you to give solipsism a wider berth than you do leaves me ambivalent about my efforts to that effect, because, as I say, I am not interested in a fencing match. I am not interested in being a mere disturbance.

Warm regards, Wayne

Wayne A. Hershberger
Professor of Psychology
Department of Psychology
Northern Illinois University
DeKalb IL 60115

Work: (815) 753-7097
Home: (815) 758-3747
Bitnet: tj0wah1@niu

Date: Wed, 13 Nov 1991 15:45:40 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: phonologies and metaphors

[from Joel Judd]

Thanks for the comments on phonology and ordering. (sideline: got to hear C. Ferguson in an informal lunch discussion today about Language and Religion. Lots of room for PCT there, judging from what I heard). Bill's definitive statement regarding level 3/4 ordering seemed quite clear now. The principle in phonetic terms(?): If you must vary a point of articulation in order to affect the phoneme, and if the phoneme cannot exist if the point of articulation does not exist, then the phoneme is of a higher level than a point of articulation.

If anyone's interested, I'm continually trying to flesh out an analogy between bike riding and language I started using a few years back. Here's what I just sent to an SLA theorist for our discussion about theory types, specifically "transition" and "property" theories, and whether both are needed to account for language acquisition (his view) or if one theory could be sufficient (my view). Continued suggestions for details regarding any of the issues discussed on the net (eg. statistical measures and understanding of single systems) would be welcome.

****

I learn how to ride a bike. Having done so, it could be said I have the knowledge 'ride a bike.' I could take a population of people with the same knowledge, quiz them on it and/or observe them riding. I could come up with a description of what the knowledge 'ride a bike' means.

We probably already agree that such a decription doesn't help me much if I haven't actually ridden yet. Sure, it tells me my feet go on the pedals (unless I'm doing stunts in which case I can put them on the handlebars, etc.), my hands on the handlebars, I move my feet in a circular, clockwise motion, and so on. But until I actually DO something--try to balance, try to pedal--such "knowledge" means little. So a description doesn't help learning a whole lot--it certainly doesn't give an account of it.

However, what about one bicyclist? How much does the decription really help us understand an individual's experience? Does it tell us about the [unique] nature of Bob's bicycling experience--the fact he uses a $2000 pro-racing bike in velodromes--as opposed to Frances, who rides a $200 huffy mountain bike off-road? What does 'bicycling' mean to each of them, as well as to the five-year old who just learned and the person riding for 60 years? And ignoring interpersonal differences (many of which AREN'T important) what about INTRApersonal differences? EVERY TIME I get on a bike, the situation is at least a little different. I know how to 'ride a bike,' but how do I know to begin braking sooner when it's raining, to actually turn the wheel a little to the right and THEN turn it left in order to turn left; why do I go around the pothole, and stay to the right side of the street, and downshift when going up an incline? Are all these going to be included in a description of 'ride a bike'? Then it's going to be pretty darn long. Are we going to stick to "universals" (everyone sits on a seat, grasps handlebars, rides forward...)? If so, then what does this tell me about one's own experience?

You can probably guess I would now ask "What does it mean to say I know how to speak English?" The upshot is that I have a hard time appreciating a description of knowledge apart from its implementation in an environment by

an organism. This doesn't mean we don't "know" things. I just don't see the attempt to describe this as a state of being as possible. This is not simply negating the performance/competence (or similar) distinction, I don't think, it's interpreting behavior in a fundamentally different way.
******
So long for now
==========
Date: Wed, 13 Nov 1991 16:51:28 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: information

[Martin Taylor 911113 14:00]
(Bill Powers 911112.0700

>The "reduction of uncertainty" definition has been extended to the
>content of messages. This definition of uncertainty depends on thinking
>in terms of a finite or enumerable set of possible meanings, shared
>between transmitter and receiver, about which one can be uncertain. This
>is not what the Shannon-Weaver concept was about.
>
Well, according to "The Mathematical Theory of Communication" (Shannon and Weaver, 1949, which is open in front of me, it was. The receiver's uncertainty before the message "x" was H(y), and afterward it was H sub x (y), the difference being the transmitted information.

But Shannon did not restrict the concept to discrete items. All of the theorems are replicated for probability density distributions and for discrete probability distributions. Only in the discrete case is there any concept of an enumerable set of possible meanings.

>Note that the question of channel capacity, which entails a conderation of
>random noise that destroys "bits", implies *variability*, which
>is not the new meaning given to "uncertainty." The new meaning (as of the
>1950s) assumes that there is a finite universe of m possible choices, and
>that the message pins down exactly n of the dimensions, leaving an
>ambiguity -- a choice among alternatives -- exactly equal to the number
>m-n of choices not pinned down. So there is really no noise involved in
>this kind of uncertainty. I say "Buy the Honda," and you reply,
>uncertainly, "The green one or the red one?" It never occurs to you that I
>might have said "Buy the Hyundai." That is not the kind of uncertainty now
 meant
> . It's assumed, in fact, that the communication channel is perfectly
>noise-free, so that if the same message is sent 100 times, the same
>uncertainty will result 100 times.
>

No. Uncertainty implies what it seems to. If there is no variability, there is no uncertainty. And even in the 50's (by which I presume you mean McGill and Garner), there was no intrinsic restriction to a finite universe of possibilities, although I grant you that most of the perceptual experiments on "good form" and the like were based on finite sets, as were the explicit analogies between informational analysis and ANOVA. It is NOT assumed that the communication channel is noise free, and in my discussions with Bruce, that idea never occurred to me. Such a channel is physically impossible, anyway.

>Quantitative measures on a continuum, such as are represented by the
>position of a meter needle on a scale or the time-interval between two

>neural impulses, can't be divided into convenient discrete packages. The
>basic "uncertainty" in such measures isn't visible until you get to the
>quantum level. With continuous measures, the more appropriate concept is
>signal-to-noise ratio. This concerns the repeatability of messages, not a
>choice among alternative meanings.
>
SNR can be measured in units of entropy, as Shannon showed, and that fact
is really the basis for the modern theory of the ideal observer that
underlies the last 30-40 years of psychoacoustics. The variability can
be seen in two different ways, which might or might not give the same
results: (1) Supposing we have an exact provider of the quantity being
measured (e.g. 1 volt is not 1.000....0001 volts), and that provider
varies the quantity, what is the probability density for the measure
read from the meter at a given measure value (e.g. what is the probability
density for a reading of 1 volt if the true value is .995 or 1.005 or ...);
(2) Given that the quantity remains the same, what is the probability
density distribution of readings that will be taken from the meter.

The first interpretation considers only the measure that was reported,
and deals with the distribution of probabilities that the "true" value
was Q volts, whereas the second asks what is the probability distribution
of readings when the true value is exactly V volts. The first leads
to Bayesian (so-called "subjective") probabilities, the second to
classical frequentist (so-called "objective") probabilities, but both
have entropic interpretations of exactly the same kind (see Part III of
the Shannon-Weaver book, or any of many reports out of the Birdsall-Tanner
group at Michigan in the 50's and early 60's, or McGill and/or Garner
in the same era).

Shannon chose his information measure because it was additive when
intuition said it should be; if you kept communicating for 2N seconds
over a given channel you could get twice the information that you would
get in N seconds. McGill and Garner showed that this additivity could
be used to partition the uncertainty inherent in a situation in ways that
elucidated the situation very well, and Garner showed that these partitions
did predict performance in various perceptual tasks, in particular the
development of categories and the concept of "good form." A pattern had
"good form" if small variations made it into a perceptually different
category, and conversely if many widely different patterns were assigned
to the same group those patterns did not have "good form." They looked
like noise.

When an event happens, which might be a message from a cooperative partner,
or it might be an observation of a neutral world, probability distributions
change within the observer. Over some set of concepts, the uncertainty
before the event (e) was H(c), and afterwards it was H sub e (c), the
difference being the amount of information ABOUT THOSE CONCEPTS provided
by the event (e). True, that difference is just a number, but then so
are the values that are differenced between the reference and the percept
in an elemental control system, which doesn't stop them from being ABOUT
some structure in the conceptual world of the controller in which the
elemental control system is embedded.

A second partition is based on the "good form" of the event (e). If the
event is, say, that a ball fell off a table, it fell in exactly one way, but
there were many other ways it could have fallen and still been seen as
the "same" event (a few microseconds earlier, or perhaps even off the
other side of the table). All those events would have provided the
same information. They would not have been distinguished by the observer,
perhaps because the observer's sensory apparatus could not distinguish

them, perhaps because the distinction was irrelevant to the "meaning" of the event. In any case, the information communicated by the event was H(c) - H sub e (c).

From a third-party view, neither H(c) nor H sub e (c) can be known, so the third party cannot determine how much information the observer got from the event. But a third-party who knew the circumstances of the event might be able to judge the probability distribution of events that could have happened and thereby determine the maximum information that the observer could have obtained. That maximum would be based on the reduction of uncertainty from the probability distribution of events before the one that happened did happen to the state of effectively zero uncertainty (given sensory limitations) after the fact. Call this H(e). Assuming that the observer knew the same as the third party about these prior event probabilities, H(c) - H sub e (c) <= H(e). The difference is the redundancy of the event set (there is no redundancy in the event itself, only in the set of events from which it was drawn). The redundancy is another important partition of the information in the situation. It is determined by the structure of the event set, and when the events are overt communicative events, it is controlled by the syntax, lexical choice probabilities, and so forth. These constructs are constructs of what communicative acts can occur, and cannot be determined from any single communicative act. That structure does two things: it reduces the amount of information that can be passed in a communicative act, and it increased the probability that the act will be observed as belonging to the intended event set, so that its meaning is taken correctly by the observer.

I don't suppose all that is as crystal clear as I would like it to be. It is not cast in control theory terms, because the concept of feedback does not apply here. We are dealing with one connection in any of the control loops involved in a control system, and EVERY such connection is subject to the same laws. They are as inevitable as gravity, not optional. You can develop models in which some of the uncertainties are set to zero, much as in Marken's parable some of the control equation values are set to zero, but in the real world equivalence classes exist, measurements have uncertainty, disturbances occur. If you set selected uncertainties to zero, you can get good results in many cases, but when you start talking about communication and language, you had better be very careful about where and when you do it.

Simple summary: uncertainty is variability, or equivalently it represents a probability distribution that is not simply one-zero. Structure is redundancy, and generates equivalence classes that have the same meaning. The amount of information received by an observer (equivalently the recipient of a message) is determined by the changes the event makes in the observers probability structure for his internal concepts, and is not measurable by a third party. In effect, it represents the meaning of the event to the observer (quantitatively, not qualitatively).

Martin Taylor

==========

Date: Thu, 14 Nov 1991 07:50:54 -0700
From: POWERS DENISON C <powersd@TRAMP.COLORADO.EDU>

Subject: Epistemology, levels, information

[From Bill Powers (911113.0700)]

Wayne Hershberger (911112) --

>.... you are using the word accessible to refer not to the possibility
>of epistemic access but to the amount of X that is accessible or to the
>difficulty of achieving access to X, which could vary, of course, with
>the directness or complexity of the epistemic process.

No, I am referring to the possibility of PROVING epistemic success. A model may be epistemically correct, but we will never be able to prove that. I was not saying that we have complete epistemic success part of the time, and incomplete success the rest of the time. I was saying that we do not EVER know whether our models of reality are successful or not, because the only way to check for that success is to repeat the process that led to the models in the first place: there is no independent check.

Perhaps we could get to the nub of this matter sooner if you would give me one example -- any example -- of a case in which we have complete epistemic success in verifying that there is a real counterpart of a perception, any perception. We apparently agree that there are some cases in which uncertainty remains, so there is no point in dealing with them. My claim is that there is NO case in which we have reached certainty, so you should be able to demolish my claim with a single counterexample.

Joel Judd (911112) --

>The principle in phonetic terms(?): If you must vary a point of
>articulation in order to affect the phoneme, and if the phoneme cannot
>exist if the point of articulation does not exist, then the phoneme is
>of a higher level than a point of articulation.

Hmm, and I just said that we shouldn't consider articulation to be at a lower level than phonetics. Let's see if we can figure this out. You've brought in a consideration I hadn't forseen: the difference between causality and levels of perception. You have to turn the screwdriver to perceive the screw going in, so inscrewing is at a higher level than screwdriving. Why, then, do I have the feeling that they're at the same level?

Maybe I could put it this way: perception of A is at a higher level than perception of B if the perception of A depends on perception of B, and if it is necessary for perception B to change if perception A is to be affected. Now it's clearly not true that perception of a phoneme depends on perception of a point of articulation, because you can hear phonemes without saying them. But you can't perceive a phoneme without perceiving some intensity of sound, and the perceived intensities (in various frequency-selective filters) must be altered if the phomeme is to be altered. So the phoneme perception is at a higher level than the sound-intensity perception. Is that better?

Very good stuff on knowing, with epistemological connotations.

Martin Taylor (911112) --

Got to send this off -- answer to your post tommorrow. In brief, I defer to the man with the book in front of him -- and I appreciate being told firmly but kindly when my mental model has drifted far off the mark.

Naturally I still have some points to defend, even with weakened resources.

Best to all

Bill P.

==========

Date: Thu, 14 Nov 1991 11:42:15 -0600
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Re: The Blind Men and the Elephant

[from Gary Cziko 911114.1130]

Rick Marken (911111)

Rick, I like your approach to the blind men and the elephant essay. Here are a suggestion and a question.

First, the title. I don't see only blind men here. I see two blind men and one very insightful man in your story. So perhaps you need to change the title.

Second, you say in desribing the "reinforcemnt interpretation of behavior:

>Behavior seems to be selected by its consequences. This happens when
>r is constant AND d is constant (call it D). All that is changed is
>the relationship ,g, between o and d. For example, d could be a certain
>amount of food deprivation. The organism has a fixed reference for the
>amount of food (say 10 units). The deprivation causes the input food to
>be perceived at 0 units. g determines how much output must be generated to
>compensate for this deprivation to bring the food input to 10 units.
>g determines how much output must be generated to make up for x amount of
>deprivation. If g is demanding, then the animal responds more per unit
>deprivation than it would if g were lenient. This is the schedules of
>reinforcement and, probably, matching findings.
>
>o = g(D)

But as I understand this (and as I believe Bill Powers has made the point many times) operant conditioning theory says that more reinforcement is supposed to INCREASE response strength and less reinforcement is supposed to REDUCE response. But here you have it the other way around (less reward, more response; more reward, less response) which is indeed what happens and makes sense only within a control theory perspective. So your concepualization of the effect of reward here does not seem to square with operant conditioning theory as I understand it. --Gary

==========

Gary A. Cziko
University of Illinois
Educational Psychology
210 Education
1310 South 6th Street
Champaign, Illinois 61820-6990
USA

Telephone: (217) 333-4382
FAX: (217) 244-0538
Internet: g-cziko@uiuc.edu (1st choice)
Bitnet: cziko@uiucvmd (2nd choice)
N9MJZ

==========

Date: Thu, 14 Nov 1991 13:02:23 EDT
Comments: Converted from PROFS to RFC822 format by PUMP V2.2X

From: "Bill CUNNINGHAM - ATCD-GI (804" <CUNNINGB%MON1@LEAV-EMH.ARMY.MIL>
Subject: Re: information
In-Reply-To: note of 11/13/91 21:16

|Bill Cunningham 911114.1300
(Martin Taylor 911113.1400)

Best summary I've read anywhere, anytime. Thank you.

==========

Date: Thu, 14 Nov 1991 12:50:29 -0600
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Writing Backwards

[from Gary Cziko 911114.1200]

In my continuing search for portable demonstrations of Perceptual Control Theory in action, I was playing with handwriting this morning. He are some things to try. I will assume that you all write with your right hand. Lefties need to substitute left for right and vice versa.

1. Write a sentence with your right hand. Notice how easy and quickly you can do this.

2. Write a sentence with your left hand. Notice how much more difficult and slower this is. Nonetheless, if you take your time, you can probably still write quite legibly (I can, anyway).

3. Now write another sentence with your left hand, but write BACKWARDS, that is, from right to left with the letters laterally tranposed (the way DaVinci wrote, I'm told). I found this to be even more difficulty than 2. From a motor perspective, we might expect it to be easier, since the actions are the mirror image of what is done with the right hand. But I think it is harder since we simply don't have a good idea of what reversed writing should LOOK like. I found that I would often "freeze" at the beginning of a word. I did not "freeze" when I wrote in the normal direction with my left hand.

4. Now try to writing some individual words (still lefthanded and backwards and laterally reversed), but now write normally with the right hand the same words simultaneously. You should be able to write quite easily backwards this way since now you have "motor" reference level which you can use for the left hand.

5. After a little bit of this, you may find that you can just "imagine" the right hand writing normally and then write backwards with the left hand without much difficulty. But you still probably don't have a good idea of what your backwards writing should look like. So even though you're looking at your left hand write backwards, it seems (to me anyway) that the visual feedback is not used very much. It is sort of like writing normally but with your eyes closed.

6. Now, you can switch back and forth between forwards and backwards writing with your left hand. And it feels quite different. Writing normally (left to right) with my left hand, I am using primarily visual feedback since I know what it is supposed to look like. It is quite slow, but I can make it look quite good if I take my time. Writing backards with my left hand (after the using the simultaneous right hand trick) is much faster and feels just fine, but looks pretty awful since I am using proprioceptive feedback, not visual.

In addition to its use as a demonstration of some key PCT ideas, handwriting might be a good way to do research on reorganization. Only pen and paper are needed and the subject leaves a permanent record of his or her behavior with no need for fancy computers and C compilers. In addition to the above tasks, you can see reorganization in action by holding a mirror at the head of your paper and writing so it looks normal in the mirror. This makes you write upside down. This is maddeningly difficult. You can see your runaway streaks of positive feedback as you try to make a line go down and it keeps ascending faster and faster the harder you try to get it to descend. This reminds me of Martin Taylor's account of Seymour Papert learning to ride a bicycle wearing reversing prisms as eyeglasses, but it doesn't hurt nearly so much to make a mistake.--Gary

==========

Gary A. Cziko
University of Illinois
Educational Psychology
210 Education
1310 South 6th Street
Champaign, Illinois 61820-6990
USA

Telephone: (217) 333-4382
FAX: (217) 244-0538
Internet: g-cziko@uiuc.edu (1st choice)
Bitnet: cziko@uiucvmd (2nd choice)
N9MJZ

==========

==========

Date: Thu, 14 Nov 1991 22:14:43 EST
From: goldstein@SATURN.GLASSBORO.EDU
Subject: method of levels

To: CSGnet people
From: David Goldstein
Subject: the method of levels
Date: 11/14/91

As you may recall, "the method of levels" plays an important role in PCT Psychotherapy. It is a way of raising a person's awareness so that he/she can become aware of background experiences(perceptions) to the one which started out the conversation. As a result of attending to background experiences, a person's awareness rises to higher levels of perception(experience).

The method of levels is the process that Bill Powers went through which resulted in the different levels in PCT. The method of levels is a "bottom-to-top" procedure. The therapist starts where the patient's awareness is ("bottom") and by looking for background experiences, helps to move the patient's awareness higher ("top").

There is no assumption that the specific levels mentioned in Bill Power's book are the ones which will be found for a particular person when the method of levels is applied. For this reason, Bill once described the process to me as the method of relative levels. If an experience, say experience A, is "in the background" for another experience, say experience B, then A is "at a higher level" than B. Experience B is a means of achieving experience A.

CSGnet people who want to see what PCT Psychotherapy looks like when the specific levels are used as the basis of therapy are directed to Ed Ford's Freedom From Stress Book. Ed follows a

"top-to-bottom" strategy. He has people identify the important system level experiences in their life. Then for each system level concept, he has them identify the important principle level perceptions which are the means of achieving it. Then for each principle level perception, he has them identify the important program level perceptions to achieve a given principle level perception.

Those CSGnet people who have been part of CSG for a while, know that I have been attempting to apply the method of levels in clinical situations. I am sorry to say that I was not having much success with it. This frustrated me greatly. I was experiencing big error signals. I had come to the conclusion that either I was doing it wrong or the method of levels was not performing as advertised.

For the past several months, Bill Powers has been generous enough to do the method of levels with me as a way of helping me with the dilemma I just described. We communicated back and forth using our internet email addresses. The purpose of this post is to describe the kinds of things I learned as a result of doing the method of levels with Bill. I will let Bill tell you what he learned about the method of levels from doing it with me. For ease of reference, I was the "patient" and Bill was the "therapist." Actually, we went through two rounds of method of level. The first start-off topic was my reactions towards the method of level as I understood it. The second start-off topic was my reactions to the experience of playing tennnis.

What I learned:

1. Each statement which the patient makes has numerous potential background experiences which the therapist can ask the patient to address. For some reason, I used to think that the background experiences would only show up after rather large segments of conversation. Each statement goes into and through the therapist who is looking for background experiences along with the patient. The therapist is trying to simulate the patient within the computer of the therapist's own brain and body.

When Bill and I were doing the method of levels, the notational convention emerged to put the background material in brackets. [like this]. Bill started doing this. Then when I was writing my post, I would put background experiences which I noticed in brackets as I became more sensitized to what a background experience was like. [I thought this was a cute notational invention.]

2. The background experience feels more like an observation than an inference. I got the best results if I could observe the feeling or thought which was in the background. It did not feel like I was drawing a conclusion or making an inference.[Tell me what it is like to draw a conclusion or make an inference] It felt like I was making an observation. This helped to give me confidence that the background experience was something which was just as real as the topic which started the conversation. Prior to doing this method of level exercise, I had my doubts about the reality of the background experience. I was not trying to please the therapist, I was trying to observe something inside of me. Either it was there or it was not to observe.

3. The therapist is much more active in identifying the background experience than I have understood to be the case from Bill's general description of the process before we did the exercise. I don't think that, in most cases, the patient can do this on his/her own, at least in the beginning. As the process goes on, the patient does become better at identifying background stuff.

[Tell me about the therapist being more active] The therapist identifies a background experience for the pateint and asks if the pateint wants to address it or prefers a different topic or prefers to continue on the same topic? The patient does have a choice. It is too much to expect that the patient will be picking the background topics, at least in the beginning.

The length of the therapist's answer makes a big difference. If it is too short, the patient feels alone in the enterprise. [I might as well be doing this on my own.] If it is too long, the patient starts to focus too much on the therapist. [I am not here to listen to his problems.] I was not aware of the importance of the length of the therapist's remarks before the exercise.

4. The identification of a background experience also feels a lot different than receiving an interpretation from a therapist. Bill and I wound up calling this more traditional approach to therapy "psychologizing." The result of psychologizing was that I felt annoyed to have to address stuff which seemed to come out of the blue from Bill. When we were following the method of levels, I felt as though I was addressing my stuff. Psychologizing reliably resulted in blocking the flow of the conversation and progress. I am sure that giving interpretations has useful roles in therapy but I am more aware of the negative side effects it can have than I was before the exercise. I will be asking my psychoanalytic friends how they avoid the negative side effects of interpretations.

5. The method of levels is not as abstract or difficult to do as I had thought. In fact, I observed that really good ordinary conversations sometimes follow the method of levels. One person says something. The other person tunes into the background stuff and addresses it. I no longer believe it is restricted to highly intelligent, verbal adults who are intellectually oriented. In fact, since the therapist meets the patient whereever the patient's awareness is located, it probably is applicable to any age group to which other verbal therapy approaches would be attempted. The therapist has to be willing and able to adjust to the state of the patient.

6. There are no fireworks emotionally or intellectually when the level of awareness is raised. The changes feel much more subtle. The patient is not transformed into someone who the patient never was.

It is true that that the patient may become aware of stuff which he/she was formerly unaware of. As I went up levels, the feeling aspects of the experience seemed to diminish in intensity. At the lower levels the feelings were stronger and more salient part of the experience.

7. I did become aware of an internal conflict. Becoming aware of the conflict did not automatically result in the immediate resolution of the conflict. It did put me in a position to think about how I could resolve the conflict given that I now understood both sides. I did give myself a daily assignment which I carry out to help me resolve the conflict. I had not previously identified this internal conflict. I can see how it has resulted in some significant inconsistencies in the way I am/behave.

If this were a real therapy session, the therapist would probably have to spend time helping the patient figure out ways to resolve the conflict once it was identified. It is too much to expect the patient to do this on his/her own. While reorganization would suggest an answer eventually, most patients would prefer the resolution to occur sooner rather than later. That is why they come to a therapist for help.

8. The end result of the exercise was to start to examine my self-image. When we got to this point, the method of levels was more difficult to apply and of questionable utility. When you are at the highest level, there is no more background stuff [Other than the voice of God].] It was here that I decided to continue the exploration on my own and that the joint exercise has stopped.

Bill tried to introduce a new exercise at this point, one which is found in Eastern systems of thought. The point of the exercise ,it seemed, was to help the patient get above the verbal/logical levels and to realize that the patient was creating the thoughts which were occurring. I saw the point but felt that the exercise was not helping me explore my self-image.

I am now applying the self-image exercise procedure which I presented during the last CSG conference I attended. In this procedure, I pretend that I am talking to an actor who will play me in a movie about my life. My job is to give the actor general kinds of statements which will serve as guides for the actor to portray me in the movie. The general statements take the form: Be this kind of person. Don't be that kind of person. As a second step in the exercise, I order the statements from most important to follow to least important to follow. As a third step in the exercise, I write additional statements for each general statement which relates to the general statement in some way. Perhaps, the more specific statements answer the questions of how or why the actor should follow the general statement. Perhaps, the more specific statement is a specific memory which relates to the general statement.

9. I do not believe that the brief demonstrations at the CSG conferences were good enough for someone to learn how to do the method of levels. [Don't feel bad if you didn't learn, I didn't.]

Nor were the verbal descriptions good enough to learn how to do it correctly. It is one of those slippery things which is very hard to learn or teach. It is easy to think you understand it but really don't.

Dick Robertson once told me that he had done the method of levels exercise with Bill but ran into some difficulties when he went to

apply it. If I remember correctly, the difficulties occurred when Dick was working with a patient who did not seem to have a self-image.

This difficulty which Dick Robertson ran into is consistent with the idea that the method of limits has its best utility at the lower levels but not at the self-image or systems level. It is good to recognize the limits of utility of the method so that one does not "throw out the baby with the bath water."

I am now remembering that Ed Ford voiced the idea that the method of levels could not be used to resolve conflicts at the systems level. This is a really good point, I now realize.

For those CSGnet people who have read Ed's book on Freedom from Stress, I see the method of levels as playing a role in the first three steps of the counseling process: exploration (outside-the-skin world--perceptions, reference perceptions or wants), evaluation (body world--feelings/emotions/moods), and committment (why a person may or may not want to work on a problem identified in the first two steps).

10. I will try to create some more effective method of helping people learn the method of levels and present it at the next conference.

I would like to close by stating that the dilemma I was experiencing, me OR the method, was resolved. It was me AND the method!

Date: Fri, 15 Nov 1991 02:31:32 -0800
From: Jim Hess <jhess@ORION.OAC.UCI.EDU>
Subject: Re: method of levels
In-Reply-To: Your message of Thu,
14 Nov 91 22:14:43 -0500. <9111141918.aa25619@orion.oac.uci.edu>

> To: CSGnet people
> From: David Goldstein
> Subject: the method of levels
> Date: 11/14/91
>
> As you may recall, "the method of levels" plays an important role
> in PCT Psychotherapy. It is a way of raising a person's awareness
> so that he/she can become aware of background
> experiences(perceptions) to the one which started out the
> conversation. As a result of attending to background experiences,
> a person's awareness rises to higher levels of
> perception(experience).
>
> The method of levels is the process that Bill Powers went through
> which resulted in the different levels in PCT. The method of
> levels is a "bottom-to-top" procedure. The therapist starts where
> the patient's awareness is ("bottom") and by looking for
> background experiences, helps to move the patient's awareness
> higher ("top").
>
Etc.

It strikes me that there are certain similarities between this theory and what an actor would call looking for the "subtext". The

assumption is that a character in a play or a person in a social situation has certain goals which they are pursuing. These goals may be conscious, but often are not. The actions, verbal and otherwise, of the person are like moves in a game in which the person seeks to fulfill their desires by manipulation of their social environment. To make sense of the person's or character's actions, you need to understand them at the higher level of these goals. This understanding may be pursued analytically by some actors, intuitively by others. It's interesting to note that a person or character who is conscious of their goals and purues them directly is perceived as strong, dynamic, and forceful. In a play, at least, they usually play a central role in furthering the plot. George Bernard Shaw was especially fond of them; they were drawn as individuals who saw through social conventions to essential processes. They broke the rules or challenged beliefs, which Shaw saw as creative destruction.

-Jim Hess-

==========

Date: Fri, 15 Nov 1991 09:41:52 EST
From: Joseph Michael Lubin <jmlubin@PHOENIX.PRINCETON.EDU>

The following came from the
Biomechanics and Movement Science listserver <BIOMCH-L@HEARN.BITNET>

From @pucc:BIOMCH-L@HEARN.BITNET Fri Nov 15 09:01:37 1991
Received: from pucc.PRINCETON.EDU by phoenix.Princeton.EDU (4.1/1.110)
id AA01604; Fri, 15 Nov 91 09:01:33 EST
Message-Id: <9111151401.AA01604@phoenix.Princeton.EDU>
Received: from PUCC.PRINCETON.EDU by pucc.PRINCETON.EDU (IBM VM SMTP V2R1)
with BSMTP id 2123; Fri, 15 Nov 91 09:01:17 EST
Received: by PUCC (Mailer R2.08 R208007) id 0978; Fri, 15 Nov 91 09:01:14 EST
Date: Fri, 15 Nov 91 08:37:19 EDT
Reply-To: "Daniel B. Sheffer" <R1DBS@AKRONVM.BITNET>
Sender: Biomechanics and Movement Science listserver <BIOMCH-L@HEARN.BITNET>
From: "Daniel B. Sheffer" <R1DBS@AKRONVM.BITNET>
Subject: NIH Guide announcements
To: Joseph Lubin <jmlubin@phoenix>
Status: RO

Two announcements in the 15 Nov 1991 NIH Guide might be of interest to readers of biomch-l.

$$R1 BEGIN NIH-NINDS-92-01 **

CLOSED LOOP CONTROL OF FUNCTIONAL NEUROMUSCULAR STIMULATION

RFP AVAILABLE: NIH-NINDS-92-01

P.T. 34; K.W. 0745047, 0740050, 0706040, 0715140

National Institute of Neurological Disorders and Stroke

The Neural Prosthesis Program of the National Institute of Neurological Disorders and Stroke (NINDS), NIH, is developing neural prostheses based on functional neuromuscular stimulation (FNS) for the restoration of quadriplegic individuals. The principal goal of the proposed project is to enhance the utility of FNS systems for hand grasp. Specific tasks include: developing and evaluating closed-loop FNS

systems for hand grasp utilizing external force and position integrating FNS wrist stabilization, FNS elbow control and surgical procedures such as tendon transfer and arthrodesis with an FNS hand grasp system; developing a biomechanical model of the hand for use in evaluating advanced FNS systems; and investigating new techniques for programming of FNS systems and for selection of electrode sites. A research team with experience in neural prostheses, spinal cord rehabilitation, hand surgery, control theory, biomechanics and physiology will be required to successfully conduct this research. It is anticipated that one award will be made for a period of three years in August 1992.

This is not a Request for Proposals (RFP). To receive a copy of the RFP, submit a written request to the following address, and supply this office with two self-addressed mailing labels. All responsible sources shall be considered by the agency. The RFP will be issued on or about November 15, 1991, with proposals due on January 14, 1992.

Contracting Officer
Contracts Management Branch, DEA
National Institute of Neurological Disorders and Stroke
Federal Building, Room 901
7550 Wisconsin Avenue
Bethesda, MD 20892
Attention: RFP No. NIH-NINDS-92-01

==========
DANIEL B. SHEFFER, PH.D., HEAD BITNET R1DBS@AKRONVM
DEPARTMENT OF BIOMEDICAL ENG. INTERNET R1DBS@VM1.CC.UAKRON.EDU
UNIVERSITY OF AKRON TELEPHONE 216 972 6650
AKRON, OHIO 44325-0302 TELEFAX 216 374 8834
==========

Joseph Lubin jmlubin@phoenix.princeton.edu
Civil Eng. Dept. 609-799-0670
Princeton University 609-258-4598
Princeton NJ 08544

==========
Date: Fri, 15 Nov 1991 07:48:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: epistemology

[From Wayne Hershberger]

Bill Powers (911114)
>Perhaps we could get to the nub of this matter sooner if you
>would give me one example -- any example -- of a case in which
>we have complete epistemic success in verifying that there is a
>real counterpart of a perception, any perception. We apparently
>agree that there are some cases in which uncertainty remains, so
>there is no point in dealing with them. My claim is that there
>is NO case in which we have reached certainty, so you should be
>able to demolish my claim with a single counterexample.

I seem not to be making myself clear because you are looking for me in the wrong direction. My argument (actually, Hume's) is not that any particular case affords certainty but rather that EVERY CASE IS ENTIRELY UNCERTAIN (i.e., "verifying that there is a REAL

counterpart of a perception" is absolutely impossible, even as a matter of degree). Hume's arguments to this effect are called Radical Skepticism for good reason. Consequently, modern science uses a coherence, rather than a correspondence, theory of truth--where reality has no ontological status.

When asked how he discovered the laws of chemical compounding, Linus Pauling replied, "I made them up." Pauling avoided any claim of having gained epistemic access to Reality--because such a claim would be gratuitous (God's Reality is a matter of faith, and serves no scientific purpose). Rather, Pauling made up a parsimonious model which provides a very COHERENT account of the chemical phenomena in question. Any claim that such man made models correspond, in varying degrees, to some divine original is epistemically empty.

You have made up a parsimonious model of living control systems which control the value of inputs from their environments. Sometimes you have used the word "virtual" to refer to these controlled variables because they are defined by the input functions which process the input. But it would overstate the case to claim that the environment contributes nothing to the values comprising these virtual variables. That is, only by overstating the case is one misled to suppose that your model implies solipsism. Your model addresses questions of correspondence between what is inside the brain and what is outside the brain but that is physiology not philosophy.

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu

==========

Date: Fri, 15 Nov 1991 10:26:21 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: configurations and transitions

[From: Bruce Nevin (911115 0823)]

I've been thinking some more about the assignment of configuration and transition perceptions to levels 3 and 4.

I read in BCP:131 as follows:

> Nearly any perception of the first three orders, may apparently be detected at fourth level as _change_. The phenomenon I am calling here the "perception of transitions" may explicitly involve rates of change of any perception of intensity, sensation, or configuration. [In experimental work reported by] Notterman, Filion, and Mandriota (1971) . . . pure rate-of-change detection was clearly shown.

It seems to me that the configuration level is the first level at which we naively reify our perceptions. A configuration is often understood to be a "thing". (Included also are such things as distances, orientations, depth signals, size perceptions, etc.--BCP:126.)

  * Below configuration level, intensity and sensation perceptions are always attributes of a thing by which we perceive it, never things in themselves.

  * Above configuration level, events and relationships clearly involve things already perceived as such, and are not attributes by which we recognize them.

A transition may be a change of one or more of the attributes of intensity or sensation by which we perceive a configuration as such. Does it not seem that transitions of this sort are below configuration level?

A transition may also be a change in the reference signal for control of configuration, defined as body posture. Varying the signal results in the cat swiveling its head back and forth. This sort of transition is clearly above configuration level.

What other transitions are above the configuration level--changes in the reference signal input to an elementary control system (ECS hereafter) on the configuration level? Let's look a bit more closely at the sort of transitions that I have proposed are below configuration level.

We perceive something as a configuration it seems in two ways:

(A) by differences in intensity and sensation

  * From one part of it to another
  * From it to its environment (figure and ground both involve configuration perceptions)

(B) by "invariances" in those lower-level perceptions--or perhaps in the differences of (A)--as they change through time

  * As the object changes shape, size, color, etc.
  * As the object moves relative to its environment and us
  * As we move relative to the object
  * As we shift our attention from one part to another of the object
  * As the environment changes around the object

I put the word "invariances" in quotes because it glosses over the way that certain differences don't make any difference (it's still the same face), and others do. The trick is in the perceptual input function summing perceptual signals passed up to the ECS on the configuration level from ECSs below it on the intensity and sensation levels, as Bill said (BCP:125):

> When we notice a solid real object, we are experiencing the output of a very sophisticated [configuration-level] perceptual function, one capable of ignoring all changes in its input sensation signals except a certain narrow class of changes.

The output of this perceptual function provides an "invariant" signal to the comparator on the configuration level. Suppose that the transition level ought realy to be broken out into two distinct levels:

  * Lower-level transitions<1> are ignored, or compensated, or cancel one another in the configuration-level perceptual function.

  * Higher-level transition<2> perceptions control changes in the

reference signal of a configuration-level ECS in the case of postural control, and presumably in recognition of the intensity and sensation "invariants" by which a configuration is recognized.

Configuration perception depends on transitions<1> in intensity and sensation perceptions. These are differences that don't make a difference because of the way the perceptual input function of the configuration ECS sums them. If they do not match something like a recognition template (as is implicit in this conception of the neural mechanism), there is no corresponding configuration perception because the output of this function does not match the reference signal.

Query: in recognition, do we think that the reference signal of the configuration ECS varies so as to "hunt" for a possible match to its perceptual input?

The bottom line, of course, is the question how neural signals are connected in living organisms that have a configuration level of perceptual control. What is the new physiological evidence since BCP?

Bruce Nevin
bn@bbn.com

==========

Date: Fri, 15 Nov 1991 07:12:09 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Information

[From Bill Powers (911114.1100)]

Martin Taylor(911113) --

OK, now I have more time to respond to your post.

I said:

>>The "reduction of uncertainty" definition has been extended to the
>>content of messages. This definition of uncertainty depends on thinking
>>in terms of a finite or enumerable set of possible meanings, shared
>>between transmitter and receiver, about which one can be uncertain.
>>This is not what the Shannon-Weaver concept was about.

You said:
>Well, according to "The Mathematical Theory of Communication" (Shannon
>and Weaver, 1949, which is open in front of me, it was. The receiver's
>uncertainty before the message "x" was H(y), and afterward it was H sub
>x (y), the difference being the transmitted information.

With more time to reflect, I will stick to my statement at least a little longer. I said that Shannon and Weaver were not talking about uncertainty in the meaning of the message, and they were not. They were talking about uncertainty in the message itself, and as you say, they came up with a measure related to signal-to-noise ratio. If a message n bits long is sent, then the number of possible messages that could have been sent is 2^^n. The log of this number is related to information. Or, better, I should say that the word "information" has been assigned formally to mean a particular function of this quantity. This is similar to the way "intelligence" is formally defined: that which intelligence tests measure.

So over the universe of 2^^n messages that could be received, which

message was sent? If there is noise, each bit has a probability of being a 1 or a 0. This leads to a calculation of the uncertainty as to which message was received. If the signal-to-noise ratio is extremely small, we might be very confident that the transmitted message, translated from bits into ASCII, was "Rain tomorrow."

We look at the clock, and realize that it shows 24:00:00. At what time was the message *intended* to be received? That crucial piece of information is not in the message. The received message contained no significant uncertainty *as a message*. It definitely says "Rain tomorrow," with a physical uncertainty of 10E-9. But the probability, at 24:00:01, that it means it will rain today is only 0.5. Of course it may not rain at all for the next week -- that's a different kind of probability, concerning the truth of the alleged meaning. I stand by my statement: Shannon and Weaver spoke only about fidelity of transmission, not about meanings (or truth). Whatever they *thought* they were talking about. They must have assumed that the meaning of a message is self-evident in the message, if they meant what you say they did.

>SNR can be measured in units of entropy, as Shannon showed, and that
>fact is really the basis for the modern theory of the ideal observer
>that underlies the last 30-40 years of psychoacoustics.

The relationship of SNR to entropy depends on which game you're playing. If you're thinking in terms of message transmission as modulation of transmitted energy (the normal engineering context), the relationship is meaningful. But in other physical contexts it's only a metaphor. As I said in my previous comment on this subject, energy does not have to travel in the same direction that the message travels. I gave examples (telegraph with battery at receiving end; neural impulses) in which energy travels opposite to the direction of propagation of the message. With a little fanciful but realizable design, you could arrange for the NET transfer of energy to be zero. So if information transfer is not related to physical energy transfer, it isn't related to physical entropy, either. The fact that you can run measures of a message through an equation that has the same form as the entropy equation (and of course you can) thus has no deep physical significance. It's just something you can do if you find the results useful.

You ask two questions:
>(1) ... what is the probability density for a reading of 1 volt if the
>true value is .995 or 1.005 or ...); (2) Given that the quantity remains
>the same, what is the probability density distribution of readings that
>will be taken from the meter.

Neither of these questions is the one I raised. When you read a meter, the meter needle appears at a certain position on the scale. Whether or not the reading is visibly fluctuating, there is no uncertainty about its APPEARANCE. Both of your questions relate to how that appearance is connected to the "real" value of the reading, an epistemological question that I was not bringing up (but do now). The uncertainty is in the explanation of the reading, not in the reading itself. If the reading is steady, it's steady; if it's fluctuating just so, it's fluctuating just so. There is no uncertainty in either appearance. I was just making the point that the underlying space in which we perceive analog quantities is continuous, mappable onto the real numbers, not the integers. So it does not have "states" -- or the number of possible states is as infinite as the number of real numbers, as infinite as the number of different frequencies a neural signal can have. If the number of possible states is infinite, then the formally-calculated information contained in any one

state is zero.

>Shannon chose his information measure because it was additive when
>intuition said it should be; if you kept communicating for 2N seconds
>over a given channel you could get twice the information that you would
>get in N seconds.

So if I say "Roses are red" twice, I transmit twice the information that could be obtained from saying it once? In terms of bits, yes. In terms of meaning, no. From a different point of view, one could maintain that only one bit of "information" has been added by the repetition of the message: the fact that it has been duplicated. Of course this could be a code: there is no message unless it is sent twice. In that case .. well, you get the point. This is another instance of confusing message information with meaning information. There are all kinds of meanings you can get from a given message. What they are depends on everything else you know and assume about the world, including but not limited to what you know and assume about the transmitting agent (how much do you know about Tom Brokaw?). The transmitter can give you information that was not intended to be sent, such as the fact that the guy at the other end doesn't like you. Heck, you can get two meanings into one message: "Mother, John can't tell you what your visit has meant to him."

>When an event happens, which might be a message from a cooperative
>partner, or it might be an observation of a neutral world, probability
>distributions change within the observer.

I understand that by "probability distributions" you don't mean "probability calculations," but signals generated in ways that produce the same outcome that a probability calculation would produce. This means that signals can be averaged continuous functions of other signals -- probabilistic fluctuations are NOT necessarily implied. Signals, like probabilities, have one value at a time. I do not think that the world we perceive contains a significant amount of random fluctuation.

>If the event is, say, that a ball fell off a table, it fell in exactly
>one way, but there were many other ways it could have fallen and still
>been seen as the "same" event (a few microseconds earlier, or perhaps
>even off the other side of the table). All those events would have
>provided the same information. They would not have been distinguished
>by the observer, perhaps because the observer's sensory apparatus could
>not distinguish them, perhaps because the distinction was irrelevant to
>the "meaning" of the event. In any case, the information communicated by
>the event was H(c) - H sub e (c).

The events are not distinguished at the category level, but of course at other levels the observer can certainly see which side it fell off, whether it bounced, what color it was, and so on. Only at one level of perception are the different ways of falling "the same." The categorizing process, as you imply, determines the meaning at the category level. This may or may not be the category intended by someone who used this "falling" to illustrate a point. The same set of events perceived in terms of other categories illustrates "carelessness" or "gravity."

The observer, however, perceives the way it did fall and not all the ways it could have fallen: or to put that more in terms of my epistemology, the way the ball was perceived to fall is taken to be the way it did fall, and not ways in which it was not perceived to fall. The way it did fall could be known only to an observer whose perceptions have somehow been calibrated correctly against reality. To calculate the amount of

information you would have to know how the ball actually fell, without using ANYONE's a perceptual signals as a measure. This applies to category perceptions as well as to the events classed in the same category. One would have to know all the possible categories into which the actual occurrance could have been put, and what the "actual category" is.

It's beginning to creep in upon me that you and I may be having the same argument that Wayne and I are having. You say

>From a third-party view, neither H(c) nor H sub e (c) can be known, so
>the third party cannot determine how much information the observer got
>from the event. But a third-party who knew the circumstances of the
>event might be able to judge the probability distribution of events that
>could have happened and thereby determine the maximum information that
>the observer could have obtained.

But the third party has exactly the same problem in determining HIS OWN information about the event. I hope I'm not quibbling here over the adoption of a naive-realist viewpoint for the purpose of making models. I hope I'm raising a question as to whether information and probability can actually be determined. If the calculation of maximum information depends on the third party knowing ALL the ways in which the ball could have fallen, then it's not an objective measure, or even a repeatable one. The number of ways in which the ball COULD have fallen (including ways the observer hasn't thought of yet, but will, and including ways that never will occur to a particular observer but might occur to a different one) can't be determined, or is infinite, take your choice. So the amount of information in the actual case is either indeterminate or zero.

The Achilles Heel of arguments concerning probability and information is the frequent impossibility of determining the true universe of possible occurrances or choices. Both probability and information are relative to assumed or agreed
boundaries on the universe, and to what you know within that universe. As long as you confine these arguments to a small enough universe with enumerable elements, they hold water. But doing that just makes this into a game, like the Prisoner's Dilemma (which I solve by having the prisoner whip out an Uzi and escape, or by having the sentence declared unconstitutional, or some such expansion of the rules). We all have to play this game, of course -- but the question is, how willing are we to keep playing by the same rules when some reasonable expansion of the universe of discourse shows that they are arbitrary and not the only possible rules?

I agree that neural signals themselves contain statistical fluctuations due to their underlying granular character. But we have to begin with experience, not theory: the world we experience shows very little of this noise. If the noise exists, it must be very small or it must be averaged before it turns into the world of perception. The averaging processes are mathematically similar to statistical calculations: means of sums over time. So at some fundamental level we're talking about the same thing.

I feel about this subject much the way I do about Ashby's concept of "variety." It may be true that control systems "reduce the variety of essential variables" and that the means of doing so must contain an amount of "variety" that matches the "variety" of external disturbances. But how do they do that? Is it enough just to match mean variety, or does the match have to be quantitative, instant by instant? If I went to an engineer and asked for a system that would "reduce variety," I might get

anything from a control system to a hammer and nail or a pot of glue. Once we know how a control system works, we can see that it does in fact reduce variety in Ashby's sense. But variety-reduction is not what makes it work. Similarly, information flow is not what makes a transmitter and receiver work, and probability is not what makes neural circuits work. Once we understand how these things work -- have a good model, and can build imitations of them -- perhaps we can see how abstract things like variety, information, and probability can provide valid characterizations of their workings. But if we understand how a system works, what do we add by such abstract descriptions? I'm really asking -- I don't know where to draw the line. Maybe this is just the old top-down, bottom-up argument.

Best

Bill P.

==========

Date: Fri, 15 Nov 1991 12:29:49 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: Exhortation

[from Joel Judd]

To All:

I heartily exhort ye to attend the SLATE lecture this Monday evening by Gary Cziko. I promise you will find it mind expanding and perspective-changing. It is at 7:30pm in 407 Illini Union.

==========

Date: Fri, 15 Nov 1991 13:13:34 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "CHARLES W. TUCKER" <N050024@UNIVSCVM.BITNET>
Subject: THE TEST AND S-R BEHAVIORISM

FROM CHUCK TUCKER 911115.1315

A THOUGHT THAT I HAD WHILE CATCHING UP ON MY READING

I am very behind with my reading but I read a post by Rick Marken on 911025 (Yes, that is October) trying to explain to Martin (I think) the difference between the S-R and PCT view; his point was that the S-R view does not see that behavior is control. Then I thought, How could they? What we see as a disturbance they see as a stimulus. A person alters their behavior from an S-R view because the stimulus from the environment acts as a stimulus to alter it rather than to maintain a purpose. One reason that they see it this way is that they don't use THE TEST (Powers; Runkel) when they examine the conduct of a person. [I think their are other reasons too but they deal with higher levels] If this is so, then one service that the NET can perform is to inform the readers of THE TEST and to apply it when questions regarding PCT come up.

Is this a good thought or what?

Regards,

Chuck

==========

Date: Fri, 15 Nov 1991 12:31:31 PST
From: marken@AERO.ORG

Subject: Blind men and elephant & more

[From Rick Marken (911115)]

Wow, what a huge amount of wonderful stuff on the net. I can see why it is hard for people to keep up if they don't plod through this stuff every day.

I would just like to thank all those who made nice comments or suggestions about the "Blind men and the elephant" paper. I know, by the way, tht much of this is not original and I will cite all those who have noticed this stuff before. I just wanted to package these ideas in one concise paper. Bill Powers, of course, was the first to point out that the behavior of a control system can look like s-r. Tom Bourbon's models and experiments (which, of course, I plan to cite) also make the point that control can look like planned output. And Wayne has been very articulate in his descriptions of the different ways of looking at the behavior of a control system. As Gary Cziko points out, the relationship of control to operant phenomena can be confusing if you think of the main operant phenomenon as "reinforcement increases the strength of behavior. My point about reinforcement in the "blind men" paper is, I think, basically correct for some operant phenomena (like the scheduling results.) It is true, for example, that rats will press more frequently on a demanding schedule -- which seems to violate the "reinforcement increases the probability of response" concept (because they produce more behavior per reinforcement, I guess). But I think that that phenonmenon (increased probability of a response followed by a reinforcement) only works in the initial stages of learning a new behavior -- during shaping.Otherwise, the apparent phenomenon is just that behavior is controlled by its conseqiences (somehow) so if those consequences are difficult to get you get high respond-ing, if easy you get low responding. If you have two choices then the prob-ability of making those choices depends on the probability of getting the reinforing consequence after each choice. Those seem to be the kinds of things that operant conditioners play with. But I would appreciate any thoughts on a clearer way to present the appearance of "control by consequences" when looking at a control system. I think my analysis -- o = -g(D) -- is OK for many cases but I'm sure it could be improved. Perhaps by writing the equations in terms of controlled variables rather than outputs?

For those of you in my APS symposium (Joel Judd, Tom Bourbon, Bill Powers) -- I hope to have the application ready to go out by next weekend. The registration fee for the meeting has not yet been determined but I am told that it should be very close to what it was last year -- $35 for members (like Tom and I) and $50 for non-members (like Joel and Bill).

David Goldstein -- looks like a great post. I have to read it more care-fully. But it is quite timely since I was trying to explain the method of levels to my wife just last weekend. I bet you can do a much better job than I did.

Gary -- Great handwriting demos!! I agree with your interpretations of what's going on. Maybe you could publish (on e-mail) the current list of "portable demonstrations" of control theory. I think these demos will be extremely useful.

Regards

Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Fri, 15 Nov 1991 15:28:00 CST
From: "HRLOT1::SMALLWOOD" <smallwood%hrlot1.decnet@HQHSD.BROOKS.AF.MIL>
Subject: Removal From Mailing List

To Whom It May Concern:

A former employee, Tom Hancock, was on your electronic mailing list, at the following address:

HANCOCK%HRLOT1.DECNET@HQHSD.BROOKS.AF.MIL

Hancock no longer works here. Please remove his name from your electronic mailing list.

Pam Smallwood
Armstrong Laboratory Computer Systems Manager

==========

Date: Sat, 16 Nov 1991 07:16:18 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Method of levels

.LM0/.RM79/[From Bill Powers (911115.0730)]

David Goldstein (911115) --

David, thanks for the description of the method of levels in operation. I, too, learned a lot from our interchange. The most important thing I learned was that the urge to "psychologize" is almost irresistible, and that every time I failed to resist it, progress totally stopped. I would think "I'm stuck -- I don't know what to ask about now -- I don't see where he's coming from." So I would start to guess. I never guessed right unless you had just told me directly, even without realizing it, what was going on in the background. The urge to psychologize (in me) came from wanting to control the process: it didn't seem to be going anywhere, and I wanted to get it started again. So I would ask something like "when you were sad, were you also a little afraid?" (that isn't what I asked). All my vast experience suggested that this was likely. The "patient" told me, in effect, "Who cares about your vast experience? We're talking about MY experience."

In this particular case, psychologizing was doubly bad because you knew how to do it too, and we were left with both of us off the track.

The "track", as you explained it so well, consists mostly of observing, not reasoning or deducing or associating. You look for what's actually there. There's no need to guess.

I also learned that it does help (and may well be necessary) for the listener to offer as much explanation as necessary to make sure the other person understands what kind of thing to look for. The other person also has to understand that the point is not to psychologize (I feel this way

because that happened to me, because I've always felt that way, because I have anxiety, and so on). The Method is not psychotherapy in any recognizeable sense. The point is to observe (I just thought, "I hope they're following this") or to check and see if the listener's guess was correct (if you didn't catch yourself saying something like it out loud, which people often don't).

For me, the most important positive thing I learned was that the method works if you just keep doing it (until it stops working, and then it's time to quit). If neither you nor the other person grasps what the background activity is, continue asking for descriptions of or elaborations on the foreground. Sooner or later an out-of-context remark will drop that provides the hint [I'm getting bored with this]. Then you just ask the person to elaborate on that topic [Tell me about being bored -- what does boredom feel like?] .

It's also interesting that this method seems to work over email, without face-to-face contact. I've known you for many years, but that fact only leads to psychologizing, so I don't think that is the explanation. It probably would have gone much faster if we'd been face to face, but maybe not -- time for reflection may not be a bad thing.

I often missed obvious things and only thought of them later -- maybe that's another reason to leave intervals in the process. To speak generally, one thought that finally occurred to me that when a person says "this is the kind of person I am," that can equally well reflect a perception or a reference level (an observation or a self-fulfilling prophecy). If a person says "I'm not very good at sports," this can easily turn into a reference condition: if I'm not very good at sports, I'd better not try to learn tennis -- that would be out of character (to pick up on Jim Hess' remark about actors). After all, if I learned tennis I could become pretty good at a sport, and I'm not very good at sports. Self-characterization defines a reality, and we often take great pains to defend against disturbances of reality -- even if we don't like it. But that is surely verging on psychologizing, if not well over the border!

For newcomers:

The method of levels is based on an observation that we often are working on (at least) two conscious levels at once. The "foreground" level is the main thing you're paying attention to. The "background" level is occasionally apparent as thoughts, attitudes, or feelings ABOUT the foreground level. Ordinarily there is always a background process running, commenting on or reacting to the foreground material. But your attention may not be on it. The method of levels consists of a talker talking about some subject, and a listener listening for comments "meta" to what is being talked about. When the listener detects some "meta" aspect of what is going on, the listener asks the talker to check and see if it is going on. I'm indebted to Hugh Gibbons for calling this process "going meta." I suppose lawyers are well-acquainted with this process. Maybe Hugh could be persuaded to discuss an interesting session on this subject that he and I had with some of his law students last year.

When your attention shifts to the background level, and you begin to talk about it, ANOTHER background level sooner or later shows up.

The method of levels arose years ago when Kirk Sattley and I wondered how many times in succession a person could switch attention from foreground

to background (making the former background into the foreground). It turned out that this process does not involve infinite regress: it tops out. It was an interesting experience. It was even illuminating. And it turned out to be almost impossible to do without someone else to listen for background-type things that escape the talker's notice even though they are said out loud.

The connection with a hierarchical model of perception is obvious, but long ago I decided that trying to tie these foreground/background pairs to a fixed set of definitions would be needlessly limiting. If I've guessed correctly about the levels, they will show up. If I can't recognize the definitions in the presented material, so what? In fact even thinking about those definitions gets in the way of the kind of concentrated split-level listening that's required.

To find out what the background level is, the listener listens and asks. For example, while you've been reading this section, has there been some thought about it in the background (yes, you)? Have you been wondering what the point is? Have you been thinking "this is like blah blah blah"? It's pretty likely that you've had SOME background attitude or thought about what you've been reading, at the same time you were reading it (and are reading this). All you have to do is notice what it is. Then if you go on thinking about that, elaborating on it for a while in your thoughts, some new background thingie (it could be anything) may become apparent. You can probably do this with the current level and one above it. But going much farther alone is hard. You get interested in the new level and forget that you're supposed to do it again. So does the person who's listening, if he or she isn't careful. The point isn't to be found in any of the levels that are uncovered. It's to keep going. Any psychological significance in what is discovered will be self-evident.

As the listener, it's ok to guess if you've had only hints and aren't sure you get the nature of the background material. But it has to be pre-arranged with the talker that the talker is to look and see if the guess is right, and reject it if it isn't. Only the talker actually knows what's going on in there now, in present time. When the listener starts trying to lead, trying to persuade the talker that the listener's insight is correct, you're off into psychologizing and the method has stopped. It's OK to ask, on sufficient evidence, "Is it jealousy in the background?" It's not OK, if the person says "no," to then say "Not even a little bit?" If the person is jealous and doesn't want you to know, it's none of your business anyway. But the chances are that you're just wrong. You have to leave your clever insights at home. The key idea is to respect the other person's autonomy. When a background process is noticed, the talker will know for certain that it is real. All you do is follow the trail.

David, I've probably not noticed a lot of things I did right and wrong -- your comments are helpful, even though I've no plans to do this for a living. It might be helpful for you to describe your experiences with applying the method, from time to time.

The method led me to some theoretical concepts about attention. It seems that attention can focus on different levels of organization. It seems that the level within or from which awareness is working is NEVER itself the object of awareness. Awareness takes one point of view toward other activities in the brain. Going up a level brings that point of view itself into the field of awareness -- and awareness is then operating from another point of view that is itself unawared.

This apparent phenomenon raises the question of why awareness can move around like this, and why it seems beneficial to do this. I conjectured that awareness has something to do with focusing reorganization. The degree of reorganization going on depends on how much error there is, but the aspect of organization that's being changed depends on what is in attention, what is the object of awareness. This would help to ensure that reorganization is applied where it is needed, and not just at random to systems in perfectly good working order. I call this a conjecture and that's what it is, but it makes some degree of sense.

People like us, who spend a lot of time reasoning and communicating, probably operate from something like the program level, attending to lower-level things like sentences and words. We "think" a lot, but the thoughts are really background activities: they are not about thoughts, but about other things. We don't go around saying, "Hey, look at this, I'm thinking!" As a consequence, whatever reorganizing is going on is not going on at the program level: we are executing programs, but reorganizing the things that the programs are about. No reorganizing goes on at the principle or system concept level, either. We tend to get into ruts -- I into the rut of control theory, each of you into some other rut. That's why we keep demonstrating the same principles over and over, almost by rote. One set of ideas has become a fixed framework within which thoughts occur, and the same thoughts even occur over and over. Only the material about which we think changes; we apply the same algorithms to everything.

This is a comfortable and useful level, but it's not the highest level. I think it's good, now and then, to pop up to higher levels and examine your own thoughts just as phenomena, reconsidering them from the level of principles. And I think it's good to pop up now and then to the system concept level to check out the principles, vary them a little, see if they're still what you want to have there, in terms of being who you are and being in the system you're part of. It might be swell to pop up above the system concept level occasionally, but I'm not sure if that's possible. It's worth trying, though, if my conjecture about reorganization and awareness is correct.

To do these things you don't need to think "Now I'm going to be at the principle level." All you have to do is look at the programs you're running through your head as thoughts. You can't see them from the program level, AS PROGRAMS. As soon as you do see them as programs, you'll be projecting principles onto them -- that's where you have to be in order to look at the programs instead of just running them.

So you can see how I would connect this awareness-shifting phenomenon with therapy. The problem in therapy is not just to reorganize, but to reorganize at a level where changes won't just be changed right back by higher-level systems. People can stew about the same problems for years and years (even in therapy), because they're reorganizing where the symptoms of the problem are and not where the cause of the problem is. This is why the method of levels pays so little attention to what is going on in the foreground. The point is to get to a higher level, in fact to the level that is choosing conflicting goals. That's the only level where un-choosing can take place. David mentioned that there aren't many fireworks associated with the method. That's because when you're reorganizing at the right level, there's nothing else to say you shouldn't. There's no more effort involved in resolving a conflict than there was in the innocent choice that originally established it. You realize "Oh, I can't want this AND that." If there are any fireworks,

they occur AFTER the resolution has taken place, as one starts to do things differently and discovers all the adjustments to the conflict that now can be dis-adjusted.

I hope some of you listening in will grab a friend and try this out. It's a most interesting experience even if there's nothing wrong with you.

Best to all

Bill P.

==========

Date: Sat, 16 Nov 1991 15:21:16 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: ooops

[from Joel Judd]

Hey, it was a real %(#*@$(# PLEASANT experience to see a message I meant for a local ESL group go out to CSG ("Exhortation"). EXCUSE me. (Unless of course you'd like to come and hear Gary's talk). That's two e-mail goofs in as many months.

==========

Date: Sat, 16 Nov 1991 21:14:16 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Paraplegics; epistemology

[Bill Powers (911116.0800)]

Joe Lubin (911115) --

Thanks for that relay from NIH, Joe. So someone is going to get closed-loop control into muscle stimulation for quadraplegics! That is wonderful news. I remember a PBS show about that guy in Indiana who was using open-loop muscle stimulation, and even with all the lurching and staggering that resulted, the patients were so joyful that tears came to my eyes. I thought then that closed-loop control would work wonders (position pots at the joints). I even made arragements to talk to this guy (whose Russian-sounding name I've mislaid), but the person who was to make the contact never did it, and I let the matter lapse, to my shame.

I don't have a medical research team, of course, but I'm going to write to the address at NIH explaining my interest, and ask if I can serve as a consultant to help review the proposals from the CT point of view. It would be a pity if this opportunity were messed up just because someone didn't understand control of input.

Wayne Hershberger (911115) --

Comes the dawn.

>My argument (actually, Hume's) is not that any particular case affords
>certainty but rather that EVERY CASE IS ENTIRELY UNCERTAIN (i.e.,
>"verifying that there is a REAL counterpart of a perception" is
>absolutely impossible, even as a matter of degree).

I feel like a wrestler who has converted an advantage to a position flat on his back. Your arguments (especially your references to Gibson) have given me the impression that you believe there is a reality outside of perception AND that perception somehow manages to represent it

veridically. Now it seems that you're saying that human perception bears no verifiable relationship to any universe "behind" or "beyond" perception, which is of course the position I have also been taking.

Unfortunately, your language still leaves me wondering what precisely is your position on the constraints we can detect between actions and perceptions, and on the significance of models. You say

>Any claim that such man made models correspond, in varying degrees, to
>some divine original is epistemically empty.

If it weren't for that word "divine" I would be more sure of how to take what you mean -- I trust you're not accusing me of religious fervor. Would you still allow for a correspondence to a "non-divine original" -- i.e., a natural universe that exists independently of our perceptions of it? In other words, are you opposing a religious view of reality, or any view that there is (or could be) a reality more inclusive than what is perceived, whether or not we can be certain about its nature? More on this at the end.

The main conundrum comes up this way:

>Your model addresses questions of correspondence between what is inside
>the brain and what is outside the brain but that is physiology not
>philosophy.

A problem is created by talking about what is inside the brain and what is outside the brain. The problem arises when we assume that we, as conscious beings, are conscious because of activities in a brain. Allow me to elaborate.

In order for my model to be consistent with the physical model of reality (both, I quite agree, being "made up" so that physiology, too, is "made up"), there are certain relationships between physical-model variables and neural-model variables that must be assumed. The physiological neural model allows for no way of getting information from physical stimulation other than through interaction of physical variables with neural sensors. For example, in the physical model there is a made-up entity called the photon. The signals in the neural model's retina supposedly arise from absorption of photons. However, the neural signals carry no information about the origins of those photons; furthermore, there is an infinite number of different photon energies and fluxes arriving from an infinite number of directions that will yield exactly the same neural signal in any given receptor.

Given a model of physical optics and observations of reflection or emission sources, we can construct a model of the origins of the photons, and show that this model is consistent with an array of neural signals that amounts to a map of the scene toward which the eye's lens is directed. So far, so good. But if we then look at the basis for accepting the physical model -- which includes things like "lenses" and "objects" and "light rays" going through something called "three-dimensional space" -- we find that there can be no basis but observations made by the same means by which an "observation" of a photon is made. We identify objects by looking at them with our eyes; we verify that there is a photon flux by interposing a light-meter (which we see) and reading -- with our eyes -- its indicator. So from the standpoint of the neural model, the physical observations we are using to assign an external source to the visual neural signals arrive in the brain by exactly the same means as the signals we are trying to explain.

This is a not a problem if we adopt a point of view from which we can see both models, the model constituting an exterior physical world and the model constituting an interior world of signals in a brain. It is not a problem if we add to the other two models a model of a non-neural conscious observer which is not confined to a brain. It becomes a problem only when we decide that the model of the brain must be a model of ourselves, the observers and thinkers.

When we adopt that view, as I do, we can no longer take the third-part omniscient view. The hypothesis is that we, who are thinking about perception, are brains like those in the model. Therefore we must be dealing with the external world (represented by the physics model) by the same means just proposed: through neural signals. If this is true, then the physical model is NOT outside the brain. It must be located inside the brain-model, as part of that model. It is a construction existing as patterns of neural signals related not by physical constraints outside us but by abstract rules and computational processes taking place in our heads.

When we apply this reasoning to purely physical models, there isn't much difficulty except with people who insist on reifying photons, electrons, quanta, phlogiston, and so on. The real difficulty arises when the external world we are thinking about is the world of subjective reality: the world we experience directly. This is clearly not the world of the physical models. Between the physical models and this world of direct experience there are few points of contact. For the most part, physical models consist of entities and relationships that are not evident in direct experience. Here and there are points where, usually through the use of instruments but not always, a physical variable corresponds to an experiencable variable. With the unaided eye, we can perceive an approximation to what a physicist calls "distance," although using instrumentation like radar or optical range-finders we can arrive at meter-readings much more consistent with physical theory than is the direct apprehension of distance. But when it comes to functions of distance such as gravitational acceleration or potential energy, direct experience remains blind.

So where do we put this world of direct experience, with all its objects and sounds and smells and relationships and people? It is not represented in either the physical model or the (physiological) neural model. I have elected to put it into the neural model, but not in the form of neurons. It exists in the brain as a weightless, massless organization of neural signals the appearance and behavior of which is precisely the appearance and behavior of the world we experience. Certainly this assumption creates a mystery; more than one. The main one is, who or what is it that apprehends this collection of neural signals in such a way that it takes on the appearance that we experience?

The most obvious error to be made at this point is to say that this mysterious observer is the agent that imposes interpretations on the neural signals so they become objects, relationships, processes, concepts, and so on. But as everyone knows by now, that simply requires expanding the model to explain how these interpretations are made. My way of avoiding this error has been to propose the levels of perception in my model. By looking for classes of perception in the apparently real world around me, I attempted to show how neural processes can themselves create signals that contain the interpretations that are needed. While this initially seems to rob experience of some vital qualities, a close examination of any particular example of this problem shows that it does

not exist. These vital qualities can't be pinned down by direct inspection, either. When one attempts to isolate them for a close look, they lose any special quality and become just an amount of something that can be more or less present. Just like a neural signal. The only specialness that there is exists in the entire collection of neural signals, each behaving in the context of all the others.

The other function of the levels is to enumerate and classify types of perceptions ranging between what have been considered "concrete" and "abstract" perceptions. By showing how successive levels of interpretation can form a link between the concrete and the abstract, the model removes the necessity for explaining these interpretations by assigning them to a homunculus. As each new level is considered, the subject-matter with which it deals is stripped out of the homunculus and returned to the physical brain. In the end, the homunculus contains only those functions of observation that are not accomplished by the brain model. And all that is left is awareness.

We now seem to have a model that is itself a neural model that contains a physical model in a nervous system. But the nervous system is basically a physical conception: it is a subset of the physical model. Logically, if one model is contained by and contains another model, there can be only one model. But there is another answer: it is that both models are contained in direct experience, and interact with each other.

The ultimate reality, therefore, is direct experience. That is the superclass within which models exist. This leads us, finally, to the ultimate mystery.

We can divide direct experience into things we do and things that happen. Many things that happen proceed without any need for our action. Among such things we can detect consistencies and dependencies. This leads us to formulate expectencies which, when formalized, we call laws of nature.

We can also take actions, which are the set of all those experiences that we can influence by an act of will. We find that these actions, themselves experiencable, affect other experiences. We can learn to create some experiences that are not directly subject to acts of will by varying those experiences that are directly willable: the whole is an act of control. Through long experience with this kind of act of control, we have found regularities that show how we must act in order to control many kinds of experiences. The reasons for these regularities are not evident in experience -- there is no a priori basis for expecting any particular act to have any particular effect on something else.

This is where we get the idea of a natural world of regularity that lies outside the boundaries of experience. And this is why we build models, both physical and neural. With models we hope to probe into that mystery that is hinted at by these unexplained regularities. We hope to reduce the complexities of these piecemeal regularities by finding underlying simplicity; I think this is what we mean by "mechanism." In physics, simplicity is attained by imagining a hidden world of fields and particles, energy and momentum and entropy. The few kinds of variables in this world lead to the vast multiplicity of different-seeming phenomena in the world of direct experience.

The question, Wayne, that you and I have not brought out into the open and resolved between us, is whether these models constitute increasingly good approximations to something beyond experience, or whether they are simply "summaries of observations." The complex picture I have tried to

lay out here should indicate my view. Clearly I don't think that either physics or neurology is as good an approximation as is usually assumed. There is too much of the human observer entwined, unanalyzed, in all our models. The very name "particle" physics shows this. But I think that there is evidence of agency outside us (other people, for example), and evidence of relationships imposed by unseen means (e.g., other people's intentions). I think that there is structure inside the Reality Box, and that while we can never arrive at a unique representation of it, we can arrive at an equivalent representation, equivalent in the sense that our models show one way it could be constructed inside, functionally equivalent to the way it is constructed. I see no contradiction in saying that all we will ever know for certain is what our own brains present to awareness, while maintaining that uncertain knowledge is not empty.

Is the remaining problem, perhaps, what is meant by "epistemic?" If "epistemic" knowledge is certain knowledge, then the argument resolves itself: there is no such thing outside direct experience. But to say that a proposition is epistemically empty does not then mean that it is incorrect or empty of signficance, because that would say that all of experience is, with complete certainty, incorrect or empty of significance beyond itself -- an epistemic fact which, of course, we can never verify.

Best

Bill P.

==========

Date: Sun, 17 Nov 1991 15:21:54 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Paraplegics; epistemology

[Martin Taylor 911117 15:00]
(Bill Powers 911116.0800)

A brief, but important point:

Bill, in answering Wayne, says:
> But if we then look at the basis for accepting
>the physical model -- which includes things like "lenses" and "objects"
>and "light rays" going through something called "three-dimensional space"
>-- we find that there can be no basis but observations made by the same
>means by which an "observation" of a photon is made. We identify objects
>by looking at them with our eyes; we verify that there is a photon flux
>by interposing a light-meter (which we see) and reading -- with our eyes
>-- its indicator. So from the standpoint of the neural model, the
>physical observations we are using to assign an external source to the
>visual neural signals arrive in the brain by exactly the same means as
>the signals we are trying to explain.

But one actually can test the "existence" of the things detected through photons by using other senses--the acoustic effects, Bishop Berkeley's (?) kick, and so forth. These form a set of converging operations that help to reduce the set of possibilities for interpreting the perceptions obtained through one sensory system.

If that were all there was to it, the same argument could be made, but extending the notion from "photon" to "physical energy exchange phenomena" or some such. But that is not all there is to it. There is the volitional aspect of what

and how we choose to observe.

Let us presume a deceitful Nature, and a passive (multi-sensory) observer. This Nature could present us with any of an infinite number of sources for peeps (Physical Energy Exchange Phenomenon, plural) that had the same effect on our sensory organs. But when WE choose which aspects of the universe to test, and in what way, the deceit becomes much harder to sustain. That's the fundamental difference between an observational science like Astronomy and an experimental one like physics. Psychology is somewhere in between.

The difference between active observation and passive observation was clear to Gibson, who distinguished "haptic" from "tactile" perception. You can try it yourself (demonstration #473b!). Have a bunch of objects available, and a friend. Close your eyes, and have the friend take one of the objects and touch it to your open hand in various orientations and ways. What you perceive is a set of touches, some soft, some warm, some sharp, and so forth. Now have the friend place the object in your hand for you to manipulate. Now what you perceive is not a set of touches, but an object. In either situation you might be able to determine which object from the set was touching your hand, but when you yourself choose where and how the object contacts (or fails to contact) your hand, it has a completely different subjective quality.

In a paper I never published some 20 years ago, I related this effect to the power of generalization from fixed-effects versus that of random-effects ANOVA, but I guess that analogy won't be of much interest to this group (yet?).

It also probably lies at the heart of JG Taylor's theory that ALL that we perceive is based on what we have learned how to control through our behaviour. I am coming more and more to believe that the two theories (Powers and JGT) are complementary, JGT's being interpretable as the control of reorganization within Powers', and the way in which the levels are themselves constructed. This seems to be a grey area in the pure Powers approach, at least as written in the 1973 book and as discussed here in the year or so I have been following this group.

Martin Taylor

PS. Not so brief a point, I guess. Two points, and I'm sorry for the length.

==========

Date: Sun, 17 Nov 1991 20:56:27 CDT
Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: Various

From Tom Bournon [911117]
Gary Cziko -- Your handwriting demonstration (like your earlier exercise in walking while jiggling your eyeball) is nice. Elegantly simple and convincing.
David Goldstein & Bill Powers -- Your posts on the method of levels is informative. So that is what has come of the off-net traffic between the two of you. If neither of you object, I will share your posts with my students, present and former, whose primary interests are in clinical practice.
Rick Marken --
Concerning APS: Do you need anything else from any of us? Must you submit registration fees along with the proposal? (If the proposal is accepted, I will be there; if it is not, I will not attend APS and would rather not waste a deposit fee.)

Concerning blind men and elephants: The idea for the paper is solid. The ensuing discussion about reinforcement and response rates (higher numbers of bar presses on reinforcement schedules that require greater numbers of presses to receive the same number of food pellets) touches on one of the (many) glaring inconsistencies in reinforcement theory. I believe operant theorists (my term for them -- most would deny they are theorists) intend their remarks about reinforcement causing an increase in response rates to refer to a longer-term "increase," as in, "at a later time, the animal is more likely than not to press the bar, after bar pressing is reinforced." see most operant theorists speaking in those rather global terms, not in terms of moment-by-moment actions by the animal.

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402

==========

Date: Mon, 18 Nov 1991 02:56:04 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Levels; perceptual correspondences; misc

[From Bill Powers(9118.0100)]

Bruce Nevin (911115) --

Good thoughts on the levels. It's always nice to know that the concept makes sense to someone else ("Making Sense" was one title I considered for a book).

> * Below configuration level, intensity and sensation perceptions are always attributes of a thing by which we perceive it, never things in themselves.
>
> * Above configuration level, events and relationships clearly involve things already perceived as such, and are not attributes by which we recognize them.

Maybe not "things" in themselves, but they are certainly "perceptions" in themselves. It's difficult to attend to the world without noticing configurations, but it's possible at least to let the configuration (and up) aspects fade into the background -- when you're matching color chips against the paint on a wall, for example. The chips have a shape, but it's irrelevant.

>A transition may be a change of one or more of the attributes of
>intensity or sensation by which we perceive a configuration as such.
>Does it not seem that transitions of this sort are below configuration
>level?
>
>A transition may also be a change in the reference signal for control of
>configuration, defined as body posture. Varying the signal results in
>the cat swiveling its head back and forth. This sort of transition is
>clearly above configuration level.

Good, this is the way it's beginning to seem to me, too. Here are some words for perceptions in the as-yet-unnamed level just below configuration and (I think) just above sensation:

Gradient: edge: texture: discontinuity (corner): curvature: contrast: fine structure (generic name?): pattern (irregular like stucco or tweed): granularity.

A "breathy" or "graveley" or "husky" voice would seem to belong here. The feeling of surfaces -- smooth, slick, wet, sandpapery, woody, slimy, and so on. Trembling, shivering. Gritty spinach.

It isn't so much "transition" as fine structure or granularity that I get from this putative level. A "pure" sensation is just a kind, like a specific color, taste, pressure, temperature, or brightness. When sensations vary randomly or in a speckly way across space and time (the time aspect not being recognized as time-like) we get something in addition to the identity of the sensation. There's also the sense of granularity -- well, what all those words have in common. The result is something built of sensations, and that can be altered only by altering sensations -- yet it's not an invariant configuration. I think Whorf was trying to describe this level when he spoke about -- what was it, "extended properties?" Things that are indefinite in extent yet have space-filling (or surface-covering, or time-occupying) character.

> * Lower-level transitions<1> are ignored, or compensated, or cancel one
>   another in the configuration-level perceptual function.
>
> * Higher-level transition<2> perceptions control changes in the
>   reference signal of a configuration-level ECS in the case of
>   postural control, and presumably in recognition of the intensity and
>   sensation "invariants" by which a configuration is recognized.

Yes. For transitions<1> I think the transition<2> aspect isn't even sensed. Transition<1> is just a static quality. A configuration can be made of different transition<1> perceptions: a diamond-shaped area of roughness on a smooth surface, or the hand-shape on the cave wall made by blown speckles of paint. Maybe this is where you see those circles and squares that are made of scattered corners. I should think that snails and slugs would have transition<1> perceptions of where they're oozing.

>Query: in recognition, do we think that the reference signal of the
>configuration ECS varies so as to "hunt" for a possible match to its
>perceptual input?

I don't think so. Perception is a signal coming out of an input function, so it's the form of the function that determines the input conditions under which the signal will have any given magnitude. In order to accomplish what you have in mind, we'd need a new kind of control signal from above -- one that can vary the parameters of a perceptual function. We may need that kind of connection, but it's not in the official (i.e., minimum) model. I couldn't justify putting it in the model without some kind of experimental demonstration in which we can't think of any other way to account for what happens. Got any examples?

It is possible for a higher-level system to vary a reference signal in order to match it to an existing (intractable) perceptual signal. This could result when the consequence of error is some other error in higher-level, or other low-level, systems. When you open a door of unknown stickiness or springiness, you increase the reference-level for pushing effort until the door starts to swing open; then you keep the reference-push the same or even reduce it to match the sensed resistance of the door, to keep from slamming the door into the wall (which most nice

people try to avoid). If the door is nailed shut, you will limit the magnitude of the push-reference signal: you don't want to pull a muscle. If you concede that a reference signal can be compared against a perceptual signal by a relationship-type perceptual function (courtesy of imagination connection), this doesn't amount to sensing error signals, so the basic perceptual postulate (we experience only perceptual signals) is left intact. Assuming that's good.

>The bottom line, of course, is the question how neural signals are
>connected in living organisms that have a configuration level of
>perceptual control. What is the new physiological evidence since BCP?

Any graduate students out there who need a suggestion for earning a merit badge? I haven't done any systematic study of neurology for many years -- the data are scattered all over the place, obviously not organized to make life easy for a control theorist. When you approach any body of conventional knowledge from the CT standpoint, you really get a picture of how random the experimentation is.

Chuck Tucker (911115) --

>... one service that the NET can perform is to inform
>the readers of THE TEST and to apply it when questions regarding PCT
>come up.
>
>Is this a good thought or what?

Good. Definitely not what.

Rick Marken (911115) --

>But I think that that phenonmenon (increased
>probability of a response followed by a reinforcement) only works in
>the initial stages of learning a new behavior -- during shaping.

This is probably the explanation of the apparent positive-feedback phenomenon when schedules provide extremely little reinforcement. Why didn't I think of that? Tom Bourbon has set up a a control model that uses the E. coli method for adjusting its own integration factor to match the model to real behavior. This could also be done just to show how the model acquires control. The obvious thing to try, based on your idea, is to let such a model come to equilibrium (using absolute error as the criterion for a change in k), and then see how it behaves as you vary the feedback connection's proportionality factor. The prediction would be that there will be more behavior for more reinforcement when the error is very large. What's actually happening is that the randomly-changing behavior has a better chance of producing reinforcement when the schedule is easier.

In operant conditioning circles there's no distinction between reorganization and performance.

Gary Cziko --

Did you get a message from Tom Hancock's former employer telling us to remove his name from CSG-L? I couldn't tell if it was sent to me or to CSG-L (deleted some lines too quickly).

Martin Taylor (911117) --

>But one actually can test the "existence" of the things detected through
>photons by using other senses--the acoustic effects, Bishop Berkeley's
>(?) kick, and so forth. These form a set of converging operations that
>help to reduce the set of possibilities for interpreting the perceptions
>obtained through one sensory system.

There is a convergence to agreement, true, which could be suggestive that we are getting a handle on the Great Unknown. But it could also mean nothing more than that we adapt all our input processes so they achieve maximum agreement with each other. We can explore only those possibilities that are provided for by our sensory equipment. If the universe has four spatial dimensions we will never know it.

>... have the friend take one of the objects and touch it to your open
>hand in various orientations and ways. What you perceive is a set of
>touches, some soft, some warm, some sharp, and so forth. Now have the
>friend place the object in your hand for you to manipulate. Now what
>you perceive is not a set of touches, but an object.

To do this (rather neat) experiment right, I would have to forget my entire history from the time I began to see and feel objects. People who have been blind from birth and have their sight restored (congenital cataracts, etc) have great difficulty learning the correspondence between visual shapes and tactile (actually tactile-kinesthetic) shapes. To identify a triangular object visually they have to count its corners at first. This suggests to me that visual and tactile-kinesthetic space are fundamentally different from each other, and that we have to learn to map each onto the other. It's interesting that "number" seems to transcend modality.

Probably the important common factor in this mapping is the correspondence between seen and felt configurations and transitions when you are watching and feeling your own hands doing things. There is no a priori connection between joint-angle signals and retinal image-movement signals, any more than there is between image-movement and tastes. Only experience tells us that making an arm feel a certain way will make it look a certain way, or that making an image get bigger and go somewhat downward will enable us to taste something, which we later connect with the image. I think we make sense of the world by learning these arbitrary rules. As I said in my last post, I consider the existence of such discoverable arbitrary rules an indication that there's a reality other than the one we experience. But we STILL don't know whether the mapping between sensory modalities itself maps onto a reality in any one-to-one way. If all sensory transforms contained a common conformal mapping, we would never discover it.

>I am coming more and more to believe that the two theories (Powers and
>JGT) are complementary, JGT's being interpretable as the control of
>reorganization within Powers', and the way in which the levels are
>themselves constructed.

When you're ready, how about concocting an essay for us on this subject? Call it a draft of an article.

>Two points, and I'm sorry for the length.

You need to develop a conscience like mine. I repeat from an earlier post: the readers have delete keys.

Best to all,

Bill P.

==========

Date: Mon, 18 Nov 1991 14:10:58 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: information, metalanguage

[From: Bruce Nevin (911118 0932)]

(Bill Powers (911112.0700)) --

>What is the basis on which you picked out the standformers from the
>screwdrivers and wrenches? It can only be that in some sense each object
>in the toolbox looked more or less "like" the standiformer in my hand
>. . . This capacity for
>perceiving "likeness" is part of the capacity for forming categories of
>perception. Some objects looked "very like" the standiformer; some looked
>only a little like it; some didn't look like it at all. If we're to
>understand categories, we have to try to guess what sort of perceptual
>processing is needed to form this variable -- continuously variable --
>sense of categoriness.

My experience of this is that (through memory and imagination) I access higher levels of control involving the perceptions to be categorized. For a screwdriver, I imagine using this thing in my hand as a screwdriver or a standiformer. A "gadget bent at right angles at both ends, with one blade set ar right angles to the direction of the other" qualifies as a screwdriver. If I don't know about your use of screwdrivers (as I indeed don't know anything about what you do with standiformers) I will leave this one out, and I am likely to include an awl and a nut driver in the collection that I hand to you on the basis of appearance. Ask any mechanic about inept assistants. Or remember the kinds of help often given by children.

> We treat
>items that are actually different as if they were equivalent. Language
>isn't necessary for doing this. But this ability is necessary for
>language.

This recourse to program level and other higher levels through imagination and memory vitiates your argument. Language may be involved in all category-and-above perception. That is still a candidate among the options I suggested for the relationship between socially constituted conventions of language and individually constituted hierarchical control of perception.

>The machinery behind [continuously variable rather than quantal]
>activities . . . can compute in many ways. One of the ways it can
>compute we experience as language. Another is logic. Another is
>continuous analogue computation. I can't make a machine that can speak
>if you just hand me a dictionary or a grammar. You have to hand me a
>soldering iron. This is why I no longer think that language is
>bootstrapped by a metalanguage. Language and metalanguage are just two
>of the things that this machinery can do. It's easy to confuse the
>program with the computer. But to understand what we observe, we have to
>try to understand the computer, not just one example of what it can do.

I think I have not made my point clear. I am certainly not saying there is an antecedent metalanguage that bootstraps language. Something like that is in Generativist claims about Universal Grammar against which I

have inveighed. "Bootstrapping" refers to how language came to be in evolutionary process, and also how children acquire language albeit with social support that Cro-Magnon humans were bootstrapping their way into.

Martin's concept of protocols in his "layered protocol model" concerns the role of metalanguage in communication (and in part metadiscourse elements in discourse). I don't think he would agree with your dispensing with this.

If for the control of language we control perceptions of categories of words, and there are words for those categories (words like "noun," "action word" and indeed "word"), those are metalanguage words and metalanguage perceptions. If the program level is involved in control of language, as has seemed evident to you in the past but maybe you have changed your mind, then those programs are metalanguage programs. An example: you can make a relative clause only out of sentences that can be conjoined as interruptions, where a word may appear at the beginning of the interruption that is identified as the same as the word immediately preceding the interruption:

> John--said John you may remember--came by and asked for you.
> John, whom you may remember,came by and asked for you.
>
> John--*said John Mary knows that came by--asked for you.
> *John, whom Mary knows that came by--asked for you.

The second pair are unacceptable English because the following sort of transposition is not provided for on the program level for English:

> Mary knows that John came by -/-> *John Mary knows that came by.

(One can sort of say "John Mary knows came by," meaning "It was John that Mary knows came by," but not at all comfortably. Here, -/-> is an arrow with a slash through it for negation or "disallowment".)

I could multiply different kinds and levels of examples indefinitely.

Any of these metalanguage perceptions (perceptions about language that enable or are involved in or are required for the control of language) may be uttered in sentences of the language. These would be metalinguistic sentences. Most metalinguistic sentences are zeroed, i.e. never uttered overtly. Words like "the former" and "the latter" are overt metalinguistic elements of sentences.

Traditionally, "language" means the possible utterances of members of a speech community. In this empiricist sense of "language," of course, one does not normally or always use language (utterances perceivable by others) in the course of controlling language, though commentary on one's own linguistic performance is commonplace (damn! another typo!). The prospect of infinite regress prohibits "self-instruction" in this literal sense as means of implementation. Suppose we do take the empiricist avenue, regarding language as just the set of possible utterances. Ignoring problems like the fact that such a set is not well defined, we have only perceptible events (sound vibrations, marks on paper) to investigate as language. These have been demonstrated to have a cumulative structure (of constraints built upon other constraints) that is social property of the language users. People apparently learn how to correlate structures in language (linguistic perceptions) with other perceptions that they control.

The correlation cannot be a simple or direct mapping. Most of the structure in language is arbitrary--word shapes, word choice, combination by word sequence or affixation or word change (ring rang rung), word order (Achumawi is VSO, verb-subject-object), and so on and on. Language structures vary geographically and by social stratification in largely arbitrary ways, and language changes over time, all in ways that have little or nothing to do with control of nonlinguistic perceptions by the speakers of the language. There is no reasonable way, Bill, that you can say that control of language is simply an artefact of nonlinguistic perceptual control.

Fortunately, there is an alternative, and commitment to PCT compels us to undertake it. We have to try to understand "language" in terms of the control of perceptions in the perceptual hierarchy. On this view, we of course can't limit ourselves to behavioral outputs (sound vibrations, articulatory movements, marks on paper). But neither can we throw out the results of work in the empiricist view. This is because any substantial collection of discourses in a language displays a complex, systematic, and internally consistent structure of the sort described (with structural provision for variability and change). This structure is not statistically determined by the investigator over many trials, discarding data that are too discrepant to fit the curve. They are determined by language users' perceptual control for their own conformity to shared social conventions.

This is what draws us to this interface between socially constituted categories, sequences, programs, etc. and other perceptions that are not mediated by socially constituted structures. I am suggesting that all of the above (category and above) are socially constituted, learned of a piece with the learning of language. Primates and cetaecians may have programs and even principles, and still lower organisms may have categories and sequences, without language. I would bet nonetheless that these are socially learned and maintained. The existence of higher levels of control in creatures without language (if demonstrated) would *not* mean that we also control perceptions on those levels without language. Having language, we use it to help integrate and control these higher-level perceptions, and most certainly to help teach them.

Conversely, the existence and pervasiveness of language does not rule out a-linguistic control of perception at category and higher levels in us humans. Is skilled athletic performance really automatized, or is it alinguistic control of program perceptions? What we call intuition comes here perhaps. But we can't talk about it. We can't describe it. If we were able to do so, it would be by that act enmeshed in the socially constituted net of language. So a productive task, as I see it, would be to devise experimental means for stalking these nonverbal perceptions and disclosing them for modelling. Since a model doesn't use language, it provides a standpoint outside of language by which to "describe" things without using language. Even language itself. More difficult, that: stalking the perceptions involved in the control of language and disclosing them. With a built-in metalanguage so handy, it is not easy.

The proposal that language is coterminous with level 7-plus perceptual control was but one of a range of proposals about the relationship between language and non-language perception, and between the social and the personal. I'm groping for ways to find out what's right and to find out more about it, whatever it is.

(Martin Taylor 911113 14:00) --

>When an event happens, which might be a message from a cooperative partner,
>or it might be an observation of a neutral world, probability distributions
>change within the observer. Over some set of concepts, the uncertainty
>before the event (e) was H(c), and afterwards it was H sub e (c), the
>difference being the amount of information ABOUT THOSE CONCEPTS provided
>by the event (e). True, that difference is just a number, but then so
>are the values that are differenced between the reference and the percept
>in an elemental control system, which doesn't stop them from being ABOUT
>some structure in the conceptual world of the controller in which the
>elemental control system is embedded.

The "concepts" C1, C2, . . . , Cn are perceptions, each corresponding to an elementary control system (ECS). It is true that the reference signal of one such ECS may be thought of as a number, r, as also the perceptual input p and the error output e may be thought of as numbers. But a probability *across* C1, C2, . . . , Cn is not comparable to these unless you are postulating a single ECS on a higher level whose input mechanism combines signals from C1, etc. in a way that corresponds to the mathematical notion of probability. Are you claiming that?

You might be making the much weaker claim that a number H(c) or H_e(c) represents not only the probability of a particular "message" occurring but also the content of that "message". The disanalogy of probabilities to signals in an ECS applies here too. A signal in an ECS "has content" because of the position and interrelations of the ECS with others in the perceptual control hierarchy.

"Aha!" you might say, "Probabilities at different layers of my Layered Protocol model have similar interrelations in a hierarchy." But the probabilities are not elements in a control hierarchy. They are elements in a descriptive hierarchy.

>In any case, the information communicated by the event was
>H(c) - H sub e (c).

No. The difference between the later probability number and the earlier one provides one measure of the *amount* of information communicated. Only by attaching these numbers to perceptions in a model of hierarchical control could you give them meaning or content and identify the information itself that was communicated. Even then, the numbers would be a quantitative measure relevant to the communication channel between originator and recipient but (I believe) of no use to the hierarchical control of perceptions by either of them.

(Bill Powers (911114.1100)) --

>information flow is not what makes a transmitter and
>receiver work, and probability is not what makes neural circuits work.
>Once we understand how these things work -- have a good model, and can
>build imitations of them -- perhaps we can see how abstract things like
>variety, information, and probability can provide valid characterizations
>of their workings. But if we understand how a system works, what do we
>add by such abstract descriptions? I'm really asking -- I don't know
>where to draw the line. Maybe this is just the old top-down, bottom-up
>argument.

I don't think the use of probability-counts as a measure of quantity of information passed is at the top, I think it is at the outside. It is the point of view of the observer, rather than of the control system.

[From Bill Powers(9118.0100)]

>words for perceptions in the as-yet-unnamed level just below
>configuration and (I think) just above sensation:
>
>Gradient: edge: texture: discontinuity (corner): curvature: contrast:
>fine structure (generic name?): pattern (irregular like stucco or tweed):
>granularity.
>
>It isn't so much "transition" as fine structure or granularity that I get
>from this putative level.
>
>The result is
>something built of sensations, and that can be altered only by altering
>sensations -- yet it's not an invariant configuration.

I was looking at the chest of drawers in my bedroom as I was thinking about this the other day as reported in my earlier message. I was wondering what lay beneath my perception of that chest of drawers as a configuration. It's not moving, so no transitions in that sense. I saw lighter and darker brown of the oak under the old finish, shadow and shading where the frame member curves down to the panel (I was looking at one end). It seems to me that it is by contrasting sensation perceptions in adjacent parts that I become aware of shape and spatial relationship. If it's a new object and I can't rely on memory and imagination, I move around it so that these relationships do shift and change and I construct an invariant 3-D model to support these intensities, sensations, and transitions<1>. The latter may be just a simple instance of recourse to a higher level of control the better to explicate perceptions on a lower one. There are still non-moving contrasts and other relationships between adjacent intensities, sensations, and perhaps textures too, to which my transition<1> refers.

But perhaps more crucially, it seems to me that configuration perception is constructed in memory and imagination. An configuration is apt to include aspects that are not directly perceptible as intensities and sensations (and textures) at the same time as othe such aspects. There must be exploration with the senses through time, and there must be the construction of a model in memory and imagination. Intensities and sensations are more immediate, nothing is hidden.

==========

Date: Mon, 18 Nov 1991 18:54:00 GMT
From: "Franklin Pierce Law, Starbranch,HEP" <FPLC.LAW@APPLELINK.APPLE.COM>
Subject: Tickling

I am finally digging myself out from under the mountain of material that you people have generated - this is a conversation that one comes to late at one's own risk. As a lawyer, my interest is in the way that people make normative choices, distinguishing good from bad and right from wrong. Why, for example, do there seem to be two dimensions to it - good and bad, right and wrong? Control theory has given me a tool for modeling normative thinking.

I will probably have questions about that later, but now I have a question that has bothered me since I majored in physiological psych: Why is it that a person cannot tickle him or herself? I have done a fair amount of reading in cognitive psych recently and find no reference to that phenomenon. Somehow, it seems like a question the solution to which might produce insights. Any ideas.

==========

Date: Mon, 18 Nov 1991 18:54:00 GMT
From: "Franklin Pierce Law,
Starbranch,HEP" <FPLC.LAW@APPLELINK.APPLE.COM>
Subject: Tickling

I am finally digging myself out from under the mountain of material that you people have generated - this is a conversation that one comes to late at one's own risk. As a lawyer, my interest is in the way that people make normative choices, distinguishing good from bad and right from wrong. Why, for example, do there seem to be two dimensions to it - (1)good and bad (2)right and wrong? Control theory has given me a tool for modeling normative thinking.

I will probably have questions about that later, but now I have a question that has bothered me since I majored in physiological psych: Why is it that a person cannot tickle him or herself? I have done a fair amount of reading in cognitive science recently and find no reference to that phenomenon. Somehow, it seems like a question the solution to which might produce insights. Any ideas?

==========

Date: Mon, 18 Nov 1991 14:12:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Re: Tickling

One theory is that there's no surprise when you try to tickle yourself. In control theory terms, you've learned to automate the adjustment of reference signal so there's zero discrepancy. Sort of like gating.
For a phylogenetic perspective on tickling see Ben Beck's book on play.
Peters_R

==========

Date: Mon, 18 Nov 1991 17:15:41 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Levels; perceptual correspondences; misc

[Martin Taylor 911118 17:00]
(Bill Powers 911118.0100)
>
...>>=me, >=Bill
>
>>... have the friend take one of the objects and touch it to your open
>>hand in various orientations and ways. What you perceive is a set of
>>touches, some soft, some warm, some sharp, and so forth. Now have the
>>friend place the object in your hand for you to manipulate. Now what
>>you perceive is not a set of touches, but an object.
>
>To do this (rather neat) experiment right, I would have to forget my
>entire history from the time I began to see and feel objects. People who
>have been blind from birth and have their sight restored (congenital
>cataracts, etc) have great difficulty learning the correspondence between
>visual shapes and tactile (actually tactile-kinesthetic) shapes.

I don't see what difference it makes to forget your history, or what effects newly sighted people have. The point is that active, intentional touching provides the perception of an object to a normal sighted person, whereas an equivalent set of imposed sensations does not. If, indeed, newly sight people have a problem in identifying visual shapes, that's an interesting but so far as I can see irrelevant phenomenon. If anything, it might be supportive, inasmuch as they would not have had any training in voluntary shifting of the visual focus of attention by eye movement or by internal shifts of focus. If that is the case, they might be expected to treat visual sensations much like a normal person treats imposed sensations, and

not make objects out of them.
>
>
>>I am coming more and more to believe that the two theories (Powers and
>>JGT) are complementary, JGT's being interpretable as the control of
>>reorganization within Powers', and the way in which the levels are
>>themselves constructed.
>
>When you're ready, how about concocting an essay for us on this subject?
>Call it a draft of an article.

That's not a bad idea. Perhaps doing that might help me to clarify my ideas
on both theoretical positions. But don't hold your breath, because I doubt
it will happen very soon.

Martin Taylor

==========

Date: Mon, 18 Nov 1991 15:36:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Re: Various

I acknowledge reception. Thanks for bringing all my other work to a crash-
ing halt. Now all I want to do is play control theory withthe other kids.

==========

Date: Mon, 18 Nov 1991 17:53:27 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: language

[from Joel Judd]

Bruce and Bill (911118; 911112),

Conceptualizing language in terms of hierarchical perceptions can be a real
pain. I agree with Bruce's comment:
>There is no reasonable way, Bill, that you can say that control of language is
 simply an artefact of >nonlinguistic perceptual control.
I prefer to substitute the word 'incidental' for 'artefact' in order to get
rid of the leftover or accidental connotations of artefact, and maintain
the intentional nature of the control of perception, which for adult humans
so often involves language.

The bugaboo is society, if I'm understanding correctly, and its role in the
development and maintanance of perceptions. Are you [Bruce] saying that
higher levels CANNOT be perceived without language, or that we do not? You
say:
>Having language, we use it to help integrate and control these higher-level
 perceptions, and most >certainly to help teach them
But we have to experience them to pereive them. How far can LINGUISTIC
experience go? You can talk to me about "love" till you're blue in the
face, but what does it mean to me? Sure, society communicates programs and
principles and concepts, but we still need to act on them somehow. We don't
control language, but rather the perceptions which language reflects
(progressively worse, I feel, as we go higher up the hierarchy). I think
the other reason language is so intertwined with perception is because so
much of our perception relies on interaction with someone else. This is
what Bickhard (1991) is trying to say in the last part of his chapter "How
does the environment affect the person?" (Gary has this and Bickhard was
going to try the net by now)
In talking about "social realities" he differentiates redundancies of

things like inanimate objects (eg. rocks) from "constructed" redundancies like dress and language. Social encounters require "cooperation and coordination." We learn what to do in situations through the scaffolds of "situation conventions" which are made up of common assumptions among participants in a particular situation. Their importance in Bickhard's scheme is explained as follows:

"They permit and guide interaction that would otherwise be pragmatically impossible. They also, however, are a level of reality, of ontology, in themselves. They constitute the organization of scoiety and culture and thought within which the individual moves, acts, developes. Because the individual *is* his or her interacting with the world--his or her resources for such interactions--and because much of that world is constituted at the level of social ontology, much of what constitutes the *person* will be afforded and provided and constrained by the possibilities available in the situation conventions of the society in which the individual develops...a person is *not* a context independent agent who merely acts in social and cultural contexts, but rather is *constituted* in terms of his or her potentialities for acting and interacting, and those potentialities of interaction exist largely in terms of the social and cultural realities in which they participate." (p. 21-22)

Language, to repeat, just happens to be the way human beings past the age of a few months deal with the environment. The fact that perceptual processes underly language can be obscured by our dependence on it, as well as by its seeming "algorithmic" nature in many cases, to use Bickhard's term, can make it difficult to explain or understand. Again, Bickhard:

"All human societies, in fact, provide powerful institutionalized conventions for social reality interactions (he refers to previous papers where he says these conventions constitute language). Language, then, constitutes the primary means for introducing, creating, and transforming social realities. In terms of the opportunities for and constraints on further language interactions--for further conversation--language also constitutes much of the social realities upon which it interacts. Insofar as the person is largely social, it follows that the person is largely constituted at a level of language. Language not only influences the person in particular interactions, language developmentally influences the person by *constituting* the person.

This is similar to the position of the hermeneuticists...It differs, however, in that the ground for the emergence of such social and linguistic realities, and of the persons constiuting and partially constituted by them, remains itself *non-linguistic*. The limits, the horizons, of the world are *not* bounded solely by language. The ground remains that of a general agent interacting with an environment that is not necessarily all linguistically, socially, constituted. Language is largely constitutive, but is not fully constitutive, of either a person or that person's world; language *and* persons are emergent out of a non-linguistic reality." (p.23)

I think Bill's comments (911116) to Wayne apply to language (if I understand them correctly):

>We can also take actions, which are the set of all those experiences that we can influence by act of >will. We find that these actions, themselves experiencable [by the way--is that a word?], affect >other experiencies. We can learn to create some experiences that are not directly subject to acts of >will by varying those experiences that are directly willable; the whole is an act of control. >Through long experience with this kind of act of control, we have found regularities that show how >we must act in order to control many kinds of

experiences. The reasons for these regularities are >not evident in
experience--there is no a priori basis for expecting any particular act to have
any >particular effect on something else.

This sounds like a brief description of language acquisition. The problem
linguists have gotten themselves into is in searching for the "mechanism"
through the regularities Bill is talking about. It doesn't exist, except
through the interactions of language-capable organisms.

==========

Date: Mon, 18 Nov 1991 20:41:51 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: information, metalanguage

[Martin Taylor 911118 20:30]
(Bruce Nevin (911118 0932))

>
>Martin's concept of protocols in his "layered protocol model" concerns
>the role of metalanguage in communication (and in part metadiscourse
>elements in discourse). I don't think he would agree with your
>dispensing with this.
>
If it does, then I don't understand what you mean by "metalanguage". I thought
it meant something like "language about language." The different protocol
layers do nothing like this. They merely transform messages from one
type to another. Of course, it is possible to talk about SOME of the things
that go on in the conduct of a protocol, and I suppose that if one did so
one would be engaging in a form of metadialogue using metalanguage. But
I think that "meta" meta bad end there.

If you mean something else by metalanguage, then I think I have to re-read
a lot of your back postings to revise my ideas as to what you intended me
to understand.

I don't think the protocols deal with "meta" anything. They just process
message contents in an informationally efficient way, so that we can get on
and communicate as best we can about things that often language has no words
for.

"I have to use words when I talk with you" is, up to a point, true. But more
to the point, you let me know what you understood from those words, and I
try to do something else to make our understandings converge, and ... As with
all hierarchic feedback systems, the different levels have different
abstractions, and I think what we usually talk about as "language" occupies
a few of the middle levels.

Martin Taylor

PS More on "information" in a separate posting, coming soon to a Mailer
near you.

==========

Date: Tue, 19 Nov 1991 07:46:53 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Words vs. Signals

[From Bill Powers(911117.0630)]

Bruce Nevin (911116) --

>My experience of this [categorizing screwdrivers] is that (through

>memory and imagination) I access higher levels of control involving the
>perceptions to be categorized. For a screwdriver, I imagine using this
>thing in my hand as a screwdriver or a standiformer.

Surely the perceptions to be categorized are of a lower level than that categorizing process: relationships of object to screwhead, events such as twisting the screwdriver-like thing, configurations of handle, shaft, and blade, and so on. I would propose that in imagining HOW you use this thing, you are imagining at the relationship, event, configuration, and maybe sensation levels. To imagine at higher levels would be to imagine WHY you are using it.

>If I don't know about your use of screwdrivers (as I indeed don't know
>anything about what you do with standiformers) I will leave this one
>out, and I am likely to include an awl and a nut driver in the
>collection that I hand to you on the basis of appearance.

That's my point. You form a trial category and start handing me things that exemplify it. Lacking a previously-associated name for this category, and not even being sure what constitutes "likeness" to the example I showed you, you have to use the nonverbal ability to create and use categories. It will take a lot of trial and error (reorganization) before you create a perceptual category that is usefully like mine, so you can reliably pick out the screwdrivers or standiformers, and so that when you hear the noise "standiformer" you will be reminded of the right perceptual category.

me:
>> We treat
>>items that are actually different as if they were equivalent. Language
>>isn't necessary for doing this. But this ability is necessary for
>>language.

you:
>This recourse to program level and other higher levels through
>imagination and memory vitiates your argument.

I don't see this as involving higher levels. It is the category level itself that creates the perceptual category named "a screwdriver." Given a reference signal that selects this category, the associated control system can generate reference signals for establishing relationships to any number of objects: picking them up, or even just looking at them. Any resulting perception that belongs in the category will produce the same category-perception and thus match the category reference signal.

The reference signal is not a word: it's just a signal in a particular place. A signal may indicate THAT a particular word has been perceived, but the signal is not that word. Neither is the category perception a category: it's just another signal indicating THAT a member of a particular category is present. The process of categorizing can be as simple as the OR operation. Any signal among the set entering the category perceptual function will result in emission of a perceptual signal, indicating that the category is present. Clearly, the OR operation doesn't know whether the incoming signals are word-perceptions or non-word perceptions. It works the same way in any case. (I'm not proposing the OR operation as the key to category formation; this is just one possibility that illustrates the principle of nonverbal perception of categories).

>I am certainly not saying there is an antecedent metalanguage that

>bootstraps language. Something like that is in Generativist claims
>about Universal Grammar against which I have inveighed. "Bootstrapping"
>refers to how language came to be in evolutionary process, and also how
>children acquire language albeit with social support that Cro-Magnon
>humans were bootstrapping their way into.

Somewhere in here is the resolution of our apparent divergences. I am trying to talk about the basic organization of the afferent brain into perceptual functions. You are talking about products of this basic organization, created by a brain or society of brains using these perceptual functions. I have not denied that language is based on social conventions. To say that language is (among other things) a social convention is not, however, to explain language; it is to reveal another capability of the brain. I am trying to understand how it is that a brain can recognize and follow a social convention in the first place.

At a lower level, consider how children learn to use tableware. They learn (in my neck of the woods) to pick up a fork in the right or dominant hand, spear a piece of meat with it, and lift it. This, however is not a description of a function of a nervous system. The description of the abilities that underlie this learned process must involve tendon and muscle-spindle receptors, spinal motoneurons, reference signals from higher systems, muscle dynamics, damping from rate feedback, and so on. My Little Man demo is an attempt to explain how, not why, children learn to manipulate a fork. Once we have a model that can imitate that socially-imposed behavior, we have a model that can also imitate other behavior that requires use of the same underlying organization.

On the other hand, you could say that children handle tableware "because" of social conventions -- that table manners are a social creation. That is true, but it's only a history of table manners, not an explanation.

Language suggests a categorizing capability. When we ask how such a capability might work, we come up with possibilities such as the OR function. When we have on hand a workable model of a categorizing function, whether it uses OR or some much more complex set of operations (about which Martin Taylor can probably tell us more), we will be able to see how this same organization can work with signals that do not stand for linguistic entities. There is nothing that restricts such a function to use with linguistic perceptions -- a signal is a signal.

This same principle applies at every level, low or high. Once we get a hint of the kind of operation required at that level, using language as one possible set of examples, we can then see how that same kind of operation can apply to any incoming neural signals, whether or not they have to do with language.

The main issue in our discussion has nothing to do with examples of good sentences and bad sentences. It is whether enumerating the apparent rules exemplified in such examples -- be they Harris's or Chomsky's -- does not tell us what kind of organization is required for making behavior conform to those or any other rules. It is to confuse a product of the brain with the organization of the brain: an application of a basic ability with that ability itself.

Time to send this off -- I'll try to send responses to the rest of today's avalanche of mail via Roger Peters' link.

Best

Bill P.
==========
Date: Tue, 19 Nov 1991 10:03:38 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: protocols as metalanguage

[From: Bruce Nevin (911120 0723)]

(Martin Taylor 911118 20:30) --

>>(Bruce Nevin (911118 0932))
>>
>>Martin's concept of protocols in his "layered protocol model" concerns
>>the role of metalanguage in communication (and in part metadiscourse
>>elements in discourse). I don't think he would agree with your
>>dispensing with this.
>>
>If it does, then I don't understand what you mean by "metalanguage". I thought
>it meant something like "language about language." The different protocol
>layers do nothing like this. They merely transform messages from one
>type to another.
>
>I don't think the protocols deal with "meta" anything. They just process
>message contents in an informationally efficient way, so that we can get on
>and communicate as best we can about things that often language has no words
>for.

How can they transform a message from one type to another without
referring to the objects and relations in the message?

From your 1988 paper Layered Protocols I: Principles, _Int.J. Man-Machine
Studies_ 28, on p. 182:

> The protocol defines how the content is to be extracted, and may
> also ensure its formal correctness.

(I note only in passing the way you seem to identify a process of
"extracting" information "content" with a process of "transforming" one
"message type" into another. I have commented on this at length in the
past.)

Perhaps one bit of the protocol on one level of syntax (using operator
grammar rather loosely) involves a program that could be stated so:

```
  lookup <string>
  if <string> := {N} then expect {On}
                             or {Onn}
                             or {Onnn}
                             or {Ono}
                             or {Oon}
```

Statements as to what <string> constitutes a word or morpheme, about
morpheme and word classes such as nouns N and operators On, etc., and
about dependencies among these words and classes of words, etc., are
metalanguage statements. Any statement in natural language, including
metalanguage statements, may be abbreviated in one or another kind of
symbolic shorthand for compactness, computational convenience, and other
like pragmatic reasons. The fallacy is to assume:

(a) That such symbolic statements are somehow more primitive than language, rather than being conventional abbreviations from which we read out language and for the interpretation and use of which we require the "background vernacular" of language.

(b) That they constitute a system separate from language.

and therefore

(c) That while they do stand "meta" to language and are therefore technically speaking "metalinguistic," they are not "metalanguage" because they are not in natural language.

Mainstream theorizing about language, as in Generative linguistics and AI, takes the hypostasis or reification of (a) for granted in its presuppositions about the character of "semantic representations," internalized or innate grammar, mentalese, etc.

Another quote, loc cit:

> From the human point of view, a protocol is a learned habit, which functions without conscious thought.

To this I would add that it is a social convention. Continuing:

> Slight variations in protocol can alter the way a message is interpreted; a tilt of the head with a rising intonation can turn a "Yes" from agreement into a sceptical question.

Well, let's look at this one. The rising intonation of most forms of the interrogative is a reduction from performative "I ask" asserted of a disjunction:

I ask whether A or something else.
I ask: A or something else?
I ask: A?
A?

(For detailed discussion of the interrogative, see Harris, _A Grammar of English on Mathematical Principles_ pp. 331-350.)

Notice that the rising intonation appears in the first member of the disjunction "A or B" in the assertion form. (The rise-fall intonation characteristic of questions in many British dialects is similarly found for disjunctions in those same dialects.) A tilt of the head back and forth often accompanies the recitation of a list, so that the single tilt here is a reduction from the series of tilts back and forth, conforming to the reduction of the list to just the first member of the disjunction.

The reduction and the conditions for it are conventions of English--part of its "protocol" for syntax if you will. The identification of operator words and argument words and the specification of their dependencies are also matters of social convention or "protocol". The reductions may obscure the words and their semantic dependencies, especially when words are reduced to affixes or occur in zero form--hence ambiguities and other degeneracies--but they never destroy them or render them unrecoverable.

The metalanguage for doing all this is in the language itself.

>If you mean something else by metalanguage, then I think I have to re-read
>a lot of your back postings to revise my ideas as to what you intended me
>to understand.

Yes.

I believe you thought "metalanguage" meant "language about language" in some discursive sense of "talking about language," rather in the way we have been talking about language here.

Crucially, the metalanguage refers in particular to the objects and relations in language in order to specify how to "do" language. It is the "internalized grammar" as the generativists have termed it.

That is: whatever it is that "does" language, concurrently "does" the metalanguage that enables and controls language, using precisely the same resources in precisely the same ways. For language, the referents are nonverbal perceptions; for the metalanguage, the referents are perceptions of the objects and relations in language. The metalanguage has a smaller and precisely defined, specialized vocabulary; the same may be said of any specialized sublanguage of a science, although the reduction in vocabulary size is not nearly so great. Other than this difference of vocabulary and reference, the kinds of objects and relations required, and the kinds of perceptual control required, are identical.

Your choice of an example was perhaps unfortunate. You often speak of body language and the like when you write about how our account of language must encompass how people *use* language and not just the "skeletal" structure of strictly linguistic information. I must say that this involves more than language.

As I have said before, language is good at transmitting information that is in an important sense public or socially accessible. It is not good at communicating interpersonal relationships, emotions, or relationships to or attitudes toward the information that language is used to transmit. There are means of communicating that do not use language. These gestural systems are partly innate (emotion in the face for example) and partly conventionalizations of things that are innate.

The tilt of the head discussed above with lists and with interrogation is of this latter sort. Another example is nodding up and down for "yes" and sideways for "no." In some cultures these are reversed. In Greek and Turkish, perhaps other Mediterranean cultures, "no" is a lift of the head with a "tsk" click of the tongue tip, often with pursed lips. Another example is pointing with the index finger. There is the hoary story of an anthropologist getting the same word no matter what he pointed at--it turned out to be the word for finger. Many peoples point with their chins.

Manfred Clynes' notion of sentic patterns or contours may be getting at something universal about gestural communication of the former, less conventionalized sort.

We use these means, as I said, especially to communicate things like relationship, attitude, and emotional state. Our reading of them may contribute to a reinterpretation of the strictly verbal information. However, they do so no more than any other contextual clue may do, such

as the presence of rain or the perception that a particular person is preparing to leave the group. If the words speak of compliance but the body language expresses defiance in a boss- subordinate dialogue, the alert boss will interpret the words differently than if the two channels of communication are congruent. But this is not a protocol. It is the reconciliation of conflicting interpretations arising from two concurrent perceptions.

Bruce Nevin
bn@bbn.com

==========

Date: Tue, 19 Nov 1991 12:09:24 EST
From: Joseph Michael Lubin <jmlubin@PHOENIX.PRINCETON.EDU>

[From Joe Lubin 911191.1100]

(To Bill Powers)

I have a student who is interested in building a LittleMan simulation. The project is intended to be primarily a computer graphics (IRIS) project, but he has a control theory course (and you and I to steer him to something that works).

Can you send me the LittleMan code (preferably c, but I used to be able to read pascal too).

Any suggestions that you might have (hints, extensions) would be great.

==========

Date: Tue, 19 Nov 1991 13:56:42 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: constituting categories

[From: Bruce Nevin (911119 1203)]

(Bill Powers(911117.0630)) --

me (911116) --
>>My experience of this [categorizing screwdrivers] is that (through
>>memory and imagination) I access higher levels of control involving the
>>perceptions to be categorized. For a screwdriver, I imagine using this
>>thing in my hand as a screwdriver or a standiformer.

you --
>Surely the perceptions to be categorized are of a lower level than that
>categorizing process: relationships of object to screwhead, events such
>as twisting the screwdriver-like thing, configurations of handle, shaft,
>and blade, and so on. I would propose that in imagining HOW you use this
>thing, you are imagining at the relationship, event, configuration, and
>maybe sensation levels. To imagine at higher levels would be to imagine
>WHY you are using it.

If I see a standard exemplar of a familiar category, recognition is so fast I have no idea what I access. If it is a non-standard case, like your double-L shaped screwdriver, I look at some critical aspect from the point of view of using it for some programmatic purpose: Will it fit in a slot or a phillips head? Can I hold it in such a way that I can use it to apply torque or leverage sufficient to turn a screw, and concurrent pressure sufficient to keep it engaged in the matching concavity in the screw head that I am imagining? And does it seem to be fashioned with this purpose in mind? The latter rules out calling e.g.

a dime or the back of a knife blade a screwdriver, though we have all used such makeshifts to turn a screw at need.

My experience is that I access sequence and program perceptions in order to remember and imagine contexts into which a screwdriver fits. Then I try to imagine replacing the screwdriver with the thing in my hand in that context. Assuredly, these remembered and imagined perceptions of contexts and perceptions of the screwdrivers themselves include perceptions at event and configuration levels, etc. And assuredly all the perceptions being categorized are of a lower level than the categorizing process. But the categorizing process itself cannot be on the category level. The point of view from which I do the categorizing is above the category level. I have to access a perception in which I employ the screwdriver for some programmatic purpose or at least for the effecting of some sequence. I don't simply remember and imagine screwdrivers in vacuo, or in some perceptual generalization of all the sub-category perceptions that have environed all screwdrivers in my experience, as far as my memory can now recall. That would be a vague blur, seems to me. I have to go to a higher level to get an invariant across all those contexts. And the contexts themselves are on the category level--contexts of the category "screwdriver," instantiated or at least instantiable in memory and imagination on levels of perception below the category level, as we become aware if some detail fails to match: "Too thick to fit in any slot that narrow that I ever saw--bet that's a punch of some kind. Yep, look at the way the other end is all peened over." "Metal's too soft--who ever heard of a lead screwdriver?"

This seems to me nicely analogous to the relationship of the reorganization process to the location of conflict, and probably for very good reason. Not knowing the category for something, or not knowing the extension of a category, is a species of conflict.

you --
>>> We treat
>>>items that are actually different as if they were equivalent. Language
>>>isn't necessary for doing this. But this ability is necessary for
>>>language.

me --
>>This recourse to program level and other higher levels through
>>imagination and memory vitiates your argument.

you --
>I don't see this as involving higher levels. It is the category level
>itself that creates the perceptual category named "a screwdriver."

It must be a level above the category level that *creates* a new category, just as it must be a level above the level where conflict is occurring that resolves the conflict. (The highest level perceptions I believe are social inventions that are accepted under innate reference values such as (especially at an early age) those that we associate with things like authority. Hence, no infinite regress of levels. Hence also no "divine authority" for or in the highest levels.)

> Given
>a reference signal that selects this category, the associated control
>system can generate reference signals for establishing relationships to

But we are talking about how we come to be "given" a reference signal that selects this category--how we create it. I agree that once we have

it this is how it appears to work.

>The reference signal is not a word: it's just a signal in a particular
>place. A signal may indicate THAT a particular word has been perceived,
>but the signal is not that word. Neither is the category perception a
>category: it's just another signal indicating THAT a member of a
>particular category is present.

The more we pursue this, the more I come to feel that Harris is right and I was wrong to propose as an opposite extreme (in a range of possibilities) that language is identical with the higher levels of control. In Harris's view, language is a self-contained, socially constituted, self-organizing system. Objects and relations in language correlate only loosely with objects and relations in the world--the world of our perceptions, of course. The correlation is tightest in the specialized sublanguage of a science, whose vocabulary and grammar change as scientific understanding of the science's domain changes. (This is abundantly illustrated in _The Form of Information in Science_ for a subfield of immunology during a period of its history.) As Tom Ryckman says in the Boston Colloquium I sent you, picture theories of meaning and reference have a kernel of truth to them, but are far from adequate. The "language is identical with higher levels of perception" view that I proposed is a picture theory. (Today I sent you a copy of Tom's paper too, Martin.)

Unfortunately, a picture theory is much neater. Without it, it is much more difficult to find the correlation between language as a learned social inheritance and perceptual control whose higher levels seems to depend heavily upon and be strongly guided by what we say about our perceptions.

There is no such correlation of table manners or other social conventions with perception. Put it this way: language has a constitutive role for table manners. Table manners have no constitutive role for language. (Behavior conforming or not to table manners may constitute something, but that is a different matter.)

Whorfian examples of influence of language on perception would support my view that our higher levels of perceptual control depend upon and are guided by the correlation of structures in language with those perceptions. Absence or refutation of such examples would support your view that we have a fully-grown hierarchical control system up to the level of system concepts that does not require language to come to be such and which may happen to learn language given an appropriate social environment.

The color example that I quoted is equivocal. Is it the naming "green" that prejudices speakers of English, or the learned color category, which they just happen to name "green"? In the nature of things, such evidence is notoriously hard to pin down. The fish can talk about eddies and fin presses, but not about water.

Another kind of argument is to claim that it is prima facie improbable that we could develop any but the most rudimentary system concepts, principles, programs, sequences, and categories without language because language is a principle social means for teaching and learning to control these perceptions. Even perceptions below the category level would be in some cases less acute and finely tuned, insofar as a higher-level perception directs attention to certain features of lower-level perceptions as salient to the category, sequence, etc.,

while ignoring others. This role of higher to lower in directing and focussing attention is exactly parallel the business with learning and with reorganization noted above.

But I am not convinced that you are making such a strong stand for separation of language and perception, viz:

(Bill Powers (911112.0700)) --

>The choices presented to
>us in direct experience differ only infinitesimally from each other; in
>fact they lie on a continuum. The artefact that is language creates an
>artificial world in which symbols with discrete values substitute for
>perceptions with continuously-variable values.

In the stronger view I just attributed to you, you would have to say that it is the artefactual world of the category level which substitutes categories with discrete values for lower-level perceptions with continuously-variable values. In fact, should we not say that this substitution is distributed over the hierarchy, some of it taking place with each step up the hierarchy? Just yesterday I was thinking and writing about the role of memory in the step from lower-level perceptions to configuration perceptions, and surely there is something of Peckham's "subsumption with neglect" even there.

(Joel Judd (Mon, 18 Nov 1991 17:53:27 -0600)) --

>Are you [Bruce] saying that
>higher levels CANNOT be perceived without language, or that we do not?

1. Ontogenetic question: I believe we can't acquire the control of higher levels of perception that we in fact have acquired as children growing to adulthood, without language, because language is a principal means of teaching and learning these perceptions and their control.

2. Phylogenetic question: I believe we humans would not have these perceptions from category level on up in anywhere near their present sophistication and complexity as a social heritage without the many millenia of historical development that has in fact gone into them.

3. Real-time perception question: no, I am saying I don't know, and that it seems to be very difficult to determine just what is going on here (ask a fish about water). It is a presumed "yes" answer to this form of your question that I believe Bill is inveighing against. I think that our talking to ourselves interferes with our perceptual control all the time, and I am supported in this by numerous therapists, spiritual teachers, etc., but even if you take them to be impressive authorities, ad verecundiam is no more valid an argument form today than it was in Aristotle's time (despite continued popularity).

Just reading your quotes from Bickhard. I guess I should add him to the list of authorities. I agree with his attention to social reality, and the general character of his understanding of how it is grounded in psychological reality (which in PCT is all we have of physical reality).

Back to work. Lunch is over.

Bruce Nevin
bn@bbn.com

==========
Date: Tue, 19 Nov 1991 15:16:52 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: protocols as metalanguage

[Martin Taylor 911119 14:00]
(Bruce Nevin 911118 0932 -- I think that should be 911119)

Bruce,
Now I understand better what you mean by "metalanguage" and I will accept that under some conditions the protocols do it. But you have another stage of argument to do. When you write:

>Perhaps one bit of the protocol on one level of syntax (using operator
>grammar rather loosely) involves a program that could be stated so:
>
> lookup <string>
> if <string> := {N} then expect {On}
> or {Onn}
> or {Onnn}
> or {Ono}
> or {Oon}
>
>Statements as to what <string> constitutes a word or morpheme, about
>morpheme and word classes such as nouns N and operators On, etc., and
>about dependencies among these words and classes of words, etc., are
>metalanguage statements.
>
I see an engine that interprets the metalanguage statements (your "program") and executes them. The statements could have been any legal metalanguage statements, and the same processing engine could have interpreted them and instructed the protocol coder/decoder what to do as a consequence. I grant the possibility that this happens, but I do not (at present) accept its necessity.

The "engine" that interprets the metalanguage is not itself a metalanguage object, is it? If so, where, if ever, does the recursion stop? I prefer to see it as a nondescript (literally, for once!) entity that accepts certain input data and executes certain ouput behaviour, which may well be dependent on temporal and local context.

I see most of the behaviour of the protocol coders/decoders in the same light (please excuse my continued implicit use of the term "code" for the transformation operation). Much of what they do is analogous to the behaviour of the engine, not of its program. The analogue to the program is the data that is to be transformed. But some of what they do is indeed analogous to the program--we call that the content of the Model in the protocol node--and that part presumably is what you are talking about when you use the term "metalanguage." If so, then I would say that in some protocol nodes it hardly exists, and in others it is quite important.

The non-metalanguage parts of a protocol node's behaviour may execute through the weight structures in a neural network, through hardwired gates of some unknown (to us) logic, or by some other means. But it is not necessary that they be simulated by another nondescript engine operating on a programming language description of their behaviour, which is what I now interpret you to mean by "metalanguage."

Metalanguage is to the protocols as an aircraft cockpit simulator is to a flying aircraft. The behaviour of the simulator is described in all sorts

of rules that alter the displays as a consequence of the pilot's manipulation of the controls and simulated external disturbances. The real aircraft flies as a consequence of the flow of air over its wings (and not because Bernoulli wrote some equations that describe reasonably well what the wing does when the air flows over it.) And in a fly-by-wire aircraft, that flow is affected by describable rules that affect the relation of the control surfaces to the aircraft's attitude and to the pilot's manipulations of the objects in the cockpit. So metalanguage has a part to play in the flight of the aircraft, much like the part I see it playing in the actions of the protocols through which we execute language. But it isn't the whole thing.

In psychology, there is a distinction between "knowing how to do" and "being able to do" sometimes encapsulated in the terms "semantic memory" and "procedural memory." Semantic memory is accessible and can be talked about, whereas procedural memory is only usable. The content of procedural memory can be discovered only by arduous and non-guaranteed analysis (which is presumably why there are so few good teachers who were great athletes before teaching their sport, and why we do not have an agreed grammar of language after 2000 years of analysis, testing, and discussion). I take it that metalanguage is an appropriate metaphor when dealing with semantic memory, but does not apply to procedural memory.

Have I still got it wrong?

Martin

PS. The IJMMS papers were written around 6 years ago, although it is only 3 since they were published. On re-reading, I find that I still agree with most of what they say, but not all, so don't hold me to a literal reading of what I said then. In particular, we had not at that time developed the 9-element model of the protocol node, which incorporates both the semantic and the procedural operation of the protocol, working together.

==========

Date: Tue, 19 Nov 1991 15:30:21 EST
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Words vs. Signals

[Martin Taylor 911119 15:30]
(Bill Powers 911117.0630)
>
>
>The reference signal is not a word: it's just a signal in a particular
>place. A signal may indicate THAT a particular word has been perceived,
>but the signal is not that word. Neither is the category perception a
>category: it's just another signal indicating THAT a member of a
>particular category is present.

Well put. That's exactly the kind of point I find so hard to put across when dealing with the different levels of abstraction in communication that is expressed by the layered protocols. The result of a transformation of a protocol that receives words is not a word string. It is a signal indicating that some concept has probably been received. It is neither the word string nor the concept, though its input may evoke some percept corresponding to the words, and its output may evoke some percept corresponding to the concept. In dealing with one of my colleagues who has been working with me for several years, this point STILL has not sunk in. There seems to be a mind-set that rejects this way of seeing things.

Martin

==========

Date: Wed, 20 Nov 1991 01:10:55 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Misc

[From Mary Powers (911120)]

ANNOUNCEMENT

The 8th annual meeting of The Control Systems Group will be at Fort Lewis College, Durango, Colorado, from Wednesday afternoon, July 29 to Sunday afternoon, August 2.

As in the past, (plenary) sessions will be held mornings and evenings, with afternoons and Sunday morning open for informal discussion, computer demonstrations, reading (see below) and goofing off.

There should be no formal presentations. Any such should be brought to the meeting in written form (30 copies) for distribution to the participants, so that a brief summary presentation can be followed by *informed* questions and discussion. What we want to hear about is what you have been doing with control theory, what you are proposing to do, what problems you are having with it, what ideas you are exploring.

What we want to avoid is a detailed reproduction of your performance before your classes, clients, dissertation committees, or naive-to-control-theory colleagues at the last convention. We want to hear about these things, but - up a level.

The fee for the conference will be about $200, including 1992-93 membership and subscription to the CSG newsletter (including Closed Loop, a publication of threads from the CSG net). Conference attendees, along with spouses, offspring and significant others, are welcome to stay at the college for a few days after the meeting to enjoy a vacation in Durango, which offers beautiful scenery, Anasazi ruins, a narrow-gauge steam train ride, river rafting, mountain biking, and lots of shopping, restaurants and bars. Durango can be reached by air with connections from Denver, Albuquerque, and Phoenix.

This is a mark-your-calendar notice of the conference. Further details and registration forms will be along later on the net and in the newsletter. To receive the newsletter and Closed Loop you must join CSG ($40 per year, students $5). The address is

The Control Systems Group
73 Ridge Place CR 510
Durango CO 81301-8136

(That CR is county road, not circle).

Mary Powers

==========

[From Bill Powers (911119.1000)]

Bruce Nevin, Martin Taylor (911118) --

Bruce says to Martin:

>"Aha!" you might say, "Probabilities at different layers of my Layered
>Protocol model have similar interrelations in a hierarchy." But the
>probabilities are not elements in a control hierarchy. They are

>elements in a descriptive hierarchy.

I agree with this: Bruce has hit on the right objection to certain uses of concepts like probability, variety, information, and so on. They are conclusions in an observer, ways of characterizing a system, but are not descriptions of how the system works. I won't say this objection applies to all uses of "probability," because as Martin has pointed out, a rather simple neural computer (a leaky integrator) can generate a signal whose amplitude is proportional to the numerical result of a probability calculation. I can see that in a really complete neural model, a probability calculation might be required as a way of showing the connection between a signal treated as an average rate of occurance of impulses and the actual details of individual impulse generation and propagation. For example, if a signal is considered to be a function of a set of independently generated trains of impulses in paths reaching a given neuron, computing the signal-to-noise ratio in the neuron's output signal would require dealing with the statistics of the incoming signals (as well as intervening chemical analog processes).

There is however, the question of tractability of a model. I suppose that in principle one could construct an entire control hierarchy out of probabilistic calculations. Such a model might actually come closer to the basic mode of operation of a nervous system. But a model that ignores statistical noise and treats signals as smoothly varying frequencies is far simpler to express, and its behavior is much easier to calculate. Is the added rigor of a probabilistic treatment needed, considering the level at which we can measure and characterize behavior?

Bruce --

>... it seems to me that configuration perception is constructed in
>memory and imagination. An configuration is apt to include aspects that
>are not directly perceptible as intensities and sensations (and
>textures) at the same time as othe such aspects. There must be
>exploration with the senses through time, and there must be the
>construction of a model in memory and imagination.

I agree that memory and imagination supply missing parts of data on which configuration-perception is built. This probably holds true at all higher levels, too. The suggestion is that at least from configuration up, control is model-based (i.e., you control a model, with the model being updated by error signals from the level below, as I proposed last year). Sooner or later somebody has to make a systematic attempt to make model-based control more official, by showing experimentally how it can explain things that the direct CT model can't.

The perceptual function itself doesn't need memory in the sense of recordings of its own past outputs, and of course a whole control system is required for imagination to work. But I don't mean to close any doors.

Hugh Gibbons (911118) --

Welcome aboard. How about putting "[From Hugh Gibbons (YYMMDD.TIME)]" in the first line of your text? Otherwise the only identification of you is "Franklin Pierce Law" in one line of the header that the mailing system transmits. I know what that means but not many others will.

Tickling by someone else is a disturbance of the way your body feels, which you strongly resist. The tickling makes the way your body feels at that position different from the way you want it to feel. When you tickle

yourself, you change the reference-level for the way that part of your body feels, and are in control of the sensations. So you don't resist the change in the sensations: you're creating them.

This doesn't explain why tickling generates such a strong reaction to the disturbance. One suspects an ancient protective mechanism concerning the parts of the body that are most vulnerable. A finger in the ribs is not distinguishable, by the low-level control systems involved, from a fang in the ribs.

Martin Taylor (911118) --

Concerning recognizing an object by feel:

>I don't see what difference it makes to forget your history, or what
>effects newly sighted people have. The point is that active,
>intentional touching provides the perception of an object to a normal
>sighted person, whereas an equivalent set of imposed sensations does
>not.

My history includes learning how the configuration of my hand, as it grasps an object, corresponds with tactile sensations and visual images of the object. In the experiment you describe, the object is touched to the palm of a hand, but the subject isn't allowed (I presume) to curl the fingers to grasp the object during each touch. A normally sighted person has done this with many objects many times, and so knows how to translate the kinesthetic coordinates of the fingers and palm into equivalent visual coordinates. When the hand is held flat, only contact at the palm is involved, so there is no information about kinesthetic space obtained by moving the fingers into contact with the object. If you had to identify an object by actively and intentionally touching it only with the palm of your hand, and not with your fingers, it would take a long time to figure out what the object is by actively touching it from different directions. The map connecting joint-angles at wrist, elbow, and shoulder to visual images is nowhere near as detailed as the map connecting finger positions to hand-sized objects.

Newly-sighted people have not formed the mapping from a set of finger-joint-angle and tactile signals to a set of 3D visual signals. Therefore even if they do grasp the object, they do not sense the tactile-kinesthetic result as self-evidently the same as the visual objects they have just learned to see. This probably shows that in adult sighted people, the mapping is not built-in but is acquired through years of experience. That is why history is important.

To see whether active intention matters, you would have to have the experimenter curl the subject's passive fingers around the object and move them around on it, the same moves that the subject makes when doing the touching intentionally. Is this what was done?

Thus I think that tactile information about objects in no way "confirms" their visual appearance. It's the assumption that the seen object is the same as the felt object that is the criterion for adjusting the maps toward agreement with each other. This is a parimonious assumption, but it's a begged question, not a proof.

>If, indeed, newly sighted people have a problem in identifying visual
>shapes, that's an interesting but so far as I can see irrelevant
>phenomenon.

Perhaps my description was too terse. It isn't just that newly-sighted people have trouble identifying objects. They have trouble saying which object they are allowed to feel is the same as one of the objects they are allowed to see. They can tell you that the visual triangle is not the the same as the visual square, but not that the square object in their (concealed) hands has the same shape as the square in the array of visual objects. What is self-evident to sighted adults is not self-evident to the newly-sighted. Presumably it is not self-evident to infants, either. The shapes are very simple: circles, squares, triangles, cut out of wood.

Joel Judd (911118) --

I appreciate your comments on language, which make me feel a little more confident about mine.

Bickhard's ideas are largely compatible, but his "social ontology" goes too far for me. The reason I worry about epistemology is not to be philosophically correct, but to try to understand the full picture of what a new brain has to deal with as it becomes organized. Whatever we put into the outside world as "real", we relieve the brain of having to construct. Furthermore, if we put something into the outside world, then the problem of building a perceptual function for it becomes a "recognition" problem rather than a "construction" problem. Most research on pattern-recognition takes the view that the shape or thing is really out there, and all the brain has to do is learn to recognize it. But these approaches invariably require a "teacher" who knows what the shape or thing "really is." The teacher tells the recognizer whether it is right or wrong.

But what can play the part of the teacher in a growing brain, particularly before language is available? Perceptions are too complex for a mere right or wrong judgement to be informative. I see the brain's real problem as being one of finding transformations that are mutually consistent -- among different perceptions in the same modality and across modalities. If there is a way of understanding perception on this basis, then there is no need to posit any particular ontological world. Some sort of ontological world is still needed, to provide consistency of dependencies, but we no longer have to pretend that anyone knows what it really is. Each brain makes sense of it as best it can. And "makes sense" is an intended pun.

This puts all brains on an equal footing, removing any privileged points of view. I've always considered that to be an essential aspect of any sort of model-building: the model has to apply to the modeler, and explanations of the modeler must be drawn from within the scope of the model. Otherwise the modeler is cheating.

It isn't necessary to rely on "social ontology" to explain societies or the way individuals interact with them. In each brain there is a conception of the society composed of all the other people who are perceived. That conception determines how the person will interact socially. There are vast differences in the ways people conceive of a society, and in how they interact with the perceived society. The ontological reality of society -- which is simply part of the rest of the universe -- is whatever it is, but any one person's model of it covers only part of it and is related in an unknowable way to the "truth." It should be possible to explain language and other social conventions strictly on the basis of what one brain can know of such things, given the equipment it has. Some recourse to ontology can be had through evolution; perhaps evolution has provided "kits" for forming levels of

perception that are in fact useful in the ontological world. I can accept that, because the brain still would not have any detailed innate knowledge of the current environment. If the construction does not begin completely from scratch, it is still construction, and the criterion for accepting one construction and rejecting another must still be consistency among the behaviors of different perceptions.

If there is any convincing indication of the degree to which perception and reference levels determine our interactions with a society, it is to be found in the way people react when you tell them there is no such thing as "the system" (a subject raised some months ago). Most people reify their concepts of "the system" to the extent of objecting violently to any suggestion that they made the whole thing up themselves. At the slightest provocation they will trot out all the perceptual evidence on which they base their own perception of "the system," and will sometimes go to great lengths to get you to stop saying it doesn't exist or that it's all in their heads. They will try to get you to start believing in "the system" just as they do, and of course they want you to believe in the SAME "system" that they perceive. The fact that hardly any two people will describe "the system" in the same way simply shows that the other people have misperceived "it." Each person has a private conception, which is of course the right conception. Nothing could convince me more thoroughly that there is no ontological "system" -- or that if there is one, it doesn't much resemble anyone's perceptual model of it.

==========

Date: Wed, 20 Nov 1991 04:37:43 -0700
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Categories; Perceptions; c-code

[From Bill Powers (911120.0200)]

My son has changed the personal ID at Boulder to fit the real sender. But that logon is going to last only another month. Graduation day is December 21 or 22! I'll be proud Pop of a mechanical engineer. By then I should be communicating through Fort Lewis College, courtesy of Roger Peters, who is obviously now on the net, and welcome.

Bruce Nevin (911119) --

Category perception: I read your description like this:

>If I see a standard exemplar [Configuration etc}
>of a familiar category [Category]
>recognition [Configuration? Category?]
>is so fast [Transition] I have no idea what I access.
>If it is a non-standard case [Missing Category], like
>your double-L shaped screwdriver [Configuration, imagined],
>I look at [Relationship] some critical aspect from
>the point of view of using it for some programmatic purpose [Program]:
>Will it fit in a slot or a phillips head? [Relationship, imagined]
>Can I hold it [Relationship, imagined] in such a way that
I can use it [Relationship, imagined]
>to apply torque or leverage [Sensation]
>sufficient to turn a screw [Event], and
>concurrent [Relationship]
>pressure [Sensation]
>sufficient to keep it engaged in the matching concavity in the screw
>   head [Relationship]

that I am imagining?
>And does it seem to be fashioned with this purpose in mind? [Indefinite higher level]

Then you say

>But the categorizing process itself cannot be on the category level.
>The point of view from which I do the categorizing is above the category >level.

What the category level of perception is for is categorizing. It receives signals standing for specific lower-level perceptions and if these signals are non-zero in the right combinations, it emits a signal indicating that an element of the category is present. What the "right combinations" are is, for me, a matter of conjecture. But categorizing is done by the perceptual function of a system at the category level.

Setting a reference signal for a category-control system tells that system to perceive presence of this category. The higher system does not do this in order to produce a category: it does this in order to create whatever kind of perception it is concerned with. It doesn't care about its outputs. It cares about a perception that is composed of category and other lower-level signals.

If there is an error (no member of the category currently in perception), the error signal must result in adjustments of reference signals at the relationship level which results in actions at all lower levels that bring a member of the category into present-time perception and thus create the required category signal. If you think that's a vague statement, you're right, but I couldn't make it any vaguer and still say anything. I have never tried to design a category control system. I'm just telling you how I conceive of this sort of system, as far as I have a coherent conception.

When your point of view is IN the category level, you are not aware of categorizing as something you are doing. Instead, you see the world organized into categories. You see a particular object [Configuration] with specific nicks, dents, colors, rust spots, and soldering-iron burns, and see not that object but "a screwdriver." Its category nature pops into view right where the object is, not "in your head." THAT "is" a screwdriver, like any other screwdriver.

When you're operating from a higher level, you can attend "programmatically" (for example) to the process of categorizing. This means that you can perceive that categorizing is going on (although you might not call it that). You can also *simulate* categorizing by running a thought-program: if it has three legs and a flat surface on top, then it's a crippled cow. But that isn't category perception: it's logic. This logic is USING a number of true categories: three, legs, flat, surface, crippled, cow. Those categories are not arrived at by logic, but by categorizing: treating as equivalent all lower-level signals that various categorizing systems pay attention to (receive). The program level is taking some of these categories as inputs and operating on them as a program operates, selecting two simultaneous (intersecting) categories (crippled, cow) that complete the implication. the controlled variable is the implication and its reference-value is TRUE.

The results of categorizing are used by higher systems, but categorizing is not done by higher systems. That's how I think of it.

>I don't simply remember and imagine screwdrivers in vacuo, or in some
>perceptual generalization of all the sub-category perceptions that have
>environed all screwdrivers in my experience, as far as my memory can now
>recall.

Agreed -- we seldom remember or imagine or perceive at lower levels in
vacuo: some higher-level processes are always going on.

I don't think that categorizing is done via memory. It's a perceptual
process that at its simplest could consist of ORing a number of lower-
level signals together. A particular set of signals from various dog-
recognizers PLUS a signal from the perceptual system that perceives the
word "dogs" enter this many-way OR function, any one of them being
capable of yielding a signal in that category-perceiver. This is how
those configuration signals come to be named. If the word-perception is
omitted, it's a nameless category. It's not necessary for the sources of
signals reaching a given category-perceiver to have anything in common
(that way of dealing with categories is really logic, not category
perception), although it's possible. Similarity in some respect may
account for the fact that the signals enter the same category perceiver
(there has to be some reason for perceiving categories at all). On the
other hand, we often judge similarity simply by the fact that we have
categorized perceptions together. But the "respect" may be nothing more
than co-occurrence. A perfectly good category consists of things that are
"mine."

It may help you understand my peculiar interpretations if you remember
that I am basically an engineer: I think of all perceptions as being
generated by a DEVICE, not by an abstract condition.

>It must be a level above the category level that *creates* a new
>category, just as it must be a level above the level where conflict is
>occurring that resolves the conflict.

New perceptions, as the model is presently constituted, are not generated
by higher level systems: they result from reorganization, which is not
part of the hierarchy. The higher-level systems work with whatever
category perceivers exist: they can neither modify them nor add to them.
When the category level is modified or expanded, that is because the
reorganizing system happens to be focused on that level where the raw
materials for computing categories are available. But the reorganizing
system does not act to produce a new category because it senses that one
is needed. It causes random changes because there is error of some kind
-- any kind. The reorganizing system is dumb: it will reorganize whatever
part of the brain it is pointed at. The rate of reorganization depends
only on how much error it is sensing, not on what that error is about,
hierarchywise. When a category perception results that, through being
controlled, eliminates the error that is driving reorganization,
reorganization ceases and a new perceiving function has been added. If a
category does not contribute to lessening of error, reorganization
continues and that temporary new category disappears. Blind variation and
selective retention, the "retention" being by default and the "selection"
being determined by what lessens error.

I'm not trying to tell you I'm right; I'm just following out the logic of
the model and the way it seems to fit with my own experience.

>But we are talking about how we come to be "given" a reference signal
>that selects this category--how we create it. I agree that once we have

>it this is how it appears to work.

My answer to this, I hope, will be anticipated by now. Selecting a category is not creating it. The the ability to perceive a given category must already exist before a higher system can select or use the systemn containing it.

>The more we pursue this, the more I come to feel that Harris is right >and I was wrong to propose as an opposite extreme (in a range of >possibilities) that language is identical to the higher levels of >control.

goody.

> In Harris's view, language is a self-contained, socially constituted, >self-organizing system.

Now you have to see that Harris isn't right, either. Language can't organize itself. It can't DO anything. It's a result, not a cause. It's the product of some kind of operation or set of operations, not the operation doing the producing. Something USES a metalanguage to be about language and to define what "aboutness" is. This is like many cases in which A appears to affect B, but in reality both A and B are the effect of C, a common cause hidden in the background. In this case, C is a device that can do all the operations necessary to produce language or metalanguage, and also to do many other things.

Having had my little gloat, I have to say that we are very much closer together than would be apparent in the last paragraph. You say

>In the stronger view I just attributed to you, you would have to say >that it is the artefactual world of the category level which substitutes >categories with discrete values for lower-level perceptions with >continuously-variable values.

This is what I mean by "the operation." It is the process of substituting a symbol for a class of lower-level perceptions that results in what we call categories. The category perception is itself now a symbol, distinguished from other symbols just by not being the same signal as the signals representing the other symbols. Now the signal can be treated by higher systems -- put into a sequence, dealt with by logical or quantitative programs -- exactly as a symbol is treated.

The important thing here is that you can (and did) describe the essential function of the category level *without referring to language or any other specific use of categories*. Your description is *more general* than a linguistic description. I would say, therefore, that it gets closer to a description of a brain function, and depends much less on describing what the brain happens to do with this function.

> In fact, should we not say that this substitution is distributed over >the hierarchy, some of it taking place with each step up the hierarchy?

"Substitution" happens at every level, but not substitutions of the category kind. To see all relations of one level to a higher level as a kind of categorization is to OCCUPY your own category level and project its perceptions onto the world of lower-level perceptions. If you're in the category level, everything belongs to a category: relationships, events, configurations, sensations, and intensities. Those names, in fact, are names of categories, aren't they? That's why it gets tough to

name the levels ABOVE categories. This model was put together by a human being using the levels that these names purport to indicate. Categorizing the levels is one way to understand them. Another way would be to describe the functions actually taking place at each level -- if we could.

Bruce, thanks for putting Bernie onto my C problem. He gave me a very clear answer, which I will eventually understand (having to do with operator precedence). This will remove the main last impediment to my feeling comfortable in C. I still write C code as if it were Pascal, putting "main()" last and defining everything before it's used, but that's just to avoid reorganizing any more than I have to. Please tell Bruce thanks for taking the trouble.

Martin Taylor (911119)--

>The result of a transformation of a protocol that receives words is not
>a word string. It is a signal indicating that some concept has probably
>been received. It is neither the word string nor the concept, though
>its input may evoke some percept corresponding to the words, and its
>output may evoke some percept corresponding to the concept.

It sounds as though what you mean by "protocol" is very close to what I mean by "perceptual function." We both seem to think of these whatsits as devices that do some operation on their inputs to produce an output that's an element of input to still a higher level. We both seem to think that at EACH level, the outputs are significant, not just hidden layers. When I said that you could probably elaborate on what is needed to make categories (beyond my simple OR), I was thinking that you know more about the potentials of neural nets and might have some more advanced ideas.

Do you have a reprint concerning your layered protocols that I could make sense of?

What I like about this crowd is the willingness to look for convergence and not just defend a position. This is how I think science ought to be done, and seldom is. This way, it's fun.

Joe Lubin (911119) --

I'll send you the code for the first C version of the Little Man. It's not terrifically well commented but you can probably make sense of it. It runs. I will include code for all the modules. This version isn't quite hooked up to the dynamics (you will see a commented-out line containing the function armcont(....) that is the link), but I don't think a student really would want to get into all that. The module (armdynam) that does the dynamics is included anyway, in case he would like to do just the kinesthetic systems.

It would be possible to extend the version-1 model (visual control only, all lower-level systems assumed) to add a layer that generates patterns of pointing, and (skipping some levels) to set up a logical system for

selecting which patterns. Then the little man could answer questions by writing the answers in the air. The entry point is the reference signals for the lateral visual control systems -- follow the logical variable "purpose" which toggles on '!' from the keyboard and makes these reference signals be fixed or varying.

I recommend eliminating all means of moving the target except the keyboard. That is really the most convenient. All references to functions that begin "ad...()" can then be eliminated -- look at the adcontrl module and you'll see what's going on -- selecting all sorts of hardware devices that will be irrelevant.

This code is written in Turbo C 2.0.

Best to all

Bill P.

==========

Date: Wed, 20 Nov 1991 07:41:22 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: transitions and configurations

(Martin Taylor 911118 17:00) --
> they would not have had any training in voluntary
>shifting of the visual focus of attention by eye movement or by internal
>shifts of focus. If that is the case, they might be expected to treat
>visual sensations much like a normal person treats imposed sensations, and
>not make objects out of them.

Since a configuration perception integrates intensity and sensation perceptions, making of them the various parts and aspects of the object, posture, spatial distance, etc., we explore a new configuration by observing contrasts in these lower sensations. We may observe changes (transitions<2>) as it moves or evolves, or we may induce apparent changes (transitions<2>) by moving relative to the configuration, so as to disclose differences (transitions<1>).

Since not all intensity and sensation perceptions of the various parts and aspects of a configuration are typically apparent at the same time, we typically reconstruct those that are not apparent out of memory and imagination. This may be recognition of a familiar configuration (drawers of a bureau), or it may be hypothesis about a less familiar one. Analogy therefore begins here.

Sound right?

Bruce Nevin
bn@bbn.com

==========

Date: Wed, 20 Nov 1991 13:32:26 +0100
From: Oded Maler <Oded.Maler@IRISA.FR>
Subject: A recent paper by Albus

Today I found (and copied but not read although browsed) the following paper:

J.S. Albus, Outline for a Theory of Intelligence, IEEE Trans. on Systems, Man, and Cybernetics, Vol 21, No 3, May/June 91,

pp. 473-509.

The Abstract:

Intelligence is defined as that which produces successful
behavior. intelligence is assumed to result from natural
selection. A model is proposed that integrates knowledge
from research in both natural and artificial systems. The
model sonsists of a hierarchical system architecture wherin:
1) control bandwidth decreases about an order of magnitude at
each higher level, 2) perceptual resolution of spatial and
temporal patterns contracts about an order of magnitude at
each higher level, 3) goals expand in scope and planning
horizons expand in space and time about an order of magnitude
at each higher level, and 4) models of the world and memories
of events expand their range in space and time by about an order
of magnitude at each higher level. At each level, functional
modules perform behavior generation (task decomposition planning
and execution), world modeling, sensory processing, and value
judgment. Sensory feedback control loops are closed at every
level.

To answer some-of-you's immediate question: No. He does not seem
to be aware of the existence of PCT, but he admits that "the
references cited in the bibliography are by no means a comprehensive
review of the subject, or even a set of representative pointers
to the literature."

--Oded

==========

Date: Wed, 20 Nov 1991 08:16:50 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: up the down hierarchy

[From Bill Powers (911120.0200)]

>What the category level of perception is for is categorizing.
>. . . categorizing is
>done by the perceptual function of a system at the category level.

>New perceptions, as the model is presently constituted, are not generated
>by higher level systems: they result from reorganization, which is not
>part of the hierarchy. The higher-level systems work with whatever
>category perceivers exist: they can neither modify them nor add to them.

For the creation of a new category, I see that the reorganization system
provides the point of view exterior to the category level that is
required.

But typically one learns a new category with guidance from another
person, and not simply by raw reorganization in the face of error. The
"affordance" for this guidance is above the category level. I have no
motive (goal, reference signal) for imagining and remembering
relationships like bit-to-slot and the like without recourse to a
program in which I use a screwdriver to fasten something with a screw.
It is because of the lack of such higher-level context that I as
neophyte helper hand you an awl and a chisel in the collection when you
ask for all the screwdrivers. Then you would likely respond (as I have

with my children) "No, look, a screwdriver is something you use to do this with, to turn a screw or a bolt." Oh, now I see, I would say.

For the category-labelling of a new constellation of sub-category perceptions, like the double-L screwdriver, I know that it fits because it includes crucial sub-category perceptions that the category perception "screwdriver" requires, though it lacks many that it calls for as typical but not required. What is the difference between obligatory and optional perceptions on the relationship, event, etc. levels below category? I submit that the difference is that the former are required for perceiving a screwdriver as part of a sequence or program perception, whereas the latter are not so required.

Bruce Nevin
bn@bbn.com

==========

Date: Wed, 20 Nov 1991 08:36:56 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: self-organizing system

[From: Bruce Nevin (911120 0818)]

(Bill Powers (911120.0200)) --

> Language can't
>organize itself. It can't DO anything. It's a result, not a cause. It's
>the product of some kind of operation or set of operations, not the
>operation doing the producing. Something USES a metalanguage to be about
>language and to define what "aboutness" is. This is like many cases in
>which A appears to affect B, but in reality both A and B are the effect
>of C, a common cause hidden in the background. In this case, C is a
>device that can do all the operations necessary to produce language or
>metalanguage, and also to do many other things.

What is missing here is (a) that language is a *social* product and (b) that language is a principal means by which we teach and learn the category, sequence, program, principle, and system-level perceptions that are understood to be normal and "real" by members of our culture.

No individual living control system creates a language as a product of its control of its perceptions (unless you take language acquisition as a special sense of creation).

I think you might say that the arcs and rings in the "crowd" program are self-organizing social products. None of the individual control systems involved sets out to produce a ring. Yet by their all controlling for goals that are interactive, a ring results. So it is, it seems, with language.

Perhaps you are reading too much of agency into the term "self-organizing". A canonical self-organizing system is a pile of sand. A pile of sand doesn't "do" anything either.

Bruce Nevin
bn@bbn.com

==========

Date: Wed, 20 Nov 1991 08:41:05 PST
From: marken@AERO.ORG

Subject: misc, self-organizing system

[From Rick Marken (911120)]

I've stayed out of this discussion of language and category perception because 1) I don't know much about it and 2) it seems pretty diffcult. But life is going pretty smoothly this week so my consciousness is looking for some 'optimum" level of conflict. Fools rush in:

Bruce Nevin (911120 0818) says:

> (a) that language is a *social* product

I agree - language must have been developed to help us control variables by cooperating with other control systems. But now it has lots more uses: self-talk (to support reasoning), poetry and other language arts, and a reified entity that keeps linguits from spending all their time writing radical political tracts.

> and (b)
>that language is a principal means by which we teach and learn the
>category, sequence, program, principle, and system-level perceptions
>that are understood to be normal and "real" by members of our culture.

Debatable. This might be true of people (it may be our principle but certainly not only way) but I think it is not true of pre-linguistic primates. To the extent that they have the machinery that can compute signals representing categories, sequences, etc then chimps and other apes, for example, should be able to learn to perceive and, possibly, control categories, sequences, etc. I think there is behavioral evidence that chimps can control categories, sequences and programs (the Kohler problem solving work). And they do these things even when they have not been taught to talk by psychologists).

>No individual living control system creates a language as a product of
>its control of its perceptions (unless you take language acquisition as
>a special sense of creation).

Probably true. Chimps vocalize to themselves -- but what they are doing probably does not count as language.

>I think you might say that the arcs and rings in the "crowd" program are
>self-organizing social products.

Why use a confusing term like self-organizing to describe a well understood phenomenon? Why not use language to help communicate. The crowd program shows that interacting purposeful systems can produce perceptual effects in observers (the arcs and rings) that are unintended side effects of their efforts to control their own perceptions of proximity to other people. It seems to me that calling this "self organizing" is no more helpful than calling it a "point attractor".

Bill Powers -- did you get the copy of APS symposium proposal that I sent via personal e-mail? Is so, can I take the lack of comments as an indication that it is perfect?

Oded -- Albus is an interesting case. His model is very similar to Bill's structurally. An interesting historical fact: Albus had a series of articles on robotics in BYTE magazine (1979) running at the same time as Bill Powers has his series in the same journal. It would be

a good exercise to read both. I did (then) and, to me, there was just
no comparison. Powers' articles were so obviously of higher quality
it was astounding. The technical difference between Albus and Powers
is clear to me now: Albus doesn't really understand that behavior is
control; thus, his is a model of how systems generate complex
behavioral outputs (even though the loops in his model are closed --
he doesn't understand what this means about the organization of
behavior). Powers, of course, is fundementally (and monumentally)
different since he shows that behavior IS control and is thus
organized around the generation of intended PERCEPTION. Albus' model will
probably get more attention tan Powers' since his basic framework is compatible
with the current zeitgeist in the engineering, behavioral and life
sciences.

Hasta Luego

Rick M.

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Wed, 20 Nov 1991 12:10:13 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: social ontology

[From: Bruce Nevin (911120 1220)]

(Bill Powers (911119.1000)) --

I think I agree with the general thrust of your comments on "social
ontology" as I understand them, viz:

>In each brain there is a
>conception of the society composed of all the other people who are
>perceived. That conception determines how the person will interact
>socially. There are vast differences in the ways people conceive of a
>society, and in how they interact with the perceived society. The
>ontological reality of society -- which is simply part of the rest of the
>universe -- is whatever it is, but any one person's model of it covers
>only part of it and is related in an unknowable way to the "truth."

However, I believe that you must take into account factors such as:

* Individuals control for their conformity to social conventions.

* Individuals overtly teach social conventions by precept.

* Individuals tacitly teach social conventions by example.

* Individuals actively seek to learn what prevailing social
  conventions are.

* Social conventions have a systematic, structural character
  (probably a byproduct of individual perceptual control).

  * Many of these structures are consistent across members of a
    community, and vary in systematically structured ways.

  * Many of these structures are consistent across communities, which
    are thereby more or less closely related, and vary in
    systematically structured ways

  * Many of these structures are consistent across time and vary in
    systematically structured ways, which history shows both the
    phylogenetic divergence of related systems of conventions and their
    convergence by areal and social-stratal interinfluence depending on
    patterns of intercourse over time.

Language provides the paramount examples, but this also applies to
gestural systems (Black American, hands extended out horizontally to
sides while arms are down and close to sides, "caution!", straight out
of West Africa via slavery, non-significant to whites; regional
differences in management of face in e.g. smiling (Birdwhistell)) and
other sets of conventions, such as your example of table manners.

>It
>should be possible to explain language and other social conventions
>strictly on the basis of what one brain can know of such things, given
>the equipment it has.

Yes, but take care not to claim or imply that each brain invents the set
of conventions. There is something there to be perceived, and it
persists through time, space, and social "space".

>. . . the brain still would not have any detailed innate
>knowledge of the current environment. If the construction does not begin
>completely from scratch, it is still construction, and the criterion for
>accepting one construction and rejecting another must still be
>consistency among the behaviors of different perceptions.

I know you are arguing here against knowledge of the environment being
hard-wired in the brain. I am urging that categorical "knowledge" about
(not of) the environment is "soft-wired" in the conventions that the brain
receives as a social heritage from parents, teachers, and peers.
The construction begins from scratch only in the sense that the
construction of the category distinction between blue and green begins
from scratch. What is innate predicts one sort of discrimination of
color chips; what is constructed predicts another. This categorization
is socially acquired and members of a speech community agree with one
another about it much more closely than members of two speech
communities where the categorization differs.

        Bruce Nevin
        bn@bbn.com
==
Date: Wed, 20 Nov 1991 12:09:34 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: the engine and metalanguage

[From: Bruce Nevin (911120 1202)]

(Martin Taylor 911119 14:00) --

>The "engine" that interprets the metalanguage is not itself a metalanguage

>object, is it? If so, where, if ever, does the recursion stop? I prefer
>to see it as a nondescript (literally, for once!) entity that accepts
>certain input data and executes certain ouput behaviour, which may well
>be dependent on temporal and local context.

It is nice in most situations to be able to say that "boss reality" is out there but all we can know of it is perceptions. This present situation may require of us more specific attributions. Yes, the (presumably neurological and biochemical) "engine" that operates the perceptual hierarchy is not itself a perception; and yes, the "engine" that "does" language, and "does" metalanguage in order to do so, is not itself entirely among the perceptions in the control hierarchy that constitute language for an individual language user. There is some hardware there.

From an exterior engineering perspective (which we are not really privileged to have, but we like to think we are), the perceptual hierarchy is all just connections and signals and mechanisms that accept inputs of one hierarchical order and produce outputs of another. Yes, that mechanism exists, presumably. But as we agree all we know of it are our perceptions of our models and theories, and our perceptions of how the former perceptions correlate with our perceptions in general.

The specific elementary control systems arranged hierarchically and in parallel for control of language and for control of the correlations of perceptions constituting language with perceptions constituting the world for us *must* include provision for both the object-language and the metalanguage aspects of language.

A metalanguage control system has no structurally privileged position in the hierarchy. A program-level control system that says "if you say a nonfirst argument before the operator and its other arguments in an interruption, and the interruption comes after words that this fronted material repeats, then you can reduce the fronted material to wh- plus a pronoun" is on the program level and takes as inputs signals from the category and sequence levels, just like any other. The fact that the program, sequences, and categories are metalanguage entities is absolutely immaterial from the exterior engineering perspective that we so much like to attribute to ourselves.

Viewing language as an object in nature, which appears to be a collective human social product, the problem of regress is handled (as I recall) as follows:

> Language contains its metalanguage<1> as a sublanguage, with restricted vocabulary ("word", "morpheme", "prior", "nearby", "same", etc.) and restricted syntax, referring to objects in relations in the language.
>
> The metalanguage<1> contains its metalanguage<2> as a sublanguage, with even more restricted vocabulary and syntax, referring to objects in relations in the metalanguage<1>.
>
> The metalanguage<2> contains its metalanguage<3> as a sublanguage, with even more restricted vocabulary and syntax, referring to objects in relations in the metalanguage<2>.
>
> The metalanguage<3> contains its metalanguage<4> as a sublanguage, with even more restricted vocabulary and syntax, referring to objects in relations in the metalanguage<3>.

The metalanguage<4> contains its metalanguage<5> by reference, using the same vocabulary and syntax to refer to objects in relations in the metalanguage<4>, and so on without need for additional resources. The semantics at this point are virtually vacuous, and all that is needed for a perceptual hierarchy to implement the regress by reference is the capacity to recognize a loop and say in effect "and so forth," which must be required for other perceptual situations.

For more detail about metalanguage<2-4> I could dig out my notes from Harris's 1986 Bampton lectures. His discussion of this point in answer to a question did not make it into the book _Language and Information_.

The same considerations apply to the perceptual control hierarchy for language. There is no requirement for infinite regress.

Bruce

==========

Date: Wed, 20 Nov 1991 18:21:38 +0100
From: Oded Maler <Oded.Maler@IRISA.FR>
Subject: Re: misc, self-organizing system

[From Oded Maler 911120]

(Rick Marken, Ibid.) says:

> Albus doesn't really understand that behavior is
>control; thus, his is a model of how systems generate complex
>behavioral outputs (even though the loops in his model are closed --
>he doesn't understand what this means about the organization of
>behavior).

Rick, you are as predictable as I thought (which doesn't mean you're wrong, of course). Some of the differences in the point of view can be attributed to the fact that Albus is more interested in building systems that achieve some goals which are defined in terms of an observer external to the systems to be built (that is, the designer).

And I can help reminding you, in the context of

> Albus doesn't really understand that behavior is control

that although no commonly-accepted theory of language has been yet arrived at, it is clear that 'is' does not serve as a 1-1 mapping..

Meanwhile I've looked a bit in that paper (only first few pages) and detected several "motherhood"-tautologies as well as some very naive and shallow treatment of evolution, language, emotional states etc. It seems that some other parts of the paper are better, and have a lot in common with PCT, although it is not clear in what sense they extend his previous book (the collection of the Byte articles you mentioned).

--Oded

==========

Date: Wed, 20 Nov 1991 10:11:57 PST
From: marken@AERO.ORG

Subject: Maintenance

[From Rick Marken (911120b)]

Oded Maler (911120) says:

>Rick, you are as predictable as I thought (which doesn't mean
>you're wrong, of course).

Well, heck, Oded, I only get one or two ideas per century -- I've got to make them go a long way.

I just read a wonderfully amusing book by Kurt Vonnegut called "Hocus-pocus". In one part, his character bemoans the fact that americans seem much more eager to discover new technologies than to maintain and develop the ones we've got (the japanese do that for us). It made me realize that maybe that's what I'm doing -- just trying to maintain and develop what I think are the one or two good ideas about how people work. Maintenance -- it makes things boring (predictable) but at my age that's not such a bad thing.

Best Regards

Rick M

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Wed, 20 Nov 1991 14:30:32 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: self-organizing system

[Martin Taylor 911120 1430]
(Bruce Nevin 911120 0818)

>(Bill Powers (911120.0200)) --
>
>> Language can't
>>organize itself. It can't DO anything. It's a result, not a cause. It's
>>the product of some kind of operation or set of operations, not the
>>operation doing the producing. Something USES a metalanguage to be about
>>language and to define what "aboutness" is. This is like many cases in
>>which A appears to affect B, but in reality both A and B are the effect
>>of C, a common cause hidden in the background. In this case, C is a
>>device that can do all the operations necessary to produce language or
>>metalanguage, and also to do many other things.
>
>What is missing here is (a) that language is a *social* product and (b)
>that language is a principal means by which we teach and learn the
>category, sequence, program, principle, and system-level perceptions
>that are understood to be normal and "real" by members of our culture.
>
I think what is also missing is that the word "not" in Bill's line 2 is better read as "and". It seems strange for the guru of feedback systems

to write as if the fact that something is a result precludes it from
being a cause. Of course if A is like B because C affects both A and B
independently, then the relationship is fortuitous. But that's only partly
the case with language. C being the natural environment of things talked
about, and A and B being the ways two people talk about them, then A and B
will share common features no matter whether the people talking A and B
ever communicate. But if they do communicate, then A will affect B and
B will affect A, and the "*social* product" (self-organizing) aspect of
language will dominate the effects of the common environment. In addition,
Bruce's point (b) (the Whorfian hypothesis, to which I subscribe) will
further create a feedback loop in which both A and B affect C, which affects
each of A and B.

(I know that these A B and C are not exactly what Bill meant by the letters,
but the intent is, I think, the same).

Martin Taylor
==========
Date: Wed, 20 Nov 1991 14:44:40 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: transitions and configurations

[Martin Taylor 911120 14:35]
(Bruce Nevin 911120 0740)

>(Martin Taylor 911118 17:00) --
>>   they would not have had any training in voluntary
>>shifting of the visual focus of attention by eye movement or by internal
>>shifts of focus. If that is the case, they might be expected to treat
>>visual sensations much like a normal person treats imposed sensations, and
>>not make objects out of them.
>
>Since a configuration perception integrates intensity and sensation
>perceptions, making of them the various parts and aspects of the object,
>posture, spatial distance, etc., we explore a new configuration by
>observing contrasts in these lower sensations. We may observe changes
>(transitions<2>) as it moves or evolves, or we may induce apparent
>changes (transitions<2>) by moving relative to the configuration, so as
>to disclose differences (transitions<1>).
>
I think there is a very significant difference between (a) observing "changes
(transitions<2>) as it moves or evolves," and inducing "apparent
changes (transitions<2>) by moving relative to the configuration,"
at least insofar as we do not invoke memories of the configuration to
fill in the unobserved relationships among the parts of the configuration,
and to a large extend even when we do incorporate memory.

This thread started as a comment on the infinity of possible ways a deceitful
Nature could provide a passive observer with a given set of sensations,
and how intentional activity on the part of the observer altered both the
range of possible deceit and the subjective effect of the sensations.
Both Bill and Bruce mention perceptual history and imagination, and these
are clearly part of the overall immediate perceptual experience. Nevertheless,
the contrast between imposed and actively sought snesation is important
even when the history exists.

Martin Taylor
==========
Date: Wed, 20 Nov 1991 14:51:49 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>

Subject:	Re: the engine and metalanguage

[Martin Taylor 911120 14:50]
(Bruce Nevin 911120 1202)
>
>
>The specific elementary control systems arranged hierarchically and in
>parallel for control of language and for control of the correlations of
>perceptions constituting language with perceptions constituting the
>world for us *must* include provision for both the object-language and
>the metalanguage aspects of language.
>
So you WOULD say that the aircraft flies because Bernoulli discovered
the equations relating pressure to airflow?

Martin Taylor
==========
Date:	Wed, 20 Nov 1991 15:02:39 EST
From:	"Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject:	category etc. sans language

Rick,

The fact that simians, cetaecians, etc. may have category, sequence,
program, principle, and even system-concept perceptions without
having language does not preclude language being "a principal means"
by which we come to have them, etc. That's why I said "a" means.
Furthermore, I think it is true that their perceptions on these
levels are less developed than those of humans precisely because
they lack language. How to prove that is another matter.
I suppose because it is debatable, and you need language to
debate :-). Anyway I don't mind saying these things again.
I've been saying a number of things repeatedly in different ways
following the good example of our mentor Bill. Seems to work.

	Fusty lumbago

	Bruce Nevin
	bn@bbn.com
==========
Date:	Thu, 21 Nov 1991 09:08:35 EST
From:	"Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject:	clarifications

[From: Bruce Nevin (911121 0730)]

(Rick Marken (911120)) --

>Why use a confusing term like self-organizing to describe a well understood
>phenomenon? Why not use language to help communicate.

I had said, quoting the phrase from Ryckman and from Harris:

> In Harris's view, language is a self-contained, socially constituted,
>self-organizing system.

Bill responded (911120.0200):

>Now you have to see that Harris isn't right, either. Language can't
>organize itself. It can't DO anything. It's a result, not a cause. It's

>the product of some kind of operation or set of operations, not the
>operation doing the producing.

My rejoinder was that the term "self-organizing" does not imply any agency "doing" the organizing--just the opposite in fact. The "self-" part of the word means that the referent falls into what seems to us to be an organized state *of* *itself*, without the intervention of any agent.

Aside from quoting, I used the term precisely in order to facilitate communication with those who are accustomed to applying it to such things as sand dunes and vortices. The point was to draw an analogy between these kinds of organization, and crowd rings, and language, to say that they are all "perceptual effects in observers that are unintended side effects" of interactions of a fundamentally different order. The hexagon-based architecture of a honeycomb or a wasps' nest appears to be another instance, as we discussed some months hence.

(Martin Taylor 911120 14:50)

>So you WOULD say that the aircraft flies because Bernoulli discovered
>the equations relating pressure to airflow?

That sort of thing has been seriously proposed, I know, but it wasn't what I intended to say. Language has a loose correspondence to perception at best, and I did pack too much into a single sentence. Let me try to unpack it a bit and see if it works better.

I said (911120 1202):
>The specific elementary control systems arranged hierarchically and in
>parallel for control of language and for control of the correlations of
>perceptions constituting language with perceptions constituting the
>world for us *must* include provision for both the object-language and
>the metalanguage aspects of language.
>

1. ECSs are arranged hierarchically and in parallel for control of language. The perceptions they control (and outwardly the behavior resulting from the control of those perceptions) constitute language.

2. ECSs are arranged hierarchically and in parallel for control of other perceptions that constitute the world for us. This refers to the fact that all we know of "boss reality" is our perceptions, and does not deny the existence of something being perceived.

3. ECSs are arranged hierarchically and in parallel for control of correlations of the perceptions of [1] that constitute language with the perceptions of [2] that constitute the world for us.

4. The hierarchical and parallel arrangements of ECSs of [1-3] *must* include provision for both the object-language and the metalanguage aspects of language.

Bruce Nevin
bn@bbn.com

==========

Date: Thu, 21 Nov 1991 12:50:12 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: clarifications

[Martin Taylor 911121 13:30]
(Bruce Nevin 911121 0730)

You expand your metalanguage explanation with a 4-part statement that seem intended to form a natural, if not strictly logical sequence, each following from its predecessor. Unfortunately, and I imagine somewhat frustratingly for you, I still cannot see your point 4 as other than an a prior statement of opinion. It is an opinion I do not share, although if you change *must* to "may" I can accept it and agree with it.

The quote in question is:

>4. The hierarchical and parallel arrangements of ECSs of [1-3] *must*
> include provision for both the object-language and the metalanguage
> aspects of language.
>
Each of us, I think, lacks some fundamental background element that the other thinks to be so obvious that it goes without saying. I wonder what it might be?

I totally agree with your comments on self-organization. CSG readers might like to look at recent articles on the cybsys mailing list--an interchange on entropy and self-organization involving Cliff Joslyn, who I believe reads this list. Cliff, if you do, how should people get hold of that series of postings without joining cybsys?

Martin Taylor
==========
Date: Thu, 21 Nov 1991 15:10:49 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: provision for metalanguage

[From: Bruce Nevin (911121 1448)]

(Martin Taylor 911121 13:30) --

>You expand your metalanguage explanation with a 4-part statement that seem
>intended to form a natural, if not strictly logical sequence, each following
>from its predecessor.

No, it was intended to make explicit only the content of that one overburdened sentence because of which you confused me with Bishop Berkeley. The four-part harmonization is not complete and self-contained because the sentence it paraphrases was not. It depends upon prior context because the sentence it paraphrases did. [Query: does "self-contained" imply some agency doing the containing?]

>Unfortunately, and I imagine somewhat frustratingly
>for you, I still cannot see your point 4 as other than an a prior statement
>of opinion. It is an opinion I do not share, although if you change *must*
>to "may" I can accept it and agree with it.

Point 4 is a repetition of a prior conclusion from discussion and argument. The claim is simply that the control system must provide for both "the object-language and the metalanguage aspects of language." And the claim is that it must do so, not merely may.

The simplest demonstration of this is the existence of metalanguage utterances like "this sentence ends with a self-referential assertion in quotation marks." The control system must provide for the control of

metalanguage utterances. Therefore it must provide for metalanguage aspects of language.

Let me try the supporting argumentation a different way, and maybe you can make explicit the basis of your disagreement instead of merely saying that our opinions differ.

The utterance "green" is in the object-language. The control system must provide for the control of this utterance (down to intensities and efforts). It must provide for the control of the correlation of this utterance with the category perception that we call green. And it must provide for the control of that category perception itself. These are all object-language aspects of language that the control system must provide for. We may argue whether the category perception that we call green is a language entity (its meaning) or not. Arbitrary boundaries may have some rewarding aspects, who knows.

The utterance "word" is in the metalanguage. The control system must provide for the control of this utterance. It must provide for the control of the correlation of this utterance with the category perception that we call a word. And it must provide for the control of that category perception itself. These are all metalanguage aspects of language that the control system must provide for.

It seems to me that you want to argue that the category perception that we call "word" (the meaning of the utterance "word") is *not* in the language, hence not in the metalanguage strictu sensu. Assuredly, it bears a metalinguistic relation to things that are indisputably in the language. If so, maybe you can say something about what is rewarding about that boundary.

If your control system includes in itself a program that refers to objects and relations in language, that program bears a metalinguistic relation to those objects and relations. If that program involves a category perception with which we associate the utterance we spell "word," that category perception is in the metalanguage, or at least in the set of meanings of the metalanguage which are in it or not depending on how you feel about arbitrary boundaries as noted above. The program can be uttered in metalanguage statements of the language. (An aside: the fact that this must be so is an important but seldom-recognized constraint on possible grammars.) Such an utterance is unarguably in the metalanguage, and the control system must provide for it.

Therefore, the control system *must* provide for both object-language and metalanguage aspects of language.

If you want to argue that the metalanguage for language is external to and prior to language, you are in numerous company. This view is commonplace in mainstream linguistics, AI, cognitive psych, etc. Please study the arguments of Harris and Ryckman against this view and show how they are wrong. I have tried to summarize these arguments several times, have provided references to Harris, and have sent you the Ryckman paper.

I'm late for my train again.

Bruce Nevin
bn@bbn.com
Date: Fri, 22 Nov 1991 07:54:00 CST
From: TJ0WAH1@NIU.BITNET

Subject: epistemology

[From Wayne Hershberger]

Bill Powers (911116.0800)
>it seems that you're saying that human perception bears
>no verifiable relationship to any universe "behind" or "beyond"
>perception, which is of course the position I have also been
>taking.

Yes.

And that, I believe, is what Gibson was saying as well, or at least trying to say. And Kant too. Your prose is more lucid than most of theirs, but as I read you (Hume, Kant, Gibson, and yourself) you all seem to be motivated by this same epistemological insight. However, the other three believed that this insight also implies that this hypothetical Reality has no empirical basis. So did Plato who claimed that Reality's basis must, therefore, be rational.

You on the other hand seem to accept a Reality appearing to require neither basis. That is, the Reality to which you persistently refer, despite your above remark, appears to be neither an induction nor a deduction but rather and an abduction; you seem to pluck it out of thin air. For instance, while admitting that a perception can not be proven to be a veridical representation of Reality, you are wont to claim that neither can it be proven that perceptions do not approximate Reality to some degree. This begs the question of the Reality itself!

I would encourage you to accept the harsh implications of your own epistemological insight and not backslide, admitting through a back door what you have banished from the front. The challenge, remaining to be addressed, is what exactly is meant by the term real or true, that a perception may be identified as veridical, as opposed to illusory, AND at the same time NOT be regarded as a representation of Reality? That is, what is it that distinguishes a veridical from an illusory perception if NOT the perception's degree of correspondence to some transcendent reality--something you claim to eschew as an arbiter of truth?

In practice, we seem to use a coherence theory of truth. Laymen and scientists alike regard a perception that can not be replicated as illusory. A perception that does not survive the layman's double take is an illusion. An empirical observation that science can not replicate is no fact.

>we find that there can be no basis but observations made by the
>same means by which an "observation" of a photon is made. We
>identify objects by looking at them with our eyes; we verify
>that there is a photon flux by interposing a light-meter (which
>we see) and reading -- with our eyes -- its indicator. So from
>the standpoint of the neural model, the physical observations we
>are using to assign an external source to the visual neural
>signals arrive in the brain by exactly the same means as the
>signals we are trying to explain.

Although this is often the case, it is not necessarily the case. Martin Taylor has already addressed this matter along the lines I

would have pursued. Joe Lubin may also want to respond, so I will not dwell on the issue. Rather, let me address the flip side of this question--to which you also allude:

>a natural universe that exists independently of our perceptions
>of it?

I submit that the ability to register luminous flux with virtually any retina or photomultiplier tube provides the very sort of independence referred to here. Only those perceptions which are demonstrably replicable across observers are objective perceptions, or objects, as we are wont to say, for short. This is the type of independence required of the objects comprising our NATURAL UNIVERSE. When it does not matter who or what makes the observation (i.e., the results are independent of the particular observer) the perception is said to be objective, or to be an object. The natural order is immanent in experience--and not to be confused with some hypothetical Reality that transcends all experience.

>So where do we put this world of direct experience, with all its
>objects and sounds and smells and relationships and people?

In the phenomenal world of time and space--which Kant recognized as intuitions (meta models?).

An article, "Silicon Babies," in the recent issue of Scientific American (December,1991) said something relevant here. Speaking of robots, as Rodney A. Brooks conceives them (which is similar to the way you conceive them), the author of the piece said:

>Subsumption architecture relies largely on the nature of the
>outside world rather than sophisticated reasoning to structure
>the robot's actions. For example, if the robot encounters an
>obstacle, the important thing is to go around it....The robot
>may not need even to remember that the object is there--after
>all, it will detect the obstacle perfectly well the next time it
>approaches it (p. 128).

The expression, "outside world," in this passage obviously refers to the robot's environment. This world outside the robot is not outside the robot's realm of experience. Neither is the robot's world a re-presentation (copy or memory) IN the robot of a world actually transcending its experience such as our model of its environment. The robot merely REGISTERS its environment in its "inimitable way." Call it modeling the environment. Of course, the way a robot registers/models its environment is not actually inimitable; identical robots would register/model in the same way.

Similarly, if your expression, "a natural universe that exists independently of our perceptions of it," refers to something other than the natural order immanent in the psychophysical flux we call experience, it is surely a reference to the perceptual/conceptual models that are constructed out of that flux registering that immanent order.

>So where do we put this world of direct experience, with all its
>objects and sounds and smells and relationships and people?
>It is not represented in either the physical model or the

>(physiological) neural model.

Right you are. The objects of direct experience are not part of either of these scientific models (physics or physiology) because these objects of direct experience are themselves models--empirical if not scientific. They are the layman's perceptual models, analogous to science's conceptual models; both types seem to involve a lot of neural processing, just as your own theoretical/scientific model says. Both types of models are modeling the same natural order. They are twin born of experience. One type of model is not modeling the other type of model. (Only psychological theories such as your model are reflective, modelling the process of modelling itself.) Neither is the basis of the other. The basis of both is the natural order which

>exists in the entire collection of neural signals, each behaving
>in the context of all the others.

Strike the word neural, and I think I could buy it. That is, if this psychophysical flux has any essential characteristic, surely it involves informing and being informed--in a word, signaling. Matter which can neither influence nor be influenced by other matter, doesn't matter. But the signalling does not begin and end in the nervous system.

Warm regards, Wayne

==========

Date: Fri, 22 Nov 1991 09:39:50 EST
From: "CHARLES W. TUCKER" <N050024@UNIVSCVM.BITNET>
Subject: Following the rules and doing the impossible

To all,

The rules of the Southern Sociological Society do not allow the same name to appear on the program for the annual meeting twice in a major "role" thus this entry will appear on next years's program:

"The Myth of Social Control"

by

Control Systems Group Associates

Query: How is this possible?

Best to all,

Chuck

==========

Date: Fri, 22 Nov 1991 07:45:44 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Language and models

[From Bill Powers (911122.0600)]

Bruce Nevin (911121) --

>Since a configuration perception integrates intensity and sensation

>perceptions, making of them the various parts and aspects of the object,
>posture, spatial distance, etc., we explore a new configuration by
>observing contrasts in these lower sensations. We may observe changes
>(transitions<2>) as it moves or evolves, or we may induce apparent
>changes (transitions<2>) by moving relative to the configuration, so as
>to disclose differences (transitions<1>).

Somehow you and I seem to be out of synch by one level, in a peculiar sort of way. Maybe I'm just getting hypersensitive to different nuances of meaning and even my best friends won't tell me. But I wouldn't say it the way you did. This is because I am thinking always in terms of a model, and not just in terms of the manifestations that are to be modeled.

I wouldn't say that a configuration perception integrates intensity and sensation. I would say that a configuration perception results from the integration of intensity and sensation (where "integration" is not to be taken literally, but stands for a process I can't define yet). That's a subtle distinction, but to me, it implies that something beside the configuration perception is doing the integrating: that there's a process being carried out in neurons, no element of the process amounting to configuration perception, but the result being a signal that stands in relation to lower levels of signals as an experienced configuration stands to experienced sensations and intensities.

Now, given lack of a configuration signal, I can see the movitation* for doing those transition<2> movements to bring out different aspects of the sensation/intensity world in the hopes that SOME configuration-sensitive device will respond. There is a lot of anxiety in looking right at something and not being able to make a familiar configuration out of it: who or what is that, sitting in the chair by my bed in the moonlight? OH! It's the clothes I threw there last night! But the heart pounds anyway.

*(I meant to type "motivation", but "movitation" means the same thing and reveals how naive the term is)

I think that some of the difficulty we're having (or at least that I'm having) is illuminated by something you say (later) in a different context:

>But typically one learns a new category with guidance from another
>person, and not simply by raw reorganization in the face of error.

This makes it sound as though another person can somehow guide your reorganizing process so that it will arrive at perception of the category that the teacher wants you to experience. But if you think about how the model is organized, is there any way in which that can be, literally, what happens at any level of perception?

The configuration level is not just a convenient convention that we can indulge in or not. It is part of our inheritance. Simply saying words to a person is not going to help that person form a configuration perception in the slightest degree. Only that person's own reorganizing system can make the required changes at the right level. All that a teacher can do is present a situation, a puzzle. I've read of ethologists/linguists who have run into just this problem. They go through elaborate gestures and mouthings indicating that they want to learn how to do or say something -- and their putative informants nudge each other and giggle. This guy is obviously nuts, Joe, squirming around and waving his arms like that, and singing nonsense. Yeah, Pete, let's do it back to him: Oogie Boogie pah

pah pah!

If the student doesn't catch on that there is something that is required, something the teacher expects to happen, some consequence of not finding out what it is, the student will just look on the teaching as a performance, a TV show. Of course there is, I think, a built-in need to fill in blanks, to make sense of things in the literal meaning. It's easy to "teach" configurations to infants and toddlers because they can't stand having their worlds full of presences that they don't recognize. Teachers often see this search for form happening and take credit for it when it succeeds, but fact they have little to do with it. They just present the problem, and it's up to the child to discover or invent the perceptual mode that will solve it. And this inventing is not a matter of intellect: it's a matter of trying out of lot of changes until one of them works. There is machinery working that is not represented in perception. This is just as true at the category level as at the configuration level.

Martin Taylor (911120) also is out of synch with me, or I with him, on this matter. He says

>It seems strange for the guru of feedback systems to write as if the
>fact that something [language] is a result precludes it from being a
>cause.

The kind of causation I meant to be talking about is called, I think, "efficient" cause. A metalanguage can't do anything to influence language because both the metalanguage and the language are being caused by neural processes, no one of which could be called "language." The processes that enable language OR metalanguage to occur are invisible to us; they are not represented in perception. The input to a perceptual function "causes" its output, but only formally; the efficient cause is the reorganization-created wiring of the function itself, which alone accounts for the fact that the output signal depends on separate input signals in the manner observed. If we are unaware of the intervening mechanism, we might search for some characteristic of the input signals that would explain why they result in just that output signal -- but the search would be in vain. It is not any characteristic of the signals that will provide an explanation. We must understand the function that makes one signal depend on others.

Please correct me if I'm misnaming Aristotle's causes.

When we perceive a configuration, this is not because of contrasts or other properties of the world of sensations and intensities. It is because of the way a configuration-constructing system combines its inputs to produce an output. That is the true explanation of why contrasts, etc. contribute to the perception of configuration. If a different sort of perceptual function is applied to the contrast etc. signals, the result is not a configuration perception but a transition perception -- or perhaps something no human being has ever experienced, but every dolphin has.

At one point Bruce chided me for forgetting that we must, after all, work with a world of intensities, sensations, configurations, and so on, and that all we can really know must be cast in those terms (if I'm remembering the comment more or less accurately). I think that this misses the point of modeling. The CT model is an attempt to account for the existence of the world of intensities, sensations, and so on. It accounts for it not just by looking at the behavior of these entities,

but by proposing something that we can't observe directly: a hierarchy of neural functions that creates some signals out of others. In direct experience we see nothing of our own input functions, comparators, output function, and interconnections. The CT model is a proposition about what underlies all experience and all action. We try to corroborate the model in many ways -- by experimental test of its implications, by looking for anatomical and functional features of the nervous system that might seem to correspond to the elements of the model, and by creating simulations that are meant to reproduce cause-effect relationships at the level of ordinary observation.

I have no quarrel with observations concerning how language comes to have the forms it has; this is clearly the outcome of people interacting with each other, adapting their actions so that, working through others, they have the desired effects (as Martin says). People modify their behaviors to work in an environment where other people are acting, too. And by their deviations from what others expect, they also influence the behavior of other people, leading to others also to modify their behaviors. Language is the result of this sort of interaction. Where the interactions differ, the language comes out different.

But none of these interactions explain language. They tell us how it is, but not why it is. To understand language in the "efficient cause" sense, we must try to find the underlying model that will lead to just those modes of interaction that we observe. That is what CT is about, among other things.

Linguists approach CT through their observations of how language is used. Of course they do. Everyone approaches CT from some field of study in which there is a central phenomenon that seems (and is) important. But there is a tendency on the part of everyone coming into CT to try to make the speciality generalize to cover all cases, and this is asking too much of a specialty. Through a study of language, one might well arrive independently at a hierarchy of concepts such as intensity, sensation, configuration, etc.. But once such a hierarchy is established, the next step is to ask how we perceive and control such things. The "how" question requires a generative model, one that works in terms of neural functions and not words or syntax or verbal rules. When a generative model is found that seems to work properly, one must then ask how the SAME model (there is only one brain) might work in other contexts that are not linguistic. And of course it does work in other contexts.

It is up to each specialist who comes to CT from a specialized background to transcend that background and look for the more general meanings in this theory. The concept of controlling perceptions does not refer to any particular perceptions. The organization of a control system is context-free (just as the SR concept is). That is what makes CT of such fundamental importance regardless of the field within which it is applied. It contains principles that transcend the language in which we attempt to represent them: the principles can be demonstrated by mindless devices as well as by human behavior. To try to reduce control theory to fit any existing body of knowledge is to clip its wings.

Best

Bill P.

==========

Date: Fri, 22 Nov 1991 12:00:55 EST

From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: where language fits in

[From: Bruce Nevin (911122 0722)]

Perhaps it would help if I made explicit the aspect of our work that has been the focus of my attention.

I have distinguished perceptions that are part of language from those that are not. I will use < and > to indicate that the enclosed element is subscripted. Configuration<L> means configuration perceptions that are part of language, and Configuration<W> means those that are not (W somewhat egregiously for "world"). For example, articulator positions, formant configurations, and phonemic segments seem to be among the Configuration<L> perceptions required for language. There are then these two parallel orders of perception, linguistic and nonlinguistic:

Nonlinguistic		Linguistic
System<W>		System<L>
Principle<W>		Principle<L>
Program<W>		Program<L>
Sequence<W>		Sequence<L>
Category<W>	<--__	Category<L>
Relationship<W>	\-->	Relationship<L>
Event<W>		Event<L>
Transition<W>		Transition<L>
Configuration<W>		Configuration<L>
Sensation<W>		Sensation<L>
Intensity<W>		Intensity<L>

These partition the perceptual hierarchy into a linguistic and a nonlinguistic moiety.

The bidirectional arrow that staggers from Relationship<L> to Category<W> represents the mappings from perceptions of utterances to nonlinguistic perceptual categories. It is exactly the same kind of connection as obtains between ECSs (elementary control systems) of any two levels. (I could put in these vertical arrows between levels, but I have left them out for conciseness.) As in the normal case, we are simplifying things by ignoring the possibility of inputting e.g. configuration perceptions to the category level. The main point to be made here is that Category<L> perceptions do not map onto Category<W> perceptions. For example, the perception that "dog" is a word does not itself map onto the category perception that we call dogs. Instead, the word "dog" (controlled at Relationship<L> and lower levels) maps onto the Category<W> perception.

Of course, "word" is part of the metalanguage. Indeed, from the Category<L> level up in the linguistic partition, we are dealing with the metalanguage. To make this explicit, let's revise the tabulation of perceptions as follows:

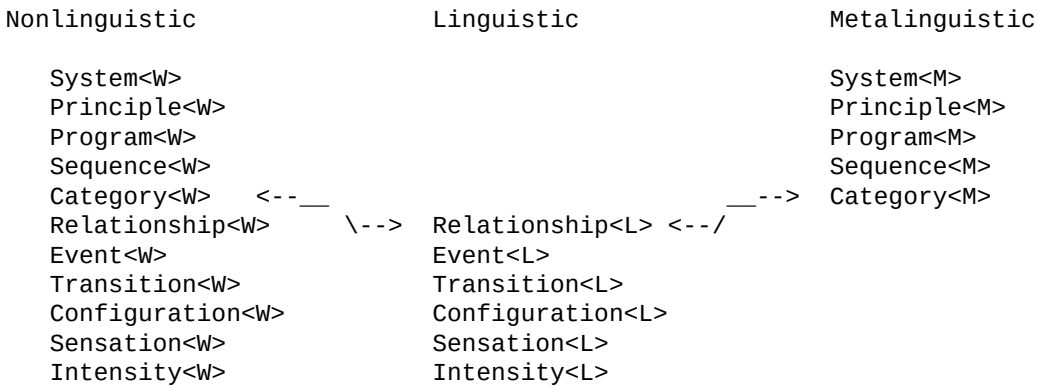

Here, subscript <M> indicates (I hope predictably) metalanguage perceptions.

Now, several of us have made the observation that not all category perceptions map onto utterances, and that we humans may have nonverbalized category perceptions just as higher mammals without language appear to have. So let us distinguish category-level and higher perceptions that do have such a correlation utterances of language--call them Category<L>, etc.--from those that do not, which remain in the nonlinguistic partition as Category<W>, etc.:

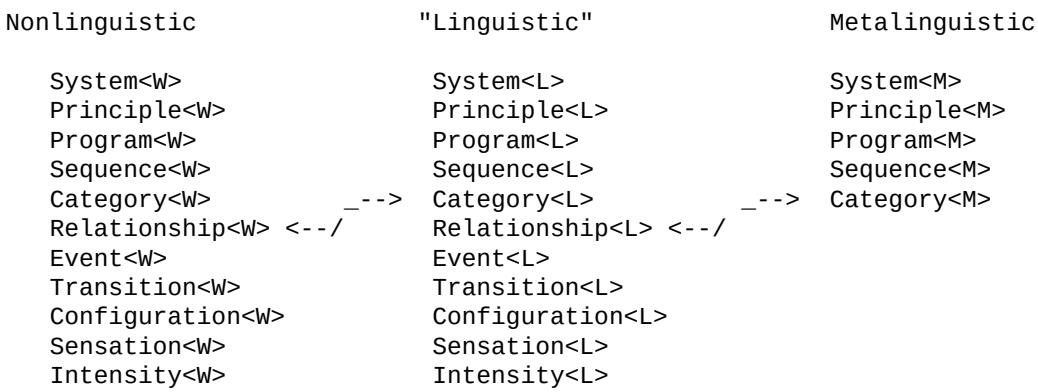

I have now put the heading "Linguistic" in quotes to mollify an anticipated reaction on the part of many that these higher levels of perception that we just moved over there are not "really" part of language. Call cultural or social perceptions together with language, if you prefer.

An important difference between e.g. Category<W> and Category<L>, I believe, is that Category<L> perceptions are conventional, socially learned, and controlled for conformity to social norms. These are characteristics that higher linguistic perceptions share with e.g. the configuration perceptions of linguistic phonology that I listed earlier as examples of Configuration<L> perceptions.

To avoid possible confusion, perhaps I should account for the fact that the bidirectional arrow on the left now appears to be reversed. I made much of the fact that we don't refer to dogs with the Category<L> perception "word". But that is now a Category<M> perception over in the Metalinguistic partition, and the arrow from Relationship<W> to Category<L> is the same one that was previously a tacit, vertical arrow between levels. This arrow on the left reminds us that the connection between Relationship<W> perceptions and category perceptions does not go

away when we move some of the latter from the <W> partition to the <L> partition. They do not go away, that is, just because these perceptions are learned of a piece with language and are the means by which we constitute social reality. Social-reality perceptions<L> are grounded in "boss reality" just as surely as perceptions<W> are.

Just to clinch the point, please note that nonlinguistic perceptions from Category<W> up are perceptions that we don't have words for. Just as soon as we have words for them, they move to the linguistic partition as Category<L> perceptions. And the relationship between utterances and their referents is in the tacit vertical arrow between Relationship<L> and Category<L>.

Now, as I have observed previously, there are metalinguistic utterances (such as the word "previously" above, and such as this parenthetical interruption). Furthermore, as I have noted, we can put our metalinguistic categories, sequences, and programs into words. Accordingly, let us now flesh out the metalinguistic partition of the language partition:

Nonlinguistic		Linguistic		Metalinguistic
System<W>		System<L>		System<M>
Principle<W>		Principle<L>		Principle<M>
Program<W>		Program<L>		Program<M>
Sequence<W>		Sequence<L>		Sequence<M>
Category<W>	_-->	Category<L>	_-->	Category<M>
Relationship<W> <-/		Relationship<L> <--/		Relationship<M>
Event<W>		Event<L>		Event<M>
Transition<W>		Transition<L>		Transition<M>
Configuration<W>		Configuration<L>		Configuration<M>
Sensation<W>		Sensation<L>		Sensation<M>
Intensity<W>		Intensity<L>		Intensity<M>

Now, these are all just perceptions in the perceptual hierarchy, so from a certain global perspective members of the three partitions have no privileged status. Just so, my presently typing fingers are "really" all just interacting concentrations of wave/particle phenomena. Less global partitions have their uses.

I want to dwell on the role of the category level in language and metalanguage. The reason I am particularly concerned with the category level is because my focus of interest is language, and categorization is what we do best with language. As the crossing arrows in the above partitions of the perceptual hierarchy emphasize, the category level is where language connects with non-language.

I certainly agree with you, Bill, that when you focus attention in the category level, everything looks like categories and categorization. When you wear shoes, the whole earth is covered with shoe leather. I do know how to move my attention around in the perceptual hierarchy a bit.

What is important to keep straight, I take it, is the difference between perception as mere observation recorded in memory, and perception resulting in outwardly observable behavior, which has been the particular concern of modelling in PCT. Often the questions or issues I have posed have concerned the former, and responses from you, Rick, and others have been in terms of the latter.

It appears that we can categorize any sort of perception. We can have

categories of system concepts. However, those category perceptions do not (I surmise) provide reference signals to the system concepts. They do not enter into the control of perception that results in changes of behavior.

In fact, I have questioned whether category perceptions enter into the perceptual *control* hierarchy at all. This was the alternative view I mentioned, in which there is no category level at all. On that view, the categorization is done by the mapping from nonlinguistic perceptions to linguistic perceptions. On that view, the mapping *across* from Relationship<L> and lower perceptions involved in the utterance "dog" to the Relationship<W> and lower perceptions of dogs *is* the categorization process.

What then of the primitive residue of Category<W> perceptions, shared with primates, etc.? On this view, I suggest that there can be no categorization without partitioning some perceptions in memory and imagination as conventional representatives of Relationship<W> and lower perceptions--nonverbal symbols. The proposal is to recognize that what we have called category perceptions (symbols, icons, labels to which humans assign words) are really just other <W> perceptions in memory, familiarly and routinely used as associative hooks for current <W> perceptions. Out of this conventionalization developed language.

An anecdote illustrates this:

I was involved in a TQM process called K-J (after the Japanese anthropologist Jiro Kawakita who developed it over a 25-year period to help him grapple with large amounts of heterogenous cultural information). There were five of us from my department. Each of us identified two weaknesses in our work process, stating them succinctly in large block letters on 3x3 yellow stickies (postit<tm>). We refined these so that each one of us was clear what the author meant and agreed that the final version stated that clearly. For each statement, each of us had grounded the statement in remembered and imagined perceptions.

We then had to group these in higher-level categories "intuitively" rather than on the basis of merely verbal connections, and we were to do this without talking about them. Thus, the fact that two slips included the word "specifications" did not mean that they should be grouped. The groupings we came up with were to be based on those remembered and imagined perceptions. Afterward, we put these categories into words.

I draw a couple of different things from this experience.

* The invention of new categories is more dynamic and rapid than the growing of new neural connections, surely, and does not necessarily involve conflict and reorganization, so far as I can tell.

* Our ability to verbalize our categories was so immediate, and so quick of social agreement, as to corroborate my subjective impression that I was inwardly processing verbal categories in order to find commonalities. "These three are about fixing the step description in the EDGE Guidelines."

* At the same time there were kinesthetic and visual configuration perceptions (abstract shapes), on the fringes of my awareness, representing aspects of the problem area and relationships among them. I seemed to "grow" these as something on which to hang key words and phrases; it was from this structure (configuration) that I drew

subsequent verbalizations; and it was by modification of this structure and its labels that I assimilated subsequent discussion and participated in the process of reaching consensus.

The KJ process involves grouping the written labels, writing labels for the groups, grouping again if required, and linking the final categories (groupings) with arrows or "anti-arrows". As this happened, I felt congruity between my abstract configuration perceptions and the graphical representation on the wall. That congruity was the perception of consensus.

It appears that I use imagined configuration perceptions to facilitate the creation of new categories by providing a structure with associative hooks.

The suggestion above was that Configuration<W>, Transition<W>, Event<W>, and Relationship<W> perceptions in memory may become conventionalized hooks for categories of Configuration<W>, Transition<W>, Event<W>, and Relationship<W> perceptions that are evanescent and particular, not conventionalized and universal (be it in real time, memory, or imagination). Higher-level perceptions take signals from these conventionalized surrogates as their perceptual input, rather than from the real-time correlates that have been mapped across to them. And the higher-level perceptions themselves are subject to conventionalization in this world of symbols, which is not a level of perception, the Category level, but rather a partition within the whole perceptual hierarchy.

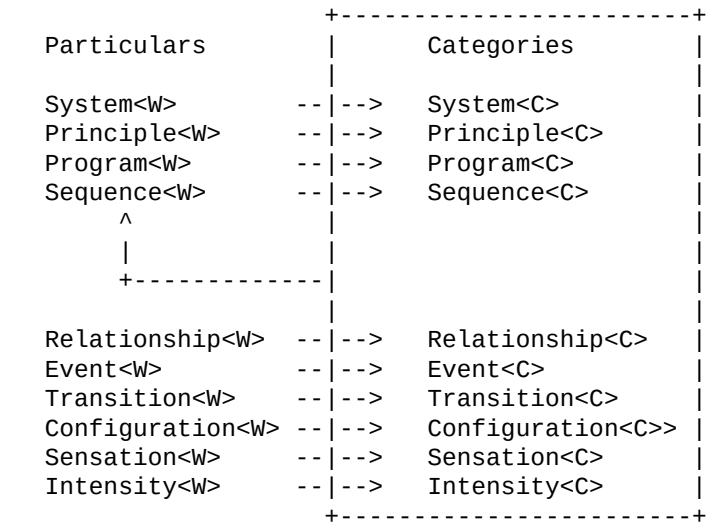

On this view, categories are a different order of perceptions in the same way that language is--as a partition in the perceptual hierarchy, rather than as a level in it.

(I agree about the difficulty finding category perceptions of sensations or intensities. They probably never get conventionalized associative hooks set up in memory because we seldom focus attention there.)

How do the language and metalanguage partitions fit into this view?

There is a mapping from the perceptual stuff of utterances to perceptions in the category partition. Sometimes a word maps onto a category perception: dog. Sometimes the mapping is more complex:

idioms and frozen expressions like "take the bull by the horns," for example, or phrases that function as single unanalyzable words in one sublanguage but not in another. The "stuff of utterances" is controlled in the linguistic partition of the hierarchy.

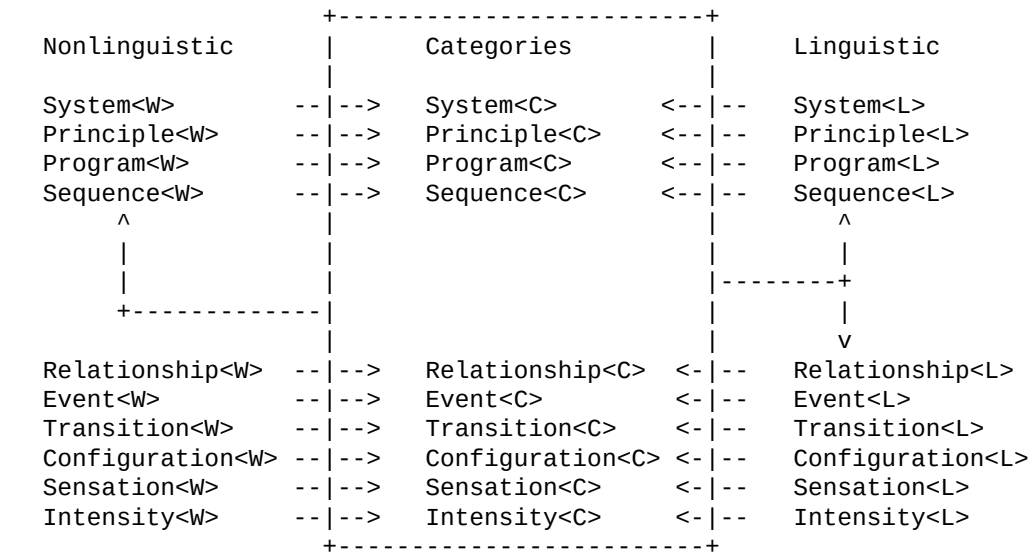

In addition (and differently from the nonlinguistic side), category perceptions of a linguistic sort are input for the Sequence<L> and higher perceptions of grammar (and of linguistics). It is in this loop that metalanguage arises. Some of the category perceptions are metalanguage perceptions.

It is by way of the category partition that we are able to associate utterances, embodying socially learned constraints whose function is the transmission of information, with real and imagined nonlinguistic perceptions.

Bruce Nevin
bn@bbn.com

==========

Date: Fri, 22 Nov 1991 10:07:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Progta ram level and language

[From Bill Powers (911122.0900)]

Don't be confused by the header on this message; I'm sending this via the Fort Lewis connection, using Roger Peters' logon. For the time being, send any direct messages via (powersd@tramp.colorado.edu) as before.

Bruce Nevin (911121) --

Self-organizing: In the CT model, things don't "just naturally" settle into an organized state like a pile of sand. The occasion for reorganization is error of some kind; the process is driven by error, and it actively alters the organization of whatever it operates on. There is indeed an agency doing the reorganizing. It works by blind (but agent-driven and closed-loop) variation and selective retention.

Point 4:

>4. The hierarchical and parallel arrangements of ECSs of [1-3] *must*
>   include provision for both the object-language and the metalanguage
>   aspects of language.

They do. They provide for them in the same way that your computer provides for writing programs with library routines, executive routines, operating systems, interface shells, word-processors, dictionaries, parsers, and the like. The same underlying computer is used for all these levels of operation: it runs programs. Its function is not to run a particular program organized in a particular way, but to run any program that can be expressed in terms of its basic computing operations: add, sub, mul, div, and .... nor, push, pop, indirect, and so on.

I don't mean that the brain's program level is organized like your computer. It's basic operations are not known, although a study of language, or mathematics, or logic, or complex skills might well give us a hint as to what the common underlying operations might be. Whatever its organization, the program level is a general-purpose system that can be used to construct language or any other program-like process. Within this one level there can be subprograms arranged hierarchically, heterarchies, logic trees, anything you can imagine. And nothing that you can't imagine, because this is where you do your imagining about such things.

You say to Martin,

>The simplest demonstration of this is the existence of metalanguage
>utterances like "this sentence ends with a self-referential assertion in
>quotation marks."

Exactly what does "self-referential" mean here (or anywhere)? What is this operation called "referring?" I claim that a sentence or a word does not "refer." It is USED TO REFER. The user is not the sentence, but some process that constructs sentences for a purpose. The process can use a sentence for the purpose of indicating something else, through translations from words to meanings. It can also use a screwdriver for the same purpose, in some instances ("Get me a bigger one of these").

It is perfectly possible for this process to construct a sentence like "This statement is false." There is no paradox there unless you believe that sentences are active agencies in themselves, that can refer to themselves without any outside direction or interpretation. The same process that constructed that sentence can construct -- I demonstrate -- "This statement is not this statement" or "Statement this is false" or "Stmt th fls." There are no rules to prevent any such constructions. Any of them could be assigned meaning. They have no meanings in themselves; they are just strings of symbols. Of course the latter ones follow no conventions, so they would not be very useful for communication.

A last example:

>A program-level control system that says "if you say a
>nonfirst argument before the operator and its other arguments in an
>interruption, and the interruption comes after words that this fronted
>material repeats, then you can reduce the fronted material to wh- plus a
>pronoun" is on the program level and takes as inputs signals from the
>category and sequence levels, just like any other.

Yes. The quoted sentence is a description of a program. It is not in itself a runnable program, but only the disassembled version of it. The actual program must take a symbol like "nonfirst" and translate it into

the operations that test a variable classed as an "argument" to see whether it is preceded by any other variable also classed as an "argument." In a language closer to the machine language, this part of the sentence could be rendered

```
 For each word in the sentence, if isargument(word) is true, use that
 word as the test argument; then call this function:

int nonfirst(testargument,sentence)
{
int i;
 for(i = 0; i < numwords(sentence); ++i)
   if(isargument(sentence[i]) == TRUE &&
      strcmp(word[i],testargument == 0)
       return FALSE;
 return TRUE;
}
```

Not that I think it's done in C, or even this way. But the program must consider each word to see if it is an argument, and then somehow scan the sentence and see if it is the first word that is classed as an argument. The operation that does this scanning and returns the status of the word is the program-level meaning of the term "nonfirst argument." The meaning of "nonfirst argument" at this level is NOT the word that turns out to be the first argument: that would be the configuration-level meaning. The program-level meaning is the implied program-like operation and its result: TRUE or FALSE.

Obviously, functions like numwords(), isargument(), and strcmp() must already exist as subprograms, and those subprograms must in turn exist as machine-language code expressed in the actual primitive operations of the program level. These same primitive operations serve to do all the translating from metalanguage to main language to primitive operations. It is perfectly possible for the metalanguage to use the same symbols as the main language, but this is not *necessary.*

I may have emitted some nonsense here and I don't claim to have worked out all the details of how a program level ought to work in relationship to words. But I hope I've come a little closer to saying why I think that the program level is something more general than language. The primitive operations can produce something structured like natural language, but they can do anything else that is programmable, with any kinds of perceptions whether they are words or not.

It would help to go up a level and look at all this talk about language, and see that it is talk about language. From what viewpoint can one see that?

I still owe Wayne a reply and will get to it. I have to send this off by 10:00 (twelve seconds from now).

Best

Bill

==========

Date: Fri, 22 Nov 1991 13:15:54 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>

Subject: fence-mending, semantic hygeine

[From: Bruce Nevin (911122 1216)]

(Bill Powers (911122.0600)) --

I had said:

>Since a configuration perception integrates intensity and sensation
>perceptions, making of them the various parts and aspects of the object,
>posture, spatial distance, etc., we explore a new configuration by
>observing contrasts in these lower sensations. We may observe changes

You said:

>I wouldn't say that a configuration perception integrates intensity and
>sensation. I would say that a configuration perception results from the
>integration of intensity and sensation . . . . That's a
>subtle distinction, but to me, it implies that something beside the
>configuration perception is doing the integrating: that there's a process
>being carried out in neurons, no element of the process amounting to
>configuration perception, but the result being a signal that stands in
>relation to lower levels of signals as an experienced configuration
>stands to experienced sensations and intensities.

I agree that active voice is clearer and more explicit than either
passive or middle voice, because the latter are ambiguous as to the
agent.

Another example of middle voice: "This is selling like hotcakes."

I should have said "since a configuration ECS integrates intensity and
sensation perceptions," etc.

You had the identical objection to Harris seeming to attribute agency to
language. Again, it was only a seeming attribution and not an actual
one. Very often what we do to make what we say and write less
cumbersome and more fluent has a cost in explicitness and precision.

>*(I meant to type "motivation", but "movitation" means the same thing and
>reveals how naive the term is)

I thought it was an error signal?

I said:

>But typically one learns a new category with guidance from another
>person, and not simply by raw reorganization in the face of error.

You:

>This makes it sound as though another person can somehow guide your
>reorganizing process so that it will arrive at perception of the category
>that the teacher wants you to experience. But if you think about how the
>model is organized, is there any way in which that can be, literally,
>what happens at any level of perception?

>It's easy
>to "teach" configurations to infants and toddlers because they can't
>stand having their worlds full of presences that they don't recognize.

>Teachers often see this search for form happening and take credit for it
>when it succeeds, but fact they have little to do with it. They just
>present the problem, and it's up to the child to discover or invent the
>perceptual mode that will solve it. And this inventing is not a matter of
>intellect: it's a matter of trying out of lot of changes until one of
>them works. There is machinery working that is not represented in
>perception. This is just as true at the category level as at the
>configuration level.

I suppose what is taught is categories of configurations, together with language about those categories, and it is from the categorization that the child gains skill in controlling the particular examples. In this, the child is *not* just stuck with discovering or inventing a new perceptual node that makes sense of lower-level perceptions. As Bruner shows, there is a *lot* of social structuring of the learning situation which plays a crucial role in the learning process.

To Martin, you say:

>A metalanguage can't do anything to influence language
>because both the metalanguage and the language are being caused by neural
>processes, no one of which could be called "language." The processes that
>enable language OR metalanguage to occur are invisible to us; they are
>not represented in perception. The input to a perceptual function
>"causes" its output, but only formally; the efficient cause is the
>reorganization-created wiring of the function itself, which alone
>accounts for the fact that the output signal depends on separate input
>signals in the manner observed. If we are unaware of the intervening
>mechanism, we might search for some characteristic of the input signals
>that would explain why they result in just that output signal -- but the
>search would be in vain. It is not any characteristic of the signals that
>will provide an explanation. We must understand the function that makes
>one signal depend on others.

Thus, a transition perception can't influence an event perception. It's really the output signal of the transition ECS that influences the event ECS. By this standard we should always refer to input and output signals rather than to perceptions. Is this the source of error for you? I certainly have no problem with that, and will try to conform. Language tends to run away on its own tracks, however, having only a loose correlation with perceptions, and an even less direct one with neural currents, so no guarantees of perfect performance.

>At one point Bruce chided me for forgetting that we must, after all, work
>with a world of intensities, sensations, configurations, and so on, and
>that all we can really know must be cast in those terms (if I'm

Gosh, Bill, I can't remember *ever* intending to chide you, and I'm sorry I gave that impression. I remember saying something like this as a caveat against misconstruing something I had said in just-prior context, on the order of "as we all know and agree all we can know of the world is our perceptions of it, but that doesn't mean we deny the existence of `boss reality'."

>there is a tendency on the part of everyone coming into CT to try to make
>the speciality generalize to cover all cases, and this is asking too much
>of a specialty.

I hope my previous long post today shows that I am not guilty of this kind of disciplinary chauvinism. What I am proposing would not be

recognized as linguistics by linguists. That means I have no predisposed audience, either in CSG or in linguistics. But this is my vocation, as what I have jocularly called an interloper--one who lopes between.

Bruce Nevin
bn@bbn.com

==========

Date: Fri, 22 Nov 1991 22:09:47 GMT
Comments: Warning -- original Sender: tag was usenet@UX1.CSO.UIUC.EDU
From: David Schweingruber <dsg1072@UXA.CSO.UIUC.EDU>
Subject: suicide information request

I am a graduate student in Dr. Clark McPhail's social psychology course at the University of Illinois at Urbana-Champaign and am studying various theoretical models dealing with suicide which I may later use during field work among support groups of people who have attempted suicide or who are suicidal.
I am becoming increasingly interested with perception control theory and wonder if anyone is doing clinical, experimental or theoretical work on suicide using PCT. I'd appreciate any information that may be helpful to me or suggestions of where to look.
I thank you in advance.

Dave Schweingruber
dsg1072@uxa.cso.uiuc.edu
(office) 217-333-6515
(home) 217-384-0884
326 Lincoln Hall
University of Illinois at Urbana-Champaign
Urbana, Ill. 61801

==========

Date: Fri, 22 Nov 1991 20:36:52 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: provision for metalanguage

[Martin Taylor 911122 20:20 --time, not vision]
(Bruce Nevin 911121 1448)
>
>
>>You expand your metalanguage explanation with a 4-part statement that seem
>>intended to form a natural, if not strictly logical sequence, each following
>>from its predecessor.
>
>No, it was intended to make explicit only the content of that one
>overburdened sentence because of which you confused me with Bishop
>Berkeley. The four-part harmonization is not complete and self-
>contained because the sentence it paraphrases was not.

I don't remember ever mentioning Berkeley in connection with anything you wrote. All I can remember in connection with him was to point out that Bill was a bit restrictive in saying passive acceptance of photons was the only check we could use on the properties of objects, and Berkeley suggested the cross-check of kicking the object. What else did you have in mind?

>
>Point 4 is a repetition of a prior conclusion from discussion and
>argument. The claim is simply that the control system must provide for
>both "the object-language and the metalanguage aspects of language."
>And the claim is that it must do so, not merely may.
>

Yes, perhaps, but in none of the prior discussion and argument have I seen any -- any at all -- evidence for that "must." Neither is there any in what follows of the posting to which I am now responding. You say, as I paraphrase, that if we want to talk about language, we must use language. OK, I have no problem with that. It's the best we can do. But I find nowhere an argument to suggest that the engine that transforms concepts into language must use a program to do so, though I have agreed that it may do so. I read (possibly wrongly) into your discussion an underlying presumption that there is some kind of general-purpose transformation engine that can be fed with different programs to do different kinds of things, possibly all to do with language, possibly not. Without that presumption, I cannot see your claim as more than a religious assertion, and with it, the claim becomes circular. One needs independent evidence that such an engine exists and is the ONLY means by which language is developed from the complex of conscious and unconscious things that go on in the brain and that can form the overt and implicit content of what we perform as symbolic behaviour. That word "ONLY" is the critical one in this discussion. It is the translation of your "must."

>
>If you want to argue that the metalanguage for language is external to
>and prior to language, you are in numerous company. This view is
>commonplace in mainstream linguistics, AI, cognitive psych, etc. Please
>study the arguments of Harris and Ryckman against this view and show how
>they are wrong. I have tried to summarize these arguments several
>times, have provided references to Harris, and have sent you the Ryckman
>paper.
>
You've totally misinterpreted me. I would argue more strongly against the people with whom you lump me than I would with you. My problem is with the concept of "metalanguage for language" as being a necessary aspect of the performing of language, in one incarnation of the term, or as being adequate to describe language, in another.

There are two issues, really: (1), which I though originally was the issue, but you put me straight a few days ago: Is any symbolic system (including natural language) adequate to describe language behaviour, and (2) Is all language behaviour produced by a general purpose machine that runs some program (that is described in natural language?)? [I apologize for the idiosyncratic punctuation, but I think it helps].

I think the evidence of 2000 years of work provides a weak argument in favour of answering (1) with a "No." The answer to (2) seems devoid of evidence.

Martin Taylor

PS This is really getting a little far from the core of CSG interest, isn't it?
==========

Date: Sat, 23 Nov 1991 07:28:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU

Subject: Epistemology and empiricism

[From Bill Powers (911123)]

Wayne Hershberger (911122) --

You are telling me now that Hume, Kant, Gibson, and I all seem to be "motivated by the same epistemological insight" -- yet:

>... the other three believed that this insight also implies that
>this hypothetical Reality has no empirical basis.

They may not have had the same attitude toward empiricism that I have. For instance, they may have been of the opinion that only empirical facts can be true and real. This would put them in good company, but it would be the company of those who customarily elevate statistical preponderances to universal certainties. I have met very few empirical facts that did not contain easily discernible uncertainty; certainty is achieved by ignoring the actual data and plunking a dot down in the middle of the scatter: that's the REAL value. Most of the time, there is a background of approximations, arbitrary assumptions, and interpretations without which empirical data would have no meaning. These assumptions, and the fact that someone is interpreting, are not mentioned in polite company.

So to say that a hypothetical reality has no empirical basis is not the indictment it might be if there were such a thing as pure empiricism untainted by human imagination and interpretation.

Even when we confine our observations to the omnipresent psychophysical flux, we see things that are contradictory. Our judgments of width and height do not agree with readings from calipers. Our judgments of straightness do not agree with straight-edges. Our judgments of relative temperature do not agree with thermometers. Our judgments of relative brightness do not agree with photomoters, and our judgment of relative color and especially color composition do not agree with spectrographs. In realms of more complex observations, we do not agree with each other about palatability, difficulty, comprehensibility, spelling, grammar, or miracles. We observe nonexistent phantom arrays created by what we know to be a single stationary flashing light. We don't even agree on a color like "green." When we do agree, the spectrograph can tell us we are looking at different colors; when we disagree, the spectrograph can tell us we are looking at the same color.

You say "An empirical observation that science can not replicate is no fact." Is there ANY empirical observation that science can, literally, "replicate?" Replication never in fact occurs: perhaps that is the ONLY replicable fact. What happens is that we make a series of meter readings that disagree with each other, and then we say that there is a Real value that lies somewhere within the range of the readings. We can't actually "replicate" a reading (in fact, if we get exactly the same reading twice in a row, we tap the meter from then on). We can't even replicate the scatter in a series of readings. We replace the scattered, variable, inexplicable individual observations with an idealization that we conceive of as the real observation. In doing this, we create precisely the reality I am talking about: a reality that we do NOT observe, but accept because it makes sense of experience. Empiricism itself leads to acceptance of a reality that underlies observation but is not the same as what we observe.

Much of what seems to be replication is a product of the human capacity for categorization. We can make observations that vary widely, yet make them appear to be the same by classing them together. If we ask 100 people, "Are you in favor of abortion?" and 60 percent of them say "no," we lump the 60 "no" answers together and say that they indicate the same opinion about abortion. In fact we don't know what question the respondents were actually answering: all we know is the question we heard ourselves asking, and what it implies to ourselves. Some of them were thinking "No, not even to save the life of the mother," while others were thinking "If it's necessary to save the life of the mother, but in general, no." Those are both "no" answers, aren't they? "No" is "No", isn't it?.

Empiricism fails as soon as you go beyond a description of a snapshot of the psychophysical flux. Take as simple a thing as a lever with the fulcrum in the middle. Pushing down on one end, you observe that the other end goes up, while the lever itself tilts. You can easily satisfy yourself that if you do push one end down, the other end will rise, and if you don't push that end down, the other end will stay where it is. There is little more to be determined, empirically, about the behavior of this lever. Now, is there any connection between your pushing on one end of the lever and the subsequent tilt of the lever and the rise of the other end? Have you given an adequate account of the lever by reporting just the facts of what happens in the psychophysical flux?

Obviously not. You have reported three facts: pushing down, tilting, and rising. Is the first fact directly influencing the other two facts? Of course not: at least one physical property of the lever, its rigidity, is needed in order for the one fact to lead to the others. Facts do not influence other facts just by existing. Given two identical-looking levers, one may behave as you expect while the other simply bends. The difference that explains the difference in behavior is not among the empirically-observed facts. It is an imagined property of the lever, deduced from its observed behavior. That property is part of the unseen reality of which I speak. Does this "rigidity" actually exist? It might. But it might also be a consequence of unseen factors such as intermolecular forces, none of which itself is "rigidity," but which together have the consequence of imparting rigidity to the lever.

The "physical" part of the "psychophysical flow" is imagined or deduced (it is deduced, then imagined). It is not observed. There is no meter that measures rigidity, or whatever factors give rise to it. And if there were a meter measuring the straightness of the lever, it would not be measuring the causes of rigidity.

Or consider another case that tells us even more about our relationship to reality: a wall switch that operates a light in the ceiling. We can easily determine that when the switch is up, the light is on, and when it is down the light is off. We can manipulate the switch and reliably observe -- for a while -- that the state of the light obediently changes. If all we care about is operating the light, we are finished.

But some of us assume that there is some connection between the switch and the light. We can't observe this connection without destroying the wall and ceiling, yet we have little doubt that there is one. This imagined connection has no basis in our empirical observation of THIS switch and THIS light, yet as a matter of faith we accept the existence of the connection. We do not accept effects at a distance, in most cases.

Does it now shake our faith if we flip the switch up and the light fails

to come on? Not at all: we deduce that the switch has failed, the bulb has burned out, or there has been a power failure. We don't observe those explanations empirically, either -- although we would like to check them out by some indirect means like turning on a floor-lamp to eliminate the general power-failure explanation. We don't actually need to visit the power plant.

We begin to suspect the switch when we wiggle it up and down and observe that the light comes on again. But now we observe an odd thing: the light now comes on when the switch is DOWN rather than up. We can still toggle the state of the light by moving the switch, but the relationship has reversed. Have the innards of the switch suddenly turned upside-down? That seems ridiculous.

Then we remember that high-school puzzle, the two-way switch. There is, we realize, ANOTHER SWITCH somewhere else that also controls this light. There is someone fiddling with that other switch!

Is this an empirical observation? No, it is a memory-based guess about a hidden reality. Can we be certain that there is another person fiddling with the other switch? No, how could we be? We haven't seen the other person or the other switch yet, and if we have to catch a plane we may never do so. It's possible that the spring in the other switch broke and let the switch flip down, with nobody operating it. It's possible that there's a relay in the circuit that short-circuited. Yet there is value even in our wrong guesses, because they are possible explanations and in other circumstances might be the correct ones. These possibilities relate not to the empirical world, the psychophysical flux, but to a world beyond what are are sensing: inside the wall, in another room we haven't visited, in a power plant we have never seen.

Most of the world within which an individual human being makes empirical observations is outside the scope of that person's perceptions, yet its imagined state forms the context within which what is observed is interpreted.

There are some aspects of the hidden reality that we strongly suspect to exist, but which we will never be able to verify. Is there really an electromagnetic flux propagating through empty space, a flux that we call "light?" There is no way to check this. We can only say that when some sort of receptor STOPS the imagined propagation, we get some sort of meter reading. There is simply no way to detect light in flight. Human reason screams at us that OF COURSE light has to propagate through space in order to reach our detectors -- but that is not and never will be an empirical fact.

Nearly all of our meter readings interfere to some extent with what is being measured; meter readings lie. An ordinary volt-ohm meter draws current when it measures voltage. That current causes the actual voltage to drop a little, and sometimes a lot. In electronics, we learn to measure the meter's resistance and calculate that of the circuit, and correct the reading to the "true" value. So the meter reading we see has to be corrected to indicate the voltage that really exists.

In order to estimate how hard a suitcase is pressing down on the rug, you have to lift it. Data from polls has to be corrected to show what the true opinions would have been if everyone had been telling the truth. When we bargain with others, we try to estimate from the offers the other is making how much that person is really willing to pay. When we see a car in the convex right-hand outside rear-view mirror, we see a label

saying "Objects appear farther than they are." When a PhD candidate fumbles a question, we make allowances for her nervousness. When an agent says "I'll give you a call if anything comes up," the actor ceases to expect a call.

We are always making adjustments to observations, denying the validity of empirical data, to bring our actions and expectations more in line with a world that underlies appearances. We are better off doing so than not doing so, even though we are sometimes mistaken in not taking literally what is before our eyes. Sometimes the phone rings and it is the agent telling the actor to report for rehearsals.

There are really no justifications for denying the existence of a reality that is different from the one we experience, even the one we experience through the use of scientific instruments. The scientific instruments themselves shout at us that there is something going on that is invisible to us. If we were strict empiricists, we would report analog meter readings in radians, not volts or pH or counts per second or RPM or pounds per square inch or quarks per cubic meter. But we do not: doing so would leave us with a world that made no sense. We have constructed an elaborate network of imagined entities and relationships that purport to live in the world on the other side of the meter readings. While this conceptual world may miss the mark, and may describe only a projection of a much larger space onto the dimensions to which the meters are sensitive, it may be correct in some respects, particularly respects having to do with derivative notions like conservation of energy or control. We will never know, of course; our meters and our sensors stand between the observer and the reality. If information is coming in to us through these channels, we still don't know what it is about. The incoming information carries no identifying labels.

Still, it pays to guess, as long as we are alert for the evidence that says we should change the guess. But today's empiricism is tomorrow's illusion, and often today's.

Over,

Bill P.

==========

Date: Sat, 23 Nov 1991 12:42:48 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: categorization

[from Joel Judd]

Last night my 2:10 year old son was in the living room as Back to the Future III started. At the beginning, M.J. Fox hides the time machine in a cave, then discovers a big black bear lives there. I asked my son, "What's that?" "I don't know," he says. [the bear rears up and roars] "What's that?" I press. "Dunno," he says again. [the bear chases Fox out of the cave into the desert. "What is it?" I ask again, as the bear moves after Fox. Finally he says, "A big vaca [spanish for 'cow']." Now he's seen a bear before in books and in real life, most recently at the St. Louis zoo in August. However, we're surrounded by farms, and we almost always pass cows or sheep going somewhere, though I can't say there are many black cows, and I've NEVER seen one stand up and roar.

Well, couldn't resist sharing part of growing up [bilingually] with y'all.

==========

Date: Sun, 24 Nov 1991 07:02:00 CST

From: TJ0WAH1@NIU.BITNET
Subject: epistemology

[From Wayne Hershberger]

Bill Powers (911123)
You use the word reality very differently in the following two passages. The first usage refers to a created reality; the other refers to a hidden reality (my underlining) which you say is unverifiable. I find the former concept very useful, but not the latter. The latter is essentially a contradiction of terms, a paradox. This paradox is readily resolved, however, using your own very cogent arguments. That is, if the words "created reality" used in the first passage below are substituted for the words "hidden reality" in the second passage below, what you then say makes sense, don't you think?

Passage #1:
>We replace the scattered, variable, inexplicable individual
>observations with an idealization that we conceive of as the
>real observation. In doing this, we create precisely the reality
>I am talking about: a reality that we do NOT observe, but accept
>because it makes sense of experience. Empiricism itself leads to
>acceptance of a reality that underlies observation but is not
>the same as what we observe.

Passage #2
>Most of the world within which an individual human being makes
>empirical observations is outside the scope of that person's
>perceptions, yet its imagined state forms the context within
>which what is observed is interpreted. There are some aspects of
>the ___hidden reality___that we strongly suspect to exist, but
>which we will never be able to verify. Is there really an
>electromagnetic flux propagating through empty space, a flux
>that we call "light?"

There is not just one created reality.

There are many created realities, some comprising perceptual "objects" (e. g., a laser beam) some comprising conceptual stuff (electromagnetic flux), both dealing with the same phenomena. When we attempt to order these created realities by their truth value, we find ourselves using parsimony and replicability as our criteria.

However, I am not sure that the conceptual stuff is any more or less empirical than the perceptual stuff. The mean of a set of data is a datum too. Further, although the arithmetic mean is not the only measure of central tendency which could be defined upon the raw data, it is very precisely constrained by its raw data.

It seems to me that our created realities transcend the raw psychophysical flux in essentially the same way that a mean transcends its data; that is, they are precisely constrained by the psychophysical flux--baring miscalculation. The input functions in your psychomodular control-system model create perceptual realities in this constrained way. That is why your

theory is not solipsistic.

Another matter:
>Empiricism fails as soon as you go beyond a description of a
>snapshot of the psychophysical flux.

On the contrary, empiricism fails when you artificially restrict
it to snapshots of the psychophysical flux.

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu

==========

Date: Sun, 24 Nov 1991 11:36:24 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Categories

[From Bill Powers (911124.0800)]

Bruce Nevin(911122) --

Bruce, you've taken the ball and are running well with it, but in a direction that I don't want to go. Your treatment of linguistics is becoming more and more complex, instead of reducing to fewer and simpler principles. Maybe

(1) Your understanding of language is far more complete and much deeper than mine, and you have begun to grasp a structuring that is actually simple but only looks complex to me, as I don't get what it is and therefore see only its complex product, or

(2) You are so focussed in the category level that you are interpreting everything as categories (partitionings, groupings, classifications, taxonomies), or

(3) some of both apply (perhaps the best choice).

My evidence for (2) is what happens to your diagram when it evolves into this:

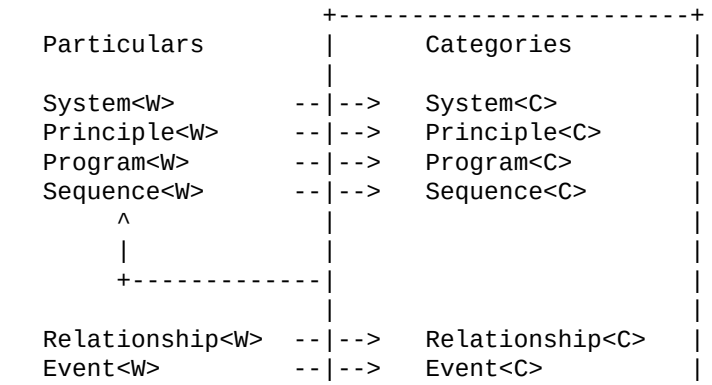

```
   Transition<W>    --|-->   Transition<C>     |
   Configuration<W> --|-->   Configuration<C>> |
   Sensation<W>     --|-->   Sensation<C>      |
   Intensity<W>     --|-->   Intensity<C>      |
                      +------------------------+
```

In effect, categories themselves have disappeared from the hierarchy. My proposed law of awareness goes like this: you are never aware OF the level that you're aware FROM. Instead, the perceptual mode of that level appears to be the main feature of the world you're examining.

I would like to propose a simpler picture, or at least a step toward one. Let's just consider the configuration level and the category level, to avoid clutter.

Category perceivers:

```
       CV1               CV2        CW0    CW1       CW2       CW3
```

Configuration perceivers:

```
                                                 ---       /\
  "noun" "word"  "square" "triangle" "line"    |   |     /__\   -----
                                                 ---
```

Quoted terms are meant to be perceptual signals indicating presence of the quoted heard or seen events/configurations of sounds. They are not classes. "Noun", for example, is perception of the spoken sound or seen letter-group.

The three figures are meant to be perceptual signals indicating presence of the respective patterns on the retina.

It's hard to draw lines this way, so I'll just list the inputs to the various configuration-perceivers:

CV1 (category-verbal-1) generates a signal when any of this set of configuration signals is present:
       {"noun" "word" "square" "triangle" "line"}

CV2 generates a signal for any of the set
      {"word" "square" "triangle" "line"}

CW0 (category-world-0) generates a signal when any of the set
     { square-perception triangle-perception line-perception }
      is present.

CW1 generates a signal when either or both of
      { square-perception "square"} is present.

CW2 generates a signal for either or both of

      { triangle-perception "triangle" }

CW3 generates a signal for either or both of

     { line-perception   "line" }

Now we can see that the following sentences indicate the same request:

1. Hand me a triangle


2. Hand me a

For each of the geometric figures, there is a large set that would be interchangeable, so most of the inputs to CW2, for example, would come from systems that perceive three-cornered figures. Only one member would be a sound or a letter-group -- the "name of" the category. However, any element of the category can evoke the category signal -- i.e., "stand for" the category. "word" can stand for {"noun" "word" "square" "triangle" "line"}. So can "noun" or "square". To narrow meanings down, higher systems specify more than one category at the same time. Then you get the lower-level set that exemplifies them all.

Given the reference-condition CW2* = nonzero, the associated category control system could produce the perceptual signal CW2 by getting lower-order systems to produce either "triangle" or /_\. If CV1 is to be satisfied at the same time, only "triangle" will work. If the logical condition "CW0 AND CW3" is to be satisfied, only ----- will work.

Note that there is no difference at all between the "verbal" and the "world" categories: they both receive configuration input signals and output a signal standing for presence of a category rather than one specific configuration. The kind of machine-language operation required at the category level is the same in either case.

Now consider the next level up, sequence. Suppose the sequence presented at the lower levels is "triangle", <image of square>, "line". From this we get (at least) three potential sequence perceptions:

S1 = CV1, CW0, CV1

S2 = CV2, CW0, CV2

S3 = CW1, CW2, CW3

and so on.

If any one of these sequences is perceived, of course the categories and configurations involved are also perceived at the same time, at a lower level. We experience all these levels at once, projected into the common space of experience. There are no labels attached saying which aspect of the experience goes with which level. At the category level, "line" and ---- (either yielding CW3) are identical; at the configuration level they are not. So in one sense a line and the name for it, "line", are equivalent. But we are also aware that they are different.

At the sequence level, we perceive S1, S2, or S3, not the corresponding triples of category signals. If the category signals are perceived as CV1, CW1, CV1, we perceive S1 -- the signal path carrying S1 carries a non-zero signal. That is all that happens at the sequence level: there is no indication of the categories as individual occurrances (although each category can be perceived, as it occurs, at the category level).

The program level then works with S1, S2, and S3. It would probably figure out that S1 and S2 are equivalent: if S1 then S2, and if S2 then S1. What these sequence-signals mean -- whether they are verbal or non-

verbal -- at lower levels is of no concern to the program level. A given program will operate in the same way whatever the source of the input signals.

Each new level of perception adds ONLY the quality or aspect of the world with which that level is concerned. It does not also carry along the lower level perceptions.

When we try to talk about levels above categories (say, principles), the higher level doing this must select words that evoke configuration signals so as to evoke category signals so as to evoke sequence signals so as to evoke program operations so as to evoke the principle perception in question. It must do this by "delegating responsibility" to the program level, which delegates it to the sequence level ... and so on. The principle level does not deal with word-configurations or non-word configurations, although perception of such configurations is certainly entailed, at the configuration level, in talking or writing and in the multileveled process of recognizing a principle.

When we try to live so that certain principles are evident in our own behavior and its effects, the same operations are applied, ultimately creating low-level perceptions that are interpreted level by level until a principle is detected -- adjustments are made until it matches the reference-principle. This can involve words and non-words, but can involve nonwords exclusively. What we term "courage" is probably made more of feelings than of words, at least in those who have it.

What HCT does, or proposes to do, is to parse the unitary-seeming world of experience into types of perceptions that can, with due attention, be discriminated from each other and which can be seen to be functionally dependent in a hierarchical fashion. It separates the thingness of words from the categoryness, and so on. It proposes that each type of perception is the business of one level of brain organization specialized to extract invariances of that type from the levels below it. This theoretical structure should apply to uses of words just as well as to uses of nonword perceptions -- in the end, they are all just perceptions.

I realize that this has been a very sketchy presentation of my view, and that I probably haven't dealt with any real linguistic problems. I'm just trying to show how different my approach is from yours (although we certainly seem in tune in most regards). Maybe yours will turn out to be simpler and better in the end than mine would be after all the complications needed to account for real language phenomena have been added. Or maybe you can find a synthesis that will satisfy both of us without discarding anything either of us thinks valuable or necessary.

I do hope your category level gets tired pretty soon.

Best

Bill P.

==========

Date: Sun, 24 Nov 1991 19:25:26 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>

Subject: Created realities

[From Bill Powers (911124.1900)]

Wayne Hershberger (911124) --

>There are many created realities, some comprising perceptual
>"objects" (e. g., a laser beam) some comprising conceptual stuff
>(electromagnetic flux), both dealing with the same phenomena.

Is this a correct diagram of what the above says?

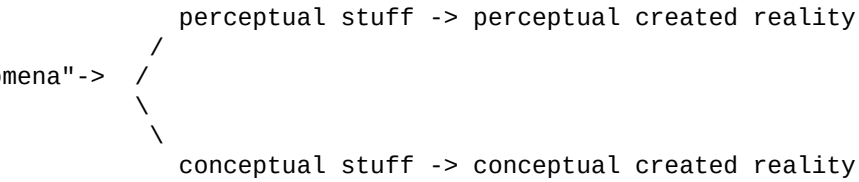

In other words,

"The same phenomena" are contained in
   Perceptual stuff, which is contained in
     perceptual created reality, and
   Conceptual stuff, which is contained in
     conceptual created reality.

Or, is it more like

 Phenomena ->
   Perceptual stuff ->
     Conceptual stuff?

Best

Bill P.

==========

Date: Sun, 24 Nov 1991 22:11:22 EST
From: goldstein@SATURN.GLASSBORO.EDU
Subject: address change

To: CSGnet people
From: David Goldstein
Subject: address change
Date: 11/24/91

Please note that as a result of some changes at Glassboro,
my internet address is: goldstein@saturn.glassboro.edu

Gary, I would appreciate it if you could make the necessary
address change.

I was informed that the address which you have for me is not
working. This is why I have probably not been receiving very
much CSGnet mail lately.

If anyone has responded to my post on the method of levels, I have
not received it because of the above problem. Please resend. Thanks.

==========

Date: Mon, 25 Nov 1991 07:54:55 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: corrigenda and addenda

[From: Bruce Nevin 911122 1833 -- back in Gloucester]

Looking over my post on linguistic vs nonlinguistic vs category perceptions from this morning, I see a need for addenda and corrigenda, as we say in the publishing biz. I have added [missing words] in square braces. I am ignoring relatively trivial omissions such as missing "with" in the paragraph before the third table and missing "them" in the paragraph after the third table. The problems develop toward the end:

--In the second paragraph after the fourth table:

. . . and categorization is [one of the things that] we do best with language. As the crossing arrows in the above partitions of the perceptual hierarchy emphasize, the category level is where language connects with non-language[, and where metalanguage connects with language].

--Two paragraphs later:

What is important to keep straight, I take it, is the difference between perception as mere observation recorded in memory, and [control of] perception resulting in . . . behavior . . . .

--Strike this paragraph. Poorly stated, too many red herrings lurking, and some issues that need to be developed more fully in autonomous context, and indeed are developed in Bill and Wayne's discussion of epistemological issues, e.g. Bill (911116.0800):

>We can divide direct experience into things we do and things that happen.

--The final table and text about it are garbled. Haste and too many work-related interruptions. The paragraph before the final table:

The "stuff of utterances" is controlled in [the lower part of] the linguistic partition of the hierarchy.

--After the above sentence, add:

This mapping (differently from the nonlinguistic side) is indicated by the bidirectional arrow from Relationship<L> to the Category partition. This is an abbreviated notation of course: if this were shown in greater detail than is possible here, that arrow would branch to each level whose categories correspond to verbal expressions.

--In the table itself, there should be two arrows instead of one on the right side. The downward-pointing arrow should be one arm of a bidirectional arrow whose other arm points leftward to the Category partition, parallel to the other arrow coming from the Category partition and upward to Sequence<L>, as follows (abbreviated):

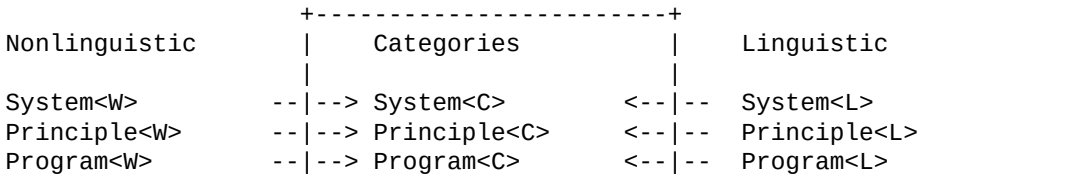

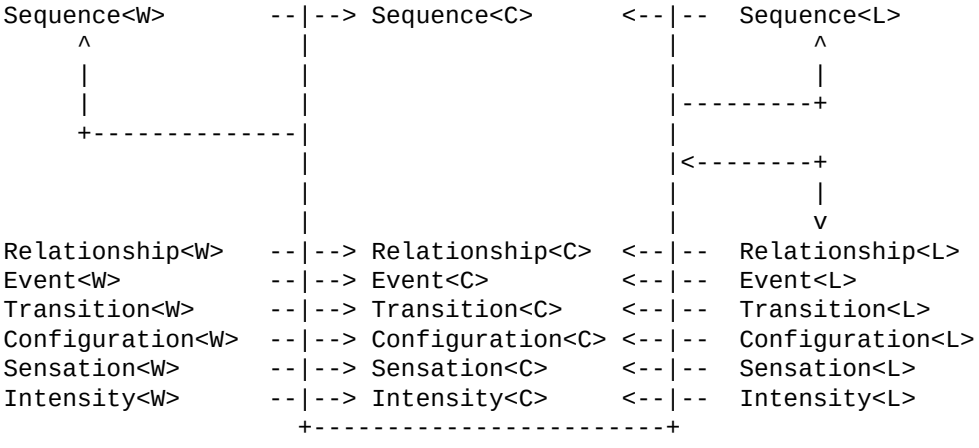

--In the paragraph following this table:

In addition [(just as on the nonlinguistic side)], category perceptions of a linguistic sort are input for the Sequence<L> and higher perceptions of grammar (and of linguistics).

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

The horizontal arrows represent associative links. What this says is that the step of categorization is a trick of associative memory that stands between the nonlinguistic perceptual control hierarchy, which uses environmental feedback for conformity to reference signals of a personal or individual nature (though perhaps chosen for what an observer might consider social ends), and the linguistic perceptual control hierarchy, which uses environmental feedback for conformity to reference signals that are learned as pre-set social conventions. In this view of things, it is the Categorization partition of the perceptual hierarchy that provides the loose coupling between language perceptions (controlled for conformity to language-as-social-product) and nonlanguage perceptions (which we construe as our physical reality), and it is the associative-memory character of categorization and of those horizontal links that makes that coupling loose.

Many of the categories are learned as social conventions. Some, perhaps many, may be created alinguistically and aculturally by the reorganization process. I believe categories may be created in memory and imagination just so as to correspond to constructions in language: chimaeras, unicorns, and "these stickies that are about fixing the step description in the EDGE guidelines." This accounts for the sometime quickness, fluidity, and ad hoc nature of categorization.

It is not clear to me that the things in the category partition form a control hierarchy, or that they need to. Perhaps the configuration apparatus is only a particular generalization of memory, associated with the nonlinguistic control hierarchy on the one hand and the linguistic control hierarchy on the other, and stratified by way of those associative links to perceptions of different levels.

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

(Bill Powers 911122.0600) --

You deny that one person can somehow guide another's creation of a new category perception. This is because in the present model the only

way to create a new category perception is the reorganization process. I have suggested a different conception of what a category perception is. From this perspective, teaching is more possible.

Here is one way that it might work. The student is attending to a set of perceptions<W>1 and wants to make sense of them. The teacher, recognizing a "teachable moment," produces an utterance perceived as linguistic construction<L>1. The relation between the two (the intended meaning of construction<L>1) is not obvious to the student. However, the socially-established structure of language is much more articulate than any structure discernible in the set of possible perceptions<W>, and is known to both teacher and student. In particular, analogies from one construction<L>1 to another construction<L>2 are much easier to make than analogies between nonlinguistic perceptions<W> are. The student is familiar with nonlinguistic correlates of some analogous construction<L>2. Using the linguistic analogy as an index, the student is able to find ways in which perceptions<W>2 are analogous to the problematic perceptions<W>1. By modifying category perceptions<C>2, the student is able to create new category perceptions<C>1 corresponding to the problematic perceptions<W>1. Diagrammatically:

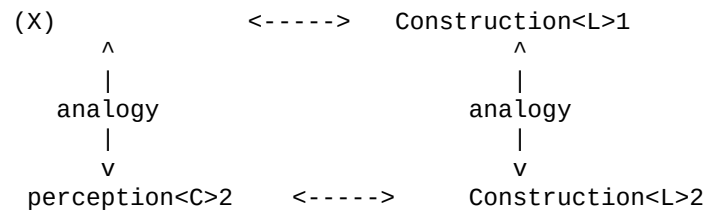

The student solves the analogic proportion for the unknown, shown here as (X), creating a new perception<C>1 to correspond with the linguistic Construction<L>1.

In fact, the teacher is very likely to produce an utterance perceived as Construction<L2>, embedded in a statement of analogy or simile. Such an assertion indicates that the student should draw an analogy between familiar perception<C>2 and novel perception<C>1.

The point is that the student is not left clueless, and can and does use the clues that are available. Motivation (a correctable, not overpowering, error signal) is necessary, of course.

A further point is the suggestion here that we do not invent new categories entirely de novo. Rather, we adapt one or more existing categories as raw material for the new one, by analogy.

But this use of analogy for teaching and learning should be possible within the present PCT model. Where does the role of analogy and metaphor fit in the present model? I propose that it is in those horizontal associative links with category perceptions--links to "real world" perceptions<W> on the one hand, and links to linguistic perceptions<L> on the other. What do you propose?

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

You say:

>It's easy to "teach" infants and toddlers because they can't stand
>having their worlds full of presences that they don't recognize.

This is a privileged case, I agree, but motivation for learning from a teacher does not cease at age 3 or 4, and conversely infants and toddlers are frequently NOT very teachable at all. While people experience "inability to recognize a presence" less and less often the older they get (until memory starts to go), there is also correspondingly less and less need to grow a new configuration perception. And people rely on categories for recognition rather than dealing with configurations directly. Even if there is something unfamiliar or peculiar about that widget, it's still a widget, or so I believe. I'm surmising that recognition and categorization depend upon associative memory horizontally across perceptual nodes, rather than upon error signals and reorganization vertically in the perceptual control hierarchy.

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

Miscategorization is a source of humor, as in the old hillbilly jokes: "Thar war a wonderful spring in that hotel room, allus fresh water when you pump this little handle? But they had this picture frame over it and it got in the way, so we pulled it off and brought it home. Grandpa's picture looks right purty in it, don't you think?" Pan to picture on wall framed by toilet seat. If you miscategorize, then you are in wrong relation to the community, and mammals don't like that. Humor tickles that error: it's reassuring and a relief because the error signal (of which the joke reminds me) is not my current error signal. I experience it only in imagination, and perhaps in memory too.

You say you've read of ethologists and linguists who couldn't get their informants to recognize and act in their role as his teachers. Unfortunately, your example reverses the roles (it's the teacher not catching on, rather than the student). I don't think your point is that he is trying to teach them, trying to get them to reorganize. You might want to try that one on me again. Also, I'd like a reference to an example of such a story. I have a hunch you may have misconstrued it. But I only have my own experience as an anthropologist and linguist to base that on. I know some stories (including my own) about tricksters, but nothing like you describe.

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

I've already responded to the comment on clipping the wings of CT. I'm trying to expand it. It can't handle language yet.

Bruce Nevin
bn@bbn.com

==========

Date: Mon, 25 Nov 1991 07:56:20 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>

Subject: some responses

[From: Bruce Nevin 911123 0643 -- Saturday in Gloucester]

(Bill Powers 911122.0900 aka Peters_R)

Self-organizing has nothing to do with reorganization. The way in which rings arise as an aggregate product of individual control with respect to other individuals is self-organizing in a way analogous to the way in which grains of sand settle into a pile as an aggregate product of their individual physical interactions with next neighbors. Language is much more complex, and has been abuilding over a much much longer span of time. (The a- of "abuilding" and "aboard" was the same word as what is now "on" in the 7th and 8th centuries AD, and this word in turn has Indo-European roots deep in prehistory.) Language is

> the collective art of expression, a summary of thousands upon thousands of individual intuitions. The individual goes lost in the collective creation, but his personal expression has left some trace in a certain give and flexibility that are inherent in all collective works of the human spirit. (Edward Sapir, _Language_)

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

Yes, we do things with language, it does not communicate or signify anything of itself, this is indisputable. Much though picturesque analogies may convey a useful sense of the integrity of language and of its development through time, language is not really a living organism.

But conversely language is not limited to the control of perceptions in human brains, articulators, and sensory organs. It really is out there in the world as a social inheritance, immanent in the using of it by fellow members of our speech community, and made so by our constant verification of our reasonable conformity to social norms (which ensures reasonable conformity to one another) in our use of it. It constitutes and depends upon conventions to which we must conform if we are to use language and indeed if it is to continue to exist at all as the aggregate product of our ongoing use of it.

What is wanted from a PCT perspective is not only an understanding of how we each reconstruct that social product as a personal resource, but also an understanding of the ways in which the particular informational capacities of language may augment the perceptual control hierarchy, may facilitate learning, and may on the other hand contribute to various confusions such as premature recognition (categorization) and closure.

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

You object to the phrase "self-referential sentence" because sentences don't refer, they are used to refer. Fine. This phrase means "sentence used to refer to itself."

By the way, the paradox is resolved in another way, by the recognition that the sentence is not yet a sentence at the point in it that is used to make the reference. Indeed, until it has been completed one can have no assurance that it ever will be a sentence.

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

The point, however, was not self-reference (the paradoxical nature of which I agree is bogus, though I have different reasons for saying so), but use of objects and relations in language to refer to other objects and relations in language. The context from which you took the self-reference example was my attempt to convey to Martin what I mean when I say "metalinguistic" and "metalanguage." It is not just discursive talk about language, as he seemed to believe, but anything that refers to language. Programs that refer to objects and relations in language are metalinguistic programs. In particular, programs used in the perceptual hierarchy for producing utterances are metalinguistic programs. What you and Martin have resisted is my claim that the latter sorts of "performative" programs are part of the language: part of its grammar. Similarly for sequences, etc.

We could step through this from a point of agreement progressively toward points of disagreement, to try to isolate what it is that occasions your resistance. I think we agree that linguistic utterances are in the language. Phonemes, syllables, words, phrases, sentences, and so on, taken as products of the linguistic analysis of utterances, are in the language too. What about perceptions of phonemes? Is language in the brain, in the perceptual control hierarchy, or is it only in the utterances, the personal/social product? The latter is essentially Harris's view, who is therefore comfortable with his agnosticism about how brains accomplish language. Notice that if we press harder, we must ask at what point in the perceptual hierarchy language ends and non-language begins. If you think this would be a fruitful exercise, let me know. I'm game.

Since that search for the limits and bases of agreement and disagreement requires dialog with you, I'll beg the question for now and talk about languistic perceptions, as I have done up to now.

ECSs that control for the understanding and production of utterances, phonemes, syllables, words, phrases, sentences, and so on are linguistic ECSs, and they control linguistic perceptions. Some of them are metalinguistic ECSs or perceptions. (Perhaps you would prefer meta linguistic-perception perceptions and meta linguistic-ECS ECSs?) My point is that the metalanguage variety of perceptions and ECSs are cut from the same fabric as the object-language variety of perceptions and ECSs. I see a close parallel to the relationship of language and metalanguage. They are different sets of perceptions and ECSs, just as metalanguage (portions of) utterances use a special metalinguistic vocabulary that is not used in object-language (portions of) utterances; and just as metalanguage words and object-language words are all words, all in the vocabulary of the language, likewise they are all perceptions and ECSs for language, all in the portions of the perceptual control hierarchy that control language. They have much more in common with one another than either has with, say, the perceptions and ECSs involved in control of canoe paddling or chess playing. The set of possible utterances includes its own metalanguage as a sublanguage. I argue that the set of possible language perceptions/ECSs likewise includes the set of metalanguage perceptions/ECSs as a proper subpart. My proposals regarding partitions of the perceptual hierarchy fit this description.

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

You say: 'It is perfectly possible for the metalanguage to use the

same symbols as the main language, but this is not "necessary".' But I was not making that claim. Metalanguage vocabulary is distinct from object-language vocabulary. I am only saying it is of the same kind. The metalanguage program ECS uses metalanguage categories to refer to objects and relations in a sentence. A different program ECS uses different categories to refer to objects and relations in a chess game. But in both cases they are programs and categories. A speaker uses object-language words to refer to the chess game, and uses metalanguage words to refer to the words of her utterance. (Examples of the latter are words like "the latter" and zeroed statements of sameness for coreference, frequently involved in various reductions.) In both cases they are words. This is the parallel I was drawing.

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

You say:

>The meaning of "nonfirst argument" . . . is NOT the word that turns
>out to be the first argument: that would be the configuration-level
>meaning. The program-level meaning is the implied program-like
>operation and its result: [in the example C-code] TRUE or FALSE.

I agree entirely. The ECS uses metalanguage categories in the program to refer to words in the sentence. Take the following sentence as an example:

John you may recognize.

The meaning of "nonfirst argument" (or equivalent) in the program is the value (the neural signal) returned by a (subroutine) program that tests the relationship of the current word, "John," to the following operator, "recognize". The referent of "nonfirst argument" (that is, of the subroutine) is that word, if the value returned is TRUE (or equivalent); it has no referent otherwise. The subroutine that tests for the required operator-argument relation is a metalanguage entity. It is in the language in precisely the same sense that a sequence detector for the sequence "John" is in the language.

Such a subroutine is required in the listener to undo the reduction in a sentence like the following:

John, whom you may recognize, will be the first to object.

By undoing the reduction, the ECSs in the linguistic perceptual hierarchy recognize that it includes the following assertion:

You may recognize John.

They also recognize the subordinate, ancillary relationship of this assertion to the main one, namely:

John will be the first to object.

Concurrently, the same ECSs are required for the speaker to produce a clarificatory interruption of the main assertion

John--said John you may recognize--will be the first to object.

and to do so in a way that is less cumbersome and disruptive to the flow of communication by reducing it to a relative clause. All

modifiers are reductions of relative clauses, which are reductions of interruptions with metalinguistic identification of the repeated word with the head word just prior to the interruption. The motivation for the reduction becomes excruciatingly obvious when multiple modifiers are adjoined in succession to the same head word. This is a simple example of some of the complexity in language that control systems manage.

Needless to say (I hope needless), I was not proposing that the example sentence "if you say a nonfirst argument before the operator" etc. was itself a program, I was merely using that sentence to indicate the need for and general character of a program as an example of how the control hierarchy must provide for metalanguage aspects of language.

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

You say: "I hope I've come a little closer to saying why I think that the program level is something more general than language." I have no quarrel with this. It is some time since I was trying on the view that all perceptions from category up were in the language. I am sure that the chimpanzee ferreting out termites with a twig is using a program perception, without language.

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

I return your question to you: from what level can you look at all this talk about language and see that it is talk about language? Is a foreigner who can't speak or understand English up a level when standing in a theater lobby or a cocktail party, knowing that talking is going on and even attending to it but not engaging with it?

Bruce Nevin
bn@bbn.com

==========

Date: Mon, 25 Nov 1991 08:58:09 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: not as complicated as all that

[From: Bruce Nevin (911125 0723)]

(Martin Taylor 911122 20:20) --

>I don't remember ever mentioning Berkeley in connection with anything
>you wrote.

This was an allusion to your charge (911120 14:50):

>So you WOULD say that the aircraft flies because Bernoulli discovered
>the equations relating pressure to airflow?

I figured one facetious and not quite accurate non sequitur deserved another. Or did you mean that seriously? BTW, it was Ben Johnson who kicked the stone, the nonverbal part of his rejoinder to Berkeley.

>I find
>nowhere an argument to suggest that the engine that transforms concepts
>into language must use a program to do so, though I have agreed that it may
>do so. I read (possibly wrongly) into your discussion an underlying
>presumption that there is some kind of general-purpose transformation

>engine that can be fed with different programs to do different kinds of
>things, possibly all to do with language, possibly not.

1. Language, understood as a set of utterances, includes an object-language and a metalanguage aspect.

If you disagree with that, please say how.

2. Language, understood as perceptions on various levels of the perceptual control hierarchy, includes perceptions for both the object-language and metalanguage aspect of those utterances. These are all perceptions of essentially the same kinds. For example, the effort/intensity, sensation, configuration, transition, event, and relationship perceptions involved in pronouncing the word "noun" are essentially of the same kind as the corresponding perceptions on the same levels involved in pronouncing the word "dog". The perceptual hierarchy provides for both aspects of utterances.

If you disagree with that, please say how.

3. Language, understood as ECSs (elementary control systems) controlling perceptions of the sorts indicated above, includes ECSs for both the object-language and metalanguage aspect of those utterances. They are all ECSs of essentially the same kinds. For example, the ECSs controlling effort/intensity, sensation, configuration, transition, event, and relationship perceptions involved in pronouncing the word "noun" are essentially of the same kind as the corresponding ECSs controlling perceptions on the same levels involved in pronouncing the word "dog". The hierarchy of ECSs provides for both aspects of utterances.

If you disagree with that, please say how.

Please note that I have not mentioned here any putative mapping from concepts to utterances, with its implicit presumption of some mentalese representation of concepts. I have been at some pains to show why I believe such a view is erroneous. But there is no need to go into that particular fox-trot just now. (I will touch on the point in a note later.)

Assuming agreement with the assertions in the preceeding paragraphs numbered 1-3, I make the following I think really quite benign proposal:

People produce certain kinds of utterances in which, for example, repeated argument words occur in zero form, or reduced to pronoun (such as the -ich or -o of "which" or "who" in a relative clause). People understand that these reduced words and their dependencies on operator words are still present in the sentence, and indeed frequently utter them in non-reduced form, or in differently reduced form, when asked to repeat what has been said.

Some of the ECSs controlling the relevant perceptions (I propose) control the words and their dependencies in completely explicit, unreduced form.<note> A decision is made, in the course of speaking, which utterance-form to give to a word--"John," for example, or the -o of "who," or zero. The ECSs responsible for this decision must must at some point refer to the word in question and determine whether or not it is "the same" as some word already uttered or planned to be uttered soon. The reference to the pair of words and the assertion of sameness are metalinguistic.

The ECSs involved are of the sort mentioned under (3) above, and the perceptions involved are of the sort mentioned under (2) above. There is no essential difference between ECSs that refer to level 1-6 perceptions corresponding to the experience of a dog, ECSs that refer to level 1-6 perceptions corresponding to the utterance "dog," and ECSs that refer to level 1-6 perceptions corresponding to the utterance "word," although the first sort are "in the world," the second sort are "in language" and the third sort are "in the metalanguage." However, to make these distinctions among them requires a point of view external to the perceptual hierarchy. Within the perceptual hierarchy, they are all cut from the same cloth.

These ECSs include program-level ECSs. They also include sequence-level ECSs. To say this is merely to espouse Perceptual Control Theory. It is not

>an argument to suggest that the engine that transforms concepts
>into language must use a program to do so

++++++++++++++++++++

Note: This aspect of perception accounts for part of the intuition of "concepts" being "translated" into words. Please observe the word "part" in the preceding sentence. My post of last week re associative memory and analogy purports to account for other parts. I subordinate this point to a footnote so as not to distract from the main point.

Bruce Nevin
bn@bbn.com

==========

Date: Mon, 25 Nov 1991 17:37:04 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: not as complicated as all that

[Martin Taylor 911125 17:15]
(Bruce Nevin 911125 0723

>>I don't remember ever mentioning Berkeley in connection with anything
>>you wrote.
>
>This was an allusion to your charge (911120 14:50):
>
>>So you WOULD say that the aircraft flies because Bernoulli discovered
>>the equations relating pressure to airflow?
>
>I figured one facetious and not quite accurate non sequitur deserved
>another. Or did you mean that seriously? BTW, it was Ben Johnson who
>kicked the stone, the nonverbal part of his rejoinder to Berkeley.
>
Thanks for the correction. But my comment about Bernoulli was as much to express my puzzlement at your apparent position as to be facetious.

>1. Language, understood as a set of utterances, includes an object-language
>and a metalanguage aspect.
>
>If you disagree with that, please say how.
>

I suppose this is at the heart of the problem. I don't see language as a set of utterances, but as a dynamic sequence of communicative acts that occur at multiple levels of abstraction, of which a set of utterances might be one. Within that set of utterances, I can't say whether I agree or disagree with the claim that it has an object-language part and a meta-language part, because the more I read of your comments on my postings and on Bill's, the less I understand what you mean. Initially, I thought your position was straightforward, because I agreed with almost everything you wrote, except for something I thought was a minor technical detail--whether the language we use can be totally described symbolically. Now it has turned into a very complex problem whose ramifications are quite obscure.

Aha! Perhaps I have it--that "set of utterances" might be the key. Do you think of language as being patterned strings of words? And is the metalanguage also patterned strings of words? If so, we are talking about two quite different concepts when we each use the term "language." In my "language" there are levels for phonemes, levels for arguments, levels for propositions, levels for words, etc. I do not think that the language of phonemes is described in sequences of phonemes, that of arguments in sequences of arguments, and so forth. As Bill says, these are entities of qualitatively different character. I think all the levels are executed through a common method, though the implementation may be quite different at the different levels. The objective of each level is to reach the conclusion that the two parties to the communication have an apparently satisfactorily similar view of what the communication was for (not necessarily of what it was; it is the intention that matters, not the action). Perhaps some of these levels can be best executed by activating a program in some more generally useful engine, while in others, the best is to activate a host of analogies that sum to an action that "won't be too far wrong." One might perhaps be able to describe the first symbolically (though it might not be performed symbolically), but I think the second would be hard to bring into the symbolic fold.
== ==

I have about 40 messages saved over the last 6-8 months that I intended to respond to, over half from you. This message is going on that stack, because I need to try to reconceive my model of what you are saying, in light of your older postings and the new ones to Bill. At present, I can't deal with whether I agree or disagree with your second and third points, because they both make claims about the object-language and the metalanguage being effectively the same, and I no longer have much idea what either concept means.

Perhaps now is the time for me to give up on following your thread until I have read Ryckman, which you kindly sent me, and instead to try to put together my long-promised essay on how I see language (Layered Protocols and CSG).

Martin
==========
Date: Tue, 26 Nov 1991 06:58:40 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: complications, implications

[From: Bruce Nevin 911126 0655]
(MMartin Taylor 911125 0723) --

Thanks, Martin. By the time you get back to this, I may even be able to express the issues more clearly and succinctly.

I don't regard language as just the collection of utterances or
the "set of sentences" as linguists have been wont to do, but I
chose that vantage as a starting point because it seems easiest
from there to see the object-language metalanguage distinction.
It is an aspect of language for which we must account, however,
namely, language as social product and social heritage.
Consider it as one circle in a Venn diagram. Some would say
this circumscribes language, and that the perceptions and
control systems in a language user are non-language. Generativists
say that the social-product aspect is an assessment of performance,
and that language "really" concerns competence incorporated in
the human nervous system in some unexplained way by a combination
of biological inheritance of innate language-specialized
mechanisms and selection of values for parameters in those mechanisms
on the basis of exposure to the social product in the environment.
This "competence" lies somehow behind the layers of perceptual
control that concern us here, which Generativists would view as
parts of a performance model. There are other views. The question
is, where is the boundary between language and non-language?
Generativists have performance mechanisms that control behavioral
outputs but are not part of "language" for them (perhaps--they
really don't talk much about mere performance). You have protocols
that control behavioral outputs in the social product but are not
part of "language" for you (if I understand you). Where is the
boundary between language and non-language in the control
system, and why is it just there?

Not asking for a response now, but posing this by way of summarizing
my view of the core of this thread, for when we take it up again.

        Bruce Nevin
        bn@bbn.com
==========
Date: Tue, 26 Nov 1991 08:29:14 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: whither categories?

[From: Bruce Nevin (911125-26 AM)]

(Bill Powers (911124.0800)) --

My aim is to account for language in PCT terms. I have tried to do so
within the present resources of the theory. The "partitions" proposed
are entirely metatheoretical, just to verify that all the bases I know
about are being covered and that they are being covered in a coherent
way. Sometimes it helps to rearrange familiar furniture, to see it
fresh.

In particular, I am not proposing a separate "language organ." As I said
in a post Monday morning, and indeed in the post to which you are
responding: these "partitions" are all equivalent from a perspective
within the perceptual hierarchy, the same kinds of ECSs controlling the
same kinds of perceptions.

Your way of distributing language perceptions among associated
non-language perceptions has obvious appeal, and (to quote a line I have
seen from you more than once) I wish I had thought of that. The issues
are not resolved, however, so continue to be patient please and bear
with me as I go through each of the two views in turn and then look for
their unity.

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

The initial partition between perceptions<W> and perceptions<L> corresponds to our distinction between language and everything else. Loosely speaking, we control for the "stuff of utterances," we control for the "stuff of nonverbal experience," and we control for correlations between the two. Control for recognition and articulation of words is completely distinct from control for nonverbal recognitions and behavioral outputs. At intermediate levels of the hierarchy we bring the two into correspondence.

In all of this, there are no additional principles.

Even when we wrestle with the arbitrary, conventional, and normative aspects of language as an inherited social product, the added complexity is outside PCT, something to be perceived in the world of interactions with other humans. No additional principles required for the theory.

But the correlation of language with non-language is a loose one. Even in the unreduced, informationally complete and explicit base forms disclosed by an operator grammar, words do not map in a simple way onto nonverbal perceptions. One reason this must be so is that language is conventionalized, nonverbal perceptions are not.

A first approach to this says that words refer to category-level perceptions. That part of categorization that is learned as an inherited social product is arbitrary, conventional and normative. Well and good, this is a PCT reincarnation of familiar ideas about a "mentalese" that gets translated into language. This justified the distinction between those category perceptions for which we have words and those for which we don't. At this point, we moved Category<L> over with language, distinct from Category<W>. The latter are, according to hypothesis, not arbitrary, conventional, normative, or learned as part inherited social products. They are invented by each living control system idiosyncratically. They are therefore not useful for communication. But some of us want them in the theory, so there they are.

Still no new principles. The apparent complexity is only metatheoretical, in the verification that we are accounting for what we know already about language. (Yes, that knowledge is open to challenge.)

But the looseness of the correlation of language with non-language is not just due to its conventionalization. There are problems of polysemy (the primer is under the finish coat of paint--but I'm pointing at the wall, or the ceiling, or referring to the other side of a wall of a box as it faces me), metaphor (the pipe collapsed, the Pope collapsed, the government collapsed, the argument collapsed), and much more. An example: extension of a word to a new range of word cooccurrences on analogy to word existing domains rather than on the basis of analogies between the several nonverbal perceptions to which the words "refer." For this, consider the reduction of "full" to the sufix -ful. In Old English, we have "careful" beside "full of care," "sinful" beside "full of sin." Later, in Middle English, we see also words like doubtful and graceful, and sure enough "full" also occurs at that later time in the phrases "full of doubt" and "full of grace," which it did not previously. But still later, we find words like resentful and forgetful, where there is no simple "full of ___" phrase. Here, you

might argue that it is the perception of someone feeling resentment that is made analogous to the perception of someone feeling doubt (or someone in a state of grace or of sin) that is the basis of the analogy. But what of the yet later words like lawful? Here, there is no sense of "full of" at all, but rather "consonant with" or "in accord with." And in words like grateful even the meaning of the first element grate- has shifted (original gratia, so that grateful has the same Latinate origin as graceful, only from two periods of borrowing into English). This shows that even at the earlier stage the basis for analogy was parallelism in the linguistic elements involved. The morpheme -ful came to be controlled as a productive suffix rather than as the second word in a compound. (To clarify the distinction: Man-child is a compound, manful is not.) Yet the ordinary meaning "full" remains today in at least some of the older coinages (sinful), though the meaning of the stem to which it is attached has shifted for others (graceful): for them, the meaning of the suffix -ful has not changed.

So there is this loose coupling of language with nonverbal perceptions that is not adequately accounted for in the existing theory.

In addition, there are some peculiarities about categorization. We seem able to invent new categories with a fluidity and speed that exceeds our ability to invent new kinds of relationships, events, transitions, and configurations. Moreover, this fluidity and speed in inventing new and arbitrary categories seems to inhere in Category<L> perceptions rather than Category<W> perceptions, so far as I can tell.

We seem to use categories a lot in associative memory. Now, granted, there is proverbially nothing more evocative of remembered experience in all its full-blown complexity than an odor. But is this really a typical example of associative memory? In my observation of categorization, I use it a lot for the kind of associative hooks that I described.

This is the point at which I removed categories from the hierarchy.

>In effect, categories themselves have disappeared from the hierarchy. My
>proposed law of awareness goes like this: you are never aware OF the
>level that you're aware FROM. Instead, the perceptual mode of that level
>appears to be the main feature of the world you're examining.

I repeat: I removed them from the hierarchy. They didn't incidentally disappear from the hierarchy as a hitherto unnoticed side effect of what I was doing. And I repeat: I observe my use of categorization as hooks in associative memory. By our maxim, I must be viewing categorization from a place other than its location. I submit that my vantage point on categorization is the perceptual hierarchy, and that when I observe categorization I am observing associative memory.

I still have not added any new principles. Associative memory has been there right along, though the theory hasn't had much to say about it.

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

A difference in theory should be testable. What testable differences does this difference make? I have suggested some, those because of which I proposed it. I contend that a simplicity argument won't hold water, since (as I have argued) I have added no new principles to the theory. I have changed only their a rrangement relative to one another.

Perhaps you can tell me some testable reason which of the following is to be preferred:

Classic PCT

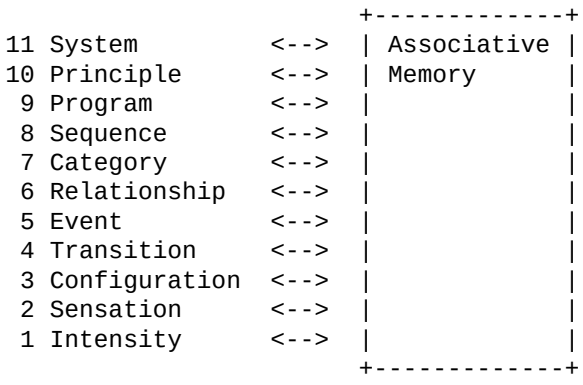

Heretical PCT

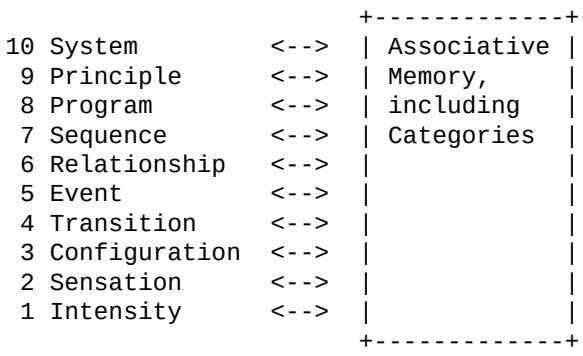

It appears to me that the heretical version reduces the complexity of the theory by identifying two things previously thought distinct, and that in doing so it helps to explain both of those things: categorization and associative memory.

This is not to say that "memories" are anything other than remembered perceptions, or that memories are located in some kind of memory organ corresponding to the box in the diagram. As before, I am partitioning functions, not functors. The horizontal arrows represent mappings from a given level into itself or into a lower level--from a perception of level n into a set of remembered perceptions of level n and level n-1.

This mapping could be by way of shared connections to higher-level ECSs just as in your proposal. There, "line" was associated with ----- by way of their shared connection to a higher-level perception CW3.

In this is the germ of an idea that categorization is simply a generalization of the change in type from any level to the next higher level.

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

In your proposal, a typical category-level perception includes both nonlinguistic perceptions<W> and linguistic perceptions<L>.
(My alinguistic Categories<W> would specify no linguistic member.) The correspondence between the two inheres in the fact that any one of them fulfills the input requirement of the category-level ECS.

>To narrow meanings down,
>higher systems specify more than one category at the same time. Then you
>get the lower-level set that exemplifies them all.

So to produce the utterance "hand me the triangle" the sequence-level ECS specifies CV1, CV2, and CW2 simultaneously for the last term in the sequence. This corresponds to a sort of feature analysis of vocabulary.

* The perception represented here by "Triangle" is a word (in the same CV1 category with all those other relationship perceptions that constitute words, and in the same CV2 category with a smaller set of words).

* The perception "triangle" can be used to refer to any of the other perceptions /\ that CW2 specifies.
  /__\

But wait a minute! Why does the linkup with other members amount to reference in the case of CW2 but not in the case of CV1 or CV2? Why the privileged status so that "triangle" can be used to refer to a triangle and not vice versa? Worse, what prevents the following from being equally so:

* The relationship perception "Triangle" can be used to refer to any of the other relationship perceptions that CV2 specifies, such as "square" and "line".

The key, you may say, is suggested in the word "use." It happens that no sequence perception specifies ("uses") the relationship perception "triangle" in this way. Well, I'm a bit queasy about that one, but OK for now.

What about those two language categories, CV1 and CV2? The difference in their ranges is interesting. To point up that difference, let's suppose there is a CV3 that includes "noun" and "word" but not the others.

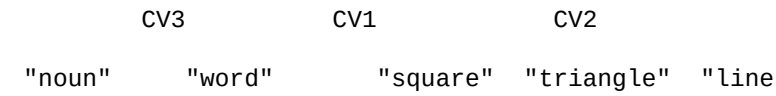

As with your examples, we have to draw in or imagine the lines that we indicated verbally. And to make a useful distinction, let's erase the line from CV2 to "word".

* The relationship perception "triangle" is a word in the object language rather than a metalanguage word (CV2).

* The relationship perception "noun" is a word in the metalanguage rather than being an object-language word (CV3).

What is the basis for these assertions? The fact that some sequence perception is going to use CV3 to refer to members of CV2--perhaps even in a sequence that includes the relationship perceptions "in" and "the" and "language" in that order. Well, that's not a big problem.

There appears to be a categorial hierarchy, in which CV1 (all words) subsumes CV2 (object-language words) and CV3 (metalanguage words). This is just like a classificatory "is-a" hierarchy such as thing, animal,

dog, collie. On your view, such hierarchies of categories are a function of the number of their I/O connections to the sequence level, and for which groups and subgroups of categories the other ends of those I/O connections terminate in the same sequence perceptions (representing the same point in the sequence). Apparent categories of categories are side effects of these connections, not represented hierarchically in the control hierarchy. The relation of metalanguage to language appears to be a special case of this.

Reverting to unhygenic vernacular for a moment, there is something most peculiar about denominating a "category" a grouping that includes an indeterminate number of like objects (triangles) and one object of a very different sort (the word "triangle"). This is the source of the queasiness I confessed earlier. OK, to assuage our sense of what the word "category" means we have to say that its technical definition in HCT includes lexical semantics. Under this technical definition, each category perception includes one word among its members. Furthermore, the I/O connections to the category ECS for this word must be such that higher-level sequence perceptions use this word to refer to the other members. Perceiving the set

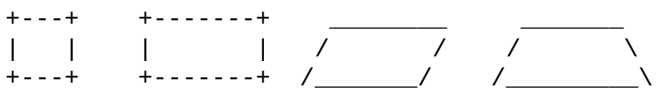

I may say they are all "quadrilaterals" or "four-sided figures".
I do not say they are all +---+s. This is because +---+ is just one
member of the set, and | | | |
the perception that   +---+   +---+
corresponds to "quadrilateral"
or "four-sided shape" is just a signal
(or equivalently a connection) from a shape-detector ECS.

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

Now: is it not the input function to an ECS that does the categorizing of lower-level perceptual inputs and determines category membership?

What if categorization is distributed throughout the hierarchy, just as associative memory is, and just as you say language is?

This would account for some differences of level that might otherwise be problematic. A morpheme like "dog" oe "die" or "-ful" is an event perception: a unitary, familiar sequence. A word like "doghouse" may be a relationship perception. A word like "grateful" may be an event while a word like "sinful" may not. In a conversation about legless frogs one migh perhaps refer to "legful" frogs. An idiom like "bought the farm" or "kicked the bucket" must be equivalent to "die" for some ECS, yet it appears to be a sequence. The geometric shapes categorized with "triangle" in your example are configurations, yet the latter is an event--or is it a relationship or sequence of tri- + angle? Try "three-sided figure" as another member of CW2.

What I am proposing is that there is no need for a category level. The input function of event or relationship ECSs for morphemes and words may accept signals from configuration ECSs like that for /\ as well as from those for the configurations and transitions of   /__\
phonology. The connection by which a word is associated with nonverbal perceptions is itself the source of categorization that we sense behind or with that word. And to the extent that there are events and

relationships for which there are no words, there can be categorizations for which there are no words.

Can this simplified theory do all the work of the erstwhile category level? I don't know. It was never clear to me what work the category level was to do.

Over to you, Bill.

Bruce
bn@bbn.com

==========

Date: Tue, 26 Nov 1991 06:29:44 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Language

[From Bill Powers (911125.0800)]

David Goldstein (911124) --

Roger on the change of address. I wonder if there has been a paucity of responses to your "method" posts because listeners are politely waiting for the linguistics thread to be finished. Be advised that none of these threads is likely to be finished until all of them are woven together. The polite will not inherit the earth. People with random comments or new threads to start, please horn in. There was a request for comments on suicide not long ago. What would CT have to say that might be helpful in this area? Clinicians?

Bruce Nevin (911123) --

>Self-organizing has nothing to do with reorganization. The way in which
>rings [in gatherings] arise as an aggregate product of individual
>control with respect to other individuals is self-organizing in a way
>analogous to the way in which grains of sand settle into a pile as an
>aggregate product of their individual physical interactions with next
>neighbors.

My understanding of how you mean this is that language is self-organizing in the sense that a community of people, each controlling for a private understanding of the community through interacting with other members of it, will converge toward a stable form of language without any external agency overseeing this process to make sure it converges in a particular way. If this is what you mean, I agree with it.

There are too many subjects floating in the air; if I don't delete the thousands of words I have just written and try to get to some point, we will just expand this discussion geometrically with every iteration. So here goes a huge chunk into the bit bucket.

Poit.

You say

>...we must ask at what point in the perceptual hierarchy language ends
>and non-language begins. If you think this would be a fruitful
>exercise, let me know.

I do. I will describe a picture of how language might fit into the hierarchy. I won't try to make it definitive -- that's your business.

Then you can say what it leaves out or why it is wrong and present your own. Then we can iterate as often as needed or until we get tired.

1. Words are perceptions. They are configurations of sounds or sights that are recognized and controlled in exactly the same way as other sounds or sights. "Noun" and "fish" and "category" are all configurations (I'll allow anything up to events), as are all other spoken or written words. The only difference between "noun" and "fish" is the way those configurations/events are used by higher-level systems. Any easily-manipulable configurations/events in any sensory modality could be used to form words.

2. Words can be perceived in relationship just as can any other perceptions of equivalent level. Words in relationship can be categorized like any other perceptions. Categories can be named: that is, a word can be used to evoke the same category-signal evoked by any other member of the category, just by arriving at the input of the category perceiver that accepts it.

3. The name of a category is a configuration/event like any other configuration/event. It exists at the configuration/event level, not at the category level. Sets of word configurations/events can be seen in relationship and be categorized. So there can be categories of words. These are linguistic categories because their elements are words.

4. Linguistic categories can be named by including in them a word-name as a configuration/event. These words are then names of metacategories. They can be perceived in relationship, and can be categorized ... und so weiter.

5. We can therefore account for the existence of nonlinguistic perceptual categories, linguistic categories, and linguistic metacatagories (to any depth) strictly by use of the same levels of perception and control used with all kinds of perceptions.

6. Linguistic categories, the elements of which are words, would be received, perceived, and controlled as parts of sequences, programs, principles, and system concepts related to and derived from words. Linguistic metacatagories, the elements of which are the names of linguistic categories (such as "noun," "verb," etc) would be received, perceived, and controlled as parts of sequences, programs, principles, and system concepts related to categories of words: the system concepts are linguistics itself.

7. The same hierarchy of perceptual and control processes, with the same basic organization, is used for control of all perceptions whether they are nonverbal or language-related. The levels are the same; the specific elementary control systems (ECS) are different.

8. I have left out any associational cross-connections between the linguistic and non-linguistic uses of perceptions. For the time being, let's stick to the extremes: all-language or non-language. These partitions can be kept apart as long as meaning doesn't come into the picture. Meaning will come into the picture where purely linguistic operations can't account for judgments about the goodness of sentences and the like.

As Chuck Tucker would say, is that systematic or what?

Judging from your comments to Martin Taylor, you should have little

objection to the above.

>... it was Ben Johnson who kicked the stone, the nonverbal part of his
>rejoinder to Berkeley.

My rejoinder to Johnson would have been "Boy, I'll bet that stone really
felt that kick!"

>You say you've read of ethologists and linguists who couldn't get
>their informants to recognize and act in their role as his teachers.

I was thinking of the ethologists and linguists as being, say, the Fiji
Islanders, and the informants as being some American tourists talking
argle-bargle and thinking they were funny.

Best

Bill P.

==========

Date: Tue, 26 Nov 1991 12:10:35 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: phonetics programs for Macintosh

Linguist List: Vol-2-819. Sat 23 Nov 1991.

Date: Thu, 21 Nov 91 09:15 PST
From: Peter Ladefoged <IDU0PNL@MVS.OAC.UCLA.EDU>
Subject: Re: 2.796 Queries: Phonetics for the Mac

With regard to programs that might help in teaching linguistics,
the following phonetics-related programs for Macintosh computers
(minimum 1 megabyte of RAM required) are available from the
UCLA Phonetics Laboratory.

Write to me (Peter Ladefoged, idu0pnl@uclamvs.bitnet) for further information

Sounds of the World's Languages
database of digitized sounds found in the world's
languages produced by native speakers;
illustrates less familiar sounds
(Hypercard 1.2 or higher;
about 25 meg. Hard disk required)

Acoustic phonetics
Hypercard tutorial on some basic concepts of
acoustic phonetics.

PlotFormants plots formant frequencies in Mel,Bark or linear
scales, with axes as in A Course in Phonetics.
draws ellipses around collections of points
showing the Standard Deviation

UCLA/Uppsala version of SoundWave
Reads MacRecorder (SoundEdit) files
Shows pitch curves, also LPC and FFT
analyses giving values of formants

A basic, nothing fancy, analysis system

Vocal 2.1 allows the user to create vocal tract shapes and assesses the acoustic results of different articulatory configurations. Somewhat dated

Draw vocal tracts
A utility program making it easy to draw phonetic diagrams.

IPAMacintalk text to speech program using IPA transcription based on the public domain program 'HyperMacintalk' (Hypercard)

MacSynth an old formant synthesizer demonstration program Not compatible with the SE/30 or MacII series

IPASounds digitized recordings of John Wells and Susan Ramsaran linked to the IPA symbols; When used with MacRecorder lets the user practice the sounds as well as learn the symbols See SPECIAL NOTE regarding payment for this program

All the above are available from:

Phonetics Lab
Department of Linguistics
UCLA
Los Angeles, CA 90024-1543

Prices Effective September 1, 1991 (subject to change without notice) Make checks payable to: Regents, University of California
(Sorry, but we have to cover our costs)

A: ___	All UCLA software (B & C)	$65.00
B: ___	Sound's of the World's Languages	$50.00
C: ___	Individual programs	$5.00/disk

SPECIAL NOTE

The program IPASounds is free to IPA members. Others must pay $20
For this one program a separate check is required
payable to the IPA (not to UCLA)

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

cross-post by Bruce Nevin bn@bbn.com
==========
Date: Tue, 26 Nov 1991 09:58:16 PST
From: marken@AERO.ORG
Subject: APS, Levels

[From Rick Marken (911125)]

Tom B, Joel J, Bill P -- the aps application has been sent. Now what if they actually accept it?

Here is a rather belated reply to an invitation to join a conversation about temporal aspects of perception.

Martin Taylor wrote:

>The major finding that I remember is that when the two events are close
>in time (<20 msec?) or far apart in time (>70 msec?) the two orders are
>easy to distinguish, but at intermediate time separations they are not.
>What this says to me is that a fast transition is a single phenomenon
>that can be incorporated as part of a configuration, whereas the slow
>transition is a sequence of configurations.

And Bill Powers said:

>Rick Marken should joint this conversation and describe again his
>experiment with recognition of perceptions of different levels.

The mis-spelling being, I'm sure, an acknowledgment that I am a product of the 60s.

My experiment is quite simple: numbers are switched on and off in two different positions on the screen. If time is progressing down the page you could represent the experiment as follows

```
2
    5
3
    1
8
    4
```

etc

The subject can control how long each number is displayed, which determines the rate at which the numbers alternate back and forth. What is perceived depends on the rate of alternation. At a very fast rate, the perception is of two, simultaneous, flickering number streams. When the rate of alternation slows to about 7 per second you perceive "movement"; the numbers appear to move back and forth. This could be called a transition or an event. But it is definitely higher than configuration since you can see the numbers in the stream quite clearly. What you cannot perceive is the order of the numbers (the sequence). To do this, you need to slow the alternation rate to about 4 per second. The perception of sequence does seem to have a verbal component; I know I am able to see the sequence when I can say which number came first and which came next. This verbal component is, I think, an artifact of using numbers. Finally, when you can perceive the sequence you still cannot see the rule that underlies the sequence of numbers. The numbers are generated using a rule like ; if the number on left was even then number on right is > 5, else number on right is < = 5. Even if you know that this rule exists, you cannot "see" it until the rate of alternation is quite slow -- like .3 per second.

I like this demo because it helps me "telescope" out some layers of perception that are ordinarily all present simultaneously in experience. It also suggests a way to look at a subject's ability to control a certain type of perception (like a sequence) in isolation from others. This would be achieved by varying the rate at which events are presented to the subject.

Anyway, I think this method has intersting possibilities. Suggestions would be most welcome.

Best regards

Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Tue, 26 Nov 1991 14:50:37 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: convergence beckons

[From: Bruce Nevin (911126 1340)]

(Bill Powers (911125.0800)) --

As you would say to Chuck: systematic, definitely not what.

I indeed have little difficulty with your summation. I think we are approaching genuine as opposed to spurious convergence here.

>language is self-organizing
>in the sense that a community of people, each controlling for a private
>understanding of the community through interacting with other members of
>it, will converge toward a stable form of language without any external
>agency overseeing this process to make sure it converges in a particular
>way. If this is what you mean, I agree with it.

Case in point.

Residual problems, now.

I suggested how categorization might be a pervasive property of input devices at each level rather than a level itself.

>I'll allow anything up to events [as word-perceptions]

The problem here is where an affix is a reduction of a word or a phrase. The examples with -ful give a hint of what is involved here. There has to be some way of going from an explicit "in accord with the law" to the reduced "lawful" and vice versa. This sounds like a program with an inverse. A great deal of paraphrase (grammatical paraphrase) falls out here. Given such a mechanism I can except that "-ful" is an event, or even that "lawful" is an event. The path I am looking for would map an event like this onto a sequence of events by way of a program. Kosher? Without that, then I have problems.

Perhaps more pointed are examples of reduction to zero, where at the program level an argument of an operator may have zero form in certain (not all) contexts of repetition. Consider:

Mary plays flute and John [plays] harp.

As word perceptions evoke associated nonverbal perceptions it is clear that a perception of John playing a harp is evoked. My sense is that the linear word-sequence is mapped onto a non-linear lattice or semilattice of word-relationships. On the one hand, this set of relationships may be linearized into word-sequences (with zeroings and

reductions) in various ways. This is still in language. On the other hand, "play" (as in musical instrument) is associated with perceptions of a different order. It is not clear to me whether that is a relationship between John and the harp or a sequence of actions John performs with the harp, perhaps even programs and principles that I perceive John as controlling, especially if I play a musical instrument and associate my experience of it with what I witness with John and his harp. (BCP does not help much here.)

The nonlinear lattice of word-dependencies, which may be linearized in various ways as sentences of a discourse or contributions to a conversation, is I think closer to the organization of nonverbal perceptions without being identical with it. It is in language and some of its constraints as to available vocabulary and what may be said "about" what (operator-argument classes) are conventional and not necessarily shared by nonverbal perceptions. These limits may help to channel or focus attention. It looks to me as though there were program functions that control this. But this would not be a program as we usually think of it, since we could not write it out in the form of linear sequences of symbols--the means for doing that being language or some formal derivative of language. So there is an interesting conceptual task there for stating this aspect of the theory without using linearized language. My Master's thesis at Penn started to wrestle with this, but what did I know in 1969!

Consider:

Mary plays flute and John, I believe, [plays] harp.

Here, "believe" is associated with a relationship between the speaker, "I," and the perception (whatever it is) associated with "play." (Again, the latter has zero form at that place in the utterance). Here are the word dependencies in the second conjunct:

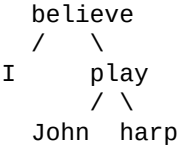

I mean by this only to signal that this sort of nesting of relationship and other perceptions must be nonverbally present. In language, we can say that the word dependencies and zeroings are controlled by a program perception. In the corresponding nonverbal perceptions I do not think that we can say this, although as mentioned language may well prime us regarding to what we should apply our attention.

It may even be that nonlinear dependency structures of words, which might be linearized as various sorts of utterances and discourses, might be our means for learning, modifying, and remembering perceptions at the highest levels, system concepts and principles. This is where something like use of language for "self instruction" might come about, though that particular phraseology is ripe for misconstrual of various sorts.

Enough for now. I hope for the good fortune of having something additional from you to take away for the long weekend, but if not I have plenty to mull over. I will be here tomorrow, then (ah! these leaps of faith!) back Monday.

Bruce

Date: Tue, 26 Nov 1991 13:57:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: epistemology

[From Wayne Hershberger]

Bill Powers (911124.1900)
Bill asks Wayne if this is his meaning:

```
>                           perceptual stuff -> perceptual
>                          /
>"The same phenomena"->   /
>                         %
>                          %
>                           conceptual stuff -> conceptual
>                           created reality
>
>In other words,
>
>"The same phenomena" are contained in
>   Perceptual stuff, which is contained in
>     perceptual created reality, and
>   Conceptual stuff, which is contained in
>     conceptual created reality.
>
>Or, is it more like
>
> Phenomena ->
>   Perceptual stuff ->
>     Conceptual stuff?
```

Neither alternative is quite correct, but the former is the closer of the two.

I would say that the natural order immanent in the psychophysical flux is realized both in the form of perceptual stuff and in the form of conceptual stuff. And when these two types of realization deal with the same phenomena, they often seem to depict a contradictory nature. As perceptual stuff, the desk at which I am sitting (an objective perception) is a static and solid object, but I may simultaneously regard it conceptually as a collection of whirling dervishes hurtling about within a confined space virtually as empty as an inflated balloon. The two types of stuff (i.e., the two types of reality) are, as I've said before, twin born of the psychophysical flux (there is nothing like reason to make sense of something). This is not to say that objective perceptions (e.g., meter readings) are not useful in testing hypotheses derived from a theory representing a putative conceptual reality, but only that the perceptions involved in such tests should NEVER directly involve the phenomenon being conceptualized. That is, you do not settle the geocentric vs. heliocentric world-view issue by watching a sunrise.

Further, I do not mean to say that there are only two levels or types of reality. For example, within the sphere I am calling perceptual stuff, are to be found such things as real illusions (e.g., Ames's window, colored shadows, etc.) which are as

different from the other perceptual stuff called real objects as
the real perceptual objects differ from the conceptual stuff.

Furthermore, speaking of realities, as if there were a limited
number of types, is misleading. Reality is but a dimension upon
which we order the truth value of our countless epistemic
creations--using replicability and parsimony as our criteria.

Warm regards, Wayne

Wayne A. Hershberger Work: (815) 753-7097
Professor of Psychology
Department of Psychology Home: (815) 758-3747
Northern Illinois University
DeKalb IL 60115 Bitnet: tj0wah1@niu

==========

Date: Tue, 26 Nov 1991 17:52:08 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Nonlinearity and dimensionality

[Martin Taylor 911126 17:30]

I want to continue an old thread in perhaps a new way, starting with a request
for Bill P. to provide a tutorial.

Bill has emphasized on many occasions that the elemental control systems (ECSs)
are scalar operators. That is, their job is to match a single input value
(a number) to a desired reference value (another number). To do this, they
try (forgive the anthropomorphism) to reduce the difference between the
current input (perception) and the reference, by emitting an error signal
that is transformed into an output that provides references for supporting
ECSs.

We have argued about whether it matters that the relation be linear between
error and output, and Bill has said something like it doesn't matter much,
because all that nonlinearity does is to introduce harmonics and complicate
the rate at which perception approaches reference, and that he has gone so
far as to use a sixth-power relation and things still work.

It does matter what sign relates the error to the reference signal provided
to a supporting ECS. If the sign is wrong, the supporting one will cause
the percept arriving at the one we are considering to move the wrong way,
increasing the error signal. Rick provides a nice demonstration that control
systems exist by changing the sign of the control-display relationship, which
completely baffles a simple 2-layer (?) model but for which humans compensate
by making a compensating internal switch of sign.

In the discussions I have read, it is implicitly assumed that the whole
system including the environment behaves smoothly, and that individual
dimensions can be found that act independently. As far as I can see,
these are necessary requirements for the simple model to work. What I
would like to see is a discussion of a simple situation in which I am
unable to see how a set of scalar ECSs can perform correctly.
I think that a carefully considered explanation of this situation would go
a long way to clearing up some of the discussions we have been having about
language.

Consider this picture (alpha-graphic pictures seem fashionable these days):

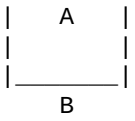

Imagine that the lines represent a wall, that you are at A and your reference condition indicates that you should be at B (equivalently, you want to drag something from A to B and the line represents a metal rod-hanger such as is found in a closet).

At some low level, there is a control system that has x-location as input, another that has y-location. Y-location shows no error, so the Y-changing ECS does nothing. X-location shows a positive error, so the X-changing ECS lowers your position -- until you come to the wall, when it can go no further. You come to an out-of-range force requirement.

Now one might think that a higher-level system could solve this problem, but I cannot see how, provided that the higher-level system works only on one scalar number as its percept and its reference. It seems to me that the interaction between X and Y is essential to the solution of this problem. X-position can be brought to its reference if and only if Y-position is moved away from its reference temporarily. Furthermore, as I have drawn the diagram, even this solution is true only if X-error is first increased.

I do not think that an appeal to reorganization is a correct way to solve this, even if it would work, because the solution is directly obvious if a vector-based control system can be used in which x-y position can be mapped onto the force-acceleration behaviour of the control system, whereas reorganization would seem to provide only a statistical solution, depending on which control-display relationships got juggled when. And even then, why should the Y-control bother to participate in the reorganization, given that its percept already matches its referent?

I fear that I am demonstrating my ignorance of PCT, but so be it. I'd like to be convinced that scalar ECSs can do the job, but I don't see how. If this one is solved, then I would like to extend the problem to the area of bifurcation and thus choice, categorization, and therefore planning. At that point, we may be in a better position to share a common language to talk about language.

Bill, please?

Martin Taylor

==========

Date: Wed, 27 Nov 1991 07:55:34 PST
From: marken@AERO.ORG
Subject: Nonlinearity and dimensionality

[From Rick M (911127)]
I never got an echo on this posting so I'm resending.

[From Rick Marken (911126)]

Martin Taylor (911126 17:30) asks a great question about how the PCT model can solve a problem that requires maving away from a goal in order to get closer.

I think you will get a couple more questions (at least) about the problem: for example, are you assuming that A can see B through the wall. Also, since reorganization is not to be the basis of getting to B (a good requirement -- with reorganization my little e. coli can solve it in no time with a very simply perceptual/output model) you must say something more about what A "knows" about getting to B.

Still, I think it would be an interesting challenge to build a PCT model that can solve your little problem. I am also looking forward to getting some suggestions about how it might be done. I think it would be nice to have A be able to get to be regardless fo where A is placed within the wall and regardless (within limits) of the shape of the wall.

What makes the problem interesting to me is that it seems to require a temporal sequence of outputs in order to produce the required sequence of inputs that solve the problem. This might require building a system that can perceive a program of events -- "if wall at x, then move to y, else z". Of course, the program would have to define the input, since A should be able to get to B even when there are changing disturbances -- like changes in the shape of the wall?

Hasta Luego

Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Wed, 27 Nov 1991 09:28:42 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: epistemology

[From: Bruce Nevin (911127 0743)]

An observation about the ongoing discussion of epistemology. I believe that in your language, Bill, you assume for yourself the perspective described by your (our) theory. This is an excellent way to test its adequacy, but perhaps deserves calling out for notice.

The theory or model requires there to be a "boss reality" there in the environment to complete the feedback loop. Without it, perceptual control is impossible. So from its perspective there has to be a reality there.

However, the theory does not describe this reality or prescribe any attributes, other than that it be present and in at least some respects stable. Capable of being modelled within the perceptual control hierarchy.

In an important sense, this environmental reality is hidden from the perceptual hierarchy. Its only access to it is proximal stimulation of intensity sensors.

In an important sense, this environmental reality is not hidden from the

perceptual hierarchy. Its model of it is presumed reasonably veridical because it in fact accomplishes perceptual control requiring feedback through the environment.

There are two senses of "model" possible here. The perceptual hierarchy may create a model of reality at higher levels of control. It is likely to use language to do this. Like any model, it is imperfect and requires periodic amendment. Because it is maintained at a high level of the control hierarchy, response to conflict is slow. The capacity for amendment is in the control hierarchy that holds the theory, not in the theory itself.

The second sense is that there is a model immanent in the perceptual control hierarchy as a whole. In its capacity to control perceptions whatever the feedback through the environment may be, the perceptual control hierarchy is a kind of reflection of the environment. Like any model, it is imperfect and requires periodic amendment. Unlike many models, it includes this capacity for amendment in itself. Speed of response to conflict varies with the level of the conflict. Call this model<1> and the other sense of "theory" model<2>

I think it was the more primitive sense of the implicit, immanent model<1> that applied when I said "capable of being modelled within the perceptual control hierarchy." Must "boss reality" be capable of being modelled in both senses? Does it have to be able to sit still for its picture to be taken, so to speak?

You, Bill, are assuming the PCT model<2> (theory) as your perspective in talking about knowability. Anything not countenanced in that model you suspect is illusory. And you are using your model<1> to test the model<2>, as indeed are we all.

We have a primitive sense that what appears to be there in the environment is real (naive realism). We can talk ourselves out of this if our model<2> (theory) calls for it. We can also ignore perceptual signals if our expectations say they are not there--if associated error at higher levels is not significant. The two cases seem to me entirely alike.

Suppose there were a physical, mechanistic basis for our primitive sense of the reality of our perceptions. This doesn't entail that this sense be articulate enough for us reliably to distinguish illusion and hallucination, something that we appear not to do. I would base this sense in the continuum of the environment outside the skin with the biochemical and biomechanical environment inside it. (An instance is the role of melatonin receptors in the suprachiasmatic nuclei (in the hypothalamus) in the maintenance of circadian rhythms, which chronobiologists are just beginning to unravel. Nice survey article in _New Scientist_ for 10/26/91 pp. 30-34.) As with other matters to do with neuropeptides and the like, awareness of this would probably have an intuitive and emotional quality rather different from the attention to perceptual signals that usually concerns us. At its peak an awareness of being part of a larger unity, perhaps.

But anything not countenanced in our model<2> or theory (the one we have adopted as our perspective) we suspect is illusory. For good reason: naive realism runs into well-known difficulties. And these intuitive apprehensions are global rather than particular. I think too that this sort of apprehension of the reality of reality is undemonstrable within the perceptual hierarchy. It is only apprehensible to it. Does that

mean it is part of the environmental feedback for the control hierarchy? I don't know.

As we all know, "boss reality" doesn't really sit still for its picture to be taken. A model<2> or theory is possible only by categorization, subsumption with neglect, conventionalization, language. Only a model<1> with its continuous, live tracking can be veridical, and that only in a limited and local sense.

Theories are models<2> of perceptions in our model<1>, which is a model or reflection of reality. All but the specific environmental feedback being tracked and controlled for is hidden from the model<1>. Potentially nothing is hidden from models<2> but their precision and accuracy are suspect. Partly this is because they are constructed using conventionalized social products such as verbalized categories. Partly it is because their responsiveness in the face of aniccha, impermanence, is too slow.

The obvious generalization is to speak of a level-1 model as most local and most accurate, a level-2 model, and so on, up to the models<2> of the system-concept level. Some idealization here about the sequential separation of levels, as we know.

Assuming of course that our model<2> of perceptual control is veridical.

I hope this seems coherent and sensible.

As students, we take your verbalizations about error signals from your comparators as indications for setting reference values in our own. Which we may do, or we may verbalize error signals in turn. A reciprocal process called communication, of course.

This is a test. This is only a test.

Bruce
bn@bbn.com

==========

Date: Tue, 26 Nov 1991 15:53:38 PST
From: marken@AERO.ORG
Subject: Nonlinearity and dimensionality

[From Rick Marken (911126)]

Martin Taylor (911126 17:30) asks a great question about how the PCT model can solve a problem that requires maving away from a goal in order to get closer.

I think you will get a couple more questions (at least) about the problem: for example, are you assuming that A can see B through the wall. Also, since reorganization is not to be the basis of getting to B (a good requirement -- with reorganization my little e. coli can solve it in no time with a very simply perceptual/output model) you must say something more about what A "knows" about getting to B.

Still, I think it would be an interesting challenge to build a PCT model that can solve your little problem. I am also looking forward to getting some suggestions about how it might be done. I think it would be nice to have A be able to get to be regardless fo where A is placed within the wall and regardless (within limits) of the shape of the wall.

What makes the problem interesting to me is that it seems to require a temporal sequence of outputs in order to produce the required sequence of inputs that solve the problem. This might require building a system that can perceive a program of events -- "if wall at x, then move to y, else z". Of course, the program would have to define the input, since A should be able to get to B even when there are changing disturbances -- like changes in the shape of the wall?

Hasta Luego

Rick

****

Richard S. Marken USMail: 10459 Holman Ave
The Aerospace Corporation Los Angeles, CA 90024
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

==========

Date: Tue, 26 Nov 1991 21:13:56 EST
From: goldstein@SATURN.GLASSBORO.EDU
Subject: PCT and suicide

To: CSGnet people
From: David Goldstein
Subject: PCT and suicide
Date: 11/26/91

I don't have the post which raised the issue of PCT and suicide. Sorry not to address the person by name or quote the exact question.

I am a licensed psychologist(NJ/PA) and currently work in an adolescent residential treatment center as the director of clinical services. On a part-time basis, I also have a part-time private practice where most of my patients are adults.

This post gives me an opportunity to think about the relationship between PCT and suicide. I will try to stick to PCT ideas applied rather than merely state what is the clinical lore which guides clinicians everyday.

1. PCT principle: Experience is the important thing for the clinician to focus on, not the "suicidal" action per se.

Implication: The clinician should try to find out what experiences the person is trying to control by the action. The method of levels can be applied here.

2. PCT principle: If the intention remains, a person will find some alternative means of fulfilling the suicidal intention.

Implication: The suicidal intent has to be changed. Simply hiding the means of self-destruction will not work. The person will find alternative means to the end.

Only the suicidal person can change the intent. From exploring why the person wants to suicide and why the person doesn't want

to suicide, the clinician may be able to encourage the person to at least suspend the suicidal goal for a period of time.

3. PCT principle: The suicidal act may be the result of the reorganization system(impulsive) or it may be the result of an acquired control system(way of life).

Implication: PCT suggests that the impulsive suicidal act reflects stonger and more chronic error signals. Therefore, they are the more serious ones. This conclusion may differ from what most clinicians would say. I have to check it out.

4. PCT principle: People want to be successful in the areas of life which are important to them.

Implication: The suicidal act suggests that the person cannot solve the important life problems and needs some help with problem solving. Once the suicidal intent is eliminated, the clinician can turn to help with problem solving.
These are my initial thoughts on the relationship between PCT and suicide. I will be interested in seeing what others come up with.

==========

Date: Tue, 26 Nov 1991 21:15:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: pROGRAM LEVEL

[From Bill Powers (911126.0800)]

Martin Taylor (911125) --

>I think all the levels are executed through a common method, though the
>implementation may be quite different at the different levels.

I haven't received your reprints yet, so I'm agog. What is the common method?

Bruce Nevin (911126) --

>The "partitions" proposed are entirely metatheoretical, just to verify
>that all the bases I know about are being covered and that they are
>being covered in a coherent way. Sometimes it helps to rearrange
>familiar furniture, to see it fresh.

Scratch paper. I understand now. Objection withdrawn.

In this post you've brought up a number of excellent and substantive issues, and what's best, I think I more or less understand what you're talking about. This gives me some hope that we will get somewhere with this before much longer. You have spurred me into unwonted paroxysms of thought.

>Even in the unreduced, informationally complete and explicit base forms
>disclosed by an operator grammar, words do not map in a simple way onto
>nonverbal perceptions.

I agree. When I read or listen to someone, there's a great deal going on in my mind -- in imagination, I suppose -- and it isn't just an echo of the words. If there's a description, I'm constructing a picture based on the words, trying to fill in what the words don't say, to make the scene more complete or to interpolate the action, make sense of it. Words don't

give a very complete picture of anything; the meaning is mostly constructed in imagination. "I noticed Harry at the station yesterday" doesn't mention that it's December, or that you had to drive a car through the snow to get there, or that you only "notice" someone at the station when you're there for some other reason, going or coming, or that there's a puff of steam when you say "That's Harry" (if you did). Or that you're a woman and Harry is your ex-husband who has sworn to kill you.

This is one reason that I don't find abstract discussions very interesting, at least above a certain degree of abstraction. As I listen to the words, all that comes to mind is more words. It's like listening to someone describing how a bridge hand was played. You can get into it, and sense the game that's going on, the strategies, the logic. But it's still all empty, unconnected to life. I want things to be connected to sight and sound and touch and feeling, to motion and action, even to relationships (and categories if there are such things) and logic and principle and system -- but to all of them, not just the upper levels.

The greatest weakness of linguistics is that it treats the language-game as though it's played in an isolation chamber under formal rules with nothing but internal consistency to worry about. It all seems to go on at one level. I hope that you and Martin and jbjg and other linguists on this net are going to do something about that eventually.

You said it yourself:

> ... language is conventionalized, nonverbal perceptions are not.

Yes. Language consists of controllable perceptions, and when we're concerned with language itself rather than what it means, the perceptions are controlled in formal and conventionalized ways that don't necessarily mesh with the rest of experience. So even if language begins as simple naming and denoting in relation to the nonverbal world, it can turn into a game played for its own sake, like mathematics. Like mathematics, language can be misused for this same reason. If the game of linguistics is played for its own sake, if it is played IN language and ABOUT language, it leaves the rest of experience to tag along behind like a barrage balloon on a string.

A hint of where I'm going: conventions are program-level stuff.

==========

All this came from a morning and afternoon spent writing various responses to your post. A lot of words. I spent a lot of time puzzling over your Heretical diagram:

Heretical PCT

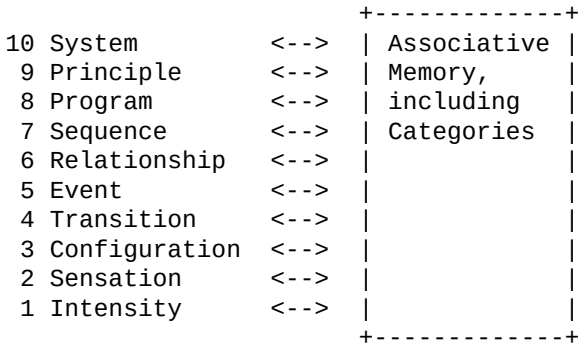

I was going to point out that this diagram refers to a set of categories, including the category called "associative memory." It was only after I gave up on this for a few hours that a different version of this diagram popped into my head, after a few lurches in other directions:

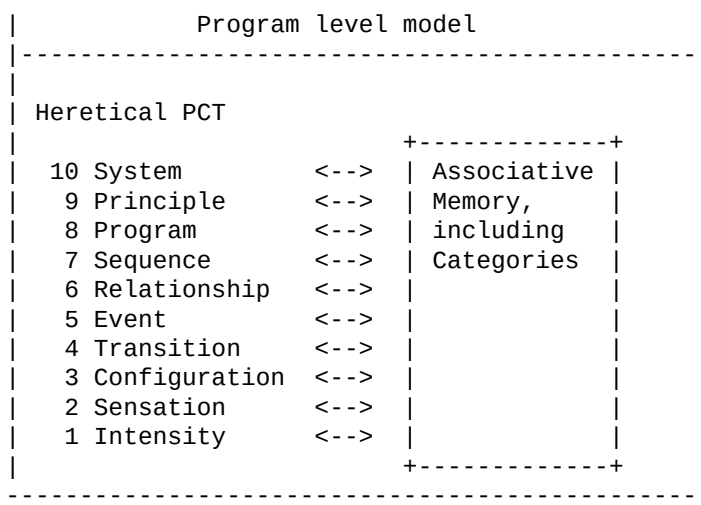

This is what I meant in the preamble when I said that "It all seems to go on at one level." The above model (or rather the logical and formal operations implied by the picture of the model) shows an ordered sequence of names of categories of perception, interacting in some way with "associative memory, including categories." So the "categories" included in associative memory are not the categories referred to by the list of names on the left.

You ventured: "What if categorization is distributed throughout the hierarchy, just as associative memory is, and just as you say language is?"

Well, it is, but not in the sense you probably meant. The operation by which a set of intensity signals is combined to produce a weighted sum that we experience as a sensation is not the categorizing process, but weighted summation. But *we perceive all sensations as belonging to a recognizeable single category: sensations.* In the same way, we can recognize a world of thinglike configurations, a world of events, a world of relationships. Each world is a recognizeable category, and within each of these worlds, we can perceive smaller categories. A nonverbal category-perceiving level lets us know that elements of each lower level belong together. That's because we have put them together, and we have done that because the result is something coherently controllable.

The point I'm working up to is that at the program level, we also have something called categories. They result from a logical process of categorization, quite different from the process of perceptual categorization. Logical categorization requires some basis on which to deduce that a set of perceptions satisfies the definition of a category.

It's at the logical, not the perceptual, level of categorization that "quadrilaterals" exist. A logical quadrilateral is "a plane figure having four sides and four angles." The terms plane, figure, four, side, and angle are names of categories that are conceivably perceivable: you can look at something and see if it is a figure, if it seems plane, if there

are sides, if the number is four, and if there are angles. But you can't look at anything and see that it is a "quadrilateral." To see that a quadrilateral exists, you must tally the sides and angles, note that the figure is plane, and compare the observations with the definition. This might be plainer if I had said "duodecahedron."

Consider the class of relationships called "between." In the array ABC, B is between A and C. This is also true of the arrays AB C and A BC. We can see instantly that B is between A and C in any of these cases, without having to reason about it.

At the logical level, B is between A and C if all three are on a straight line, B and C are to the right of A, and B is closer to A than is C: that is, the distance B - A is less than the distance C - A. Notice that in order to decide logically whether to class this relationship of B to A and C as "between," it's necessary to examine the evidence and see if it fits the definition. This is not the same process by which we recognize the "betweenness" of B in looking at ABC. It is done by symbolic reasoning, not by wetware.

This, by the way, is why graphics programs are so terrifically slow. They have to work out things by logic that the human brain sees at a blow of the eye: whether a point is inside or outside a closed figure, for example.

So, unfortunately, we use the word "category" for two quite different things: one is a matter of direct, if learned, perception: the other is a matter of symbolic reasoning.

An interlude:

>* The perception "triangle" can be used to refer to any of the other
> perceptions /\ that CW2 specifies.
>   /__\
>
>But wait a minute! Why does the linkup with other members amount to
>reference in the case of CW2 but not in the case of CV1 or CV2? Why the
>privileged status so that "triangle" can be used to refer to a triangle
>and not vice versa?

I didn't say that "triangle" refers to /_\. What I said was that either "triangle" or /_\ will give rise to the same category perception, not that they give rise to each other. Either one, therefore, will serve for satisfying the demand that CW2 be perceived.

It may be that hearing "triangle" results immediately in imagining /_\, and that seeing /_\ results immediately in imagining "triangle," but that is association, not categorization. It is not necessarily two-way. It does not amount to an assertion that these perceptions are equivalent. Seeing a car and thinking "accident" doesn't mean equivalence: it means only association. I think that's an entirely different phenomenon. Real, but different.

In trying to reconcile my notation, you conclude

>* The relationship perception "triangle" is a word in the object
> language rather than a metalanguage word (CV2).
>
>* The relationship perception "noun" is a word in the metalanguage

> rather than being an object-language word (CV3).

But this assumes that at the category level as I conceive it there are such things as metalanguage elements. When "noun" and "word" are collected into one category, this means only that the same category signal will be generated on hearing or reading either word. Only if we ask WHY such a category might have been developed do we raise the question of metalanguage. And that gets us to the levels that use categories: the levels of reason, in which we can lump sequences and programs.

Using a metalanguage is similar to going, in a dictionary, from the definition to the word. Rather than denoting the entire set of words designating a "person, place, thing, state, or quality" (and enumerating the names of members of those categories), we can say "noun." Whenever we use this substitution, we are actually manipulating all the members of this metacategory at once: in principle, if we've done this correctly, any statement involving "noun" could have the entire set of noun-words substituted for "noun" in the same assertion. This is the logic of a metalanguage, isn't it? In making an assertion about an entire set, we substitute the name of the set for the set: "Let 'noun' = {..........}.

When we say "plural noun," we are specifying two metacategories at once: the set of all words for a "person, place, thing, state, or quality" and the set of all words for multiple instances of things. Only plurals of nouns (and not plural pronouns) will satisfy both metacategories. (I get my definitions from dictionaries, not linguists, so pipe down).

However, we certainly do not ever actually make these substitutions. One result is that using language and metalanguage, we can say things about metaterms that are not in fact true of all instances of such terms. It's easy to overlook cases like the plural noun in "Siamese twins." It isn't true that plural nouns refer to separable things. Most of them, yes. But you can't separate trousers with scissors. Metalanguage is subject to the same defects as any form of generalization.

And these defects are caused by deviations of logic from experience. This brings me back to my main point.

I'm suggesting that we consider seriously the possibility that all of linguistics (or whatever part you say) is concerned with the program level. I was mistaken in accusing you of occupying the category level. Our differences concerning categories are the difference between thinking of perceptual categories and thinking of logical categories. What I think of and have diagrammed as category operations are in the realm of direct experience. It is the program level that, through subsituting names of classes for perceptual classes containing all kinds of perceptions, and substituting names of collections of words for perceptual classes containing words, creates metaterms and manipulates them according to rules appropriate to them.

I agree with your criticisms of my proposed naming process -- something better is needed than merely including the names in the classes. The kind of naming we need makes the formation of logical classes as easy as saying "Let x = any object with three legs." The elements of that statement are perceptual categories of varying sizes. But the substitution operation is the naming process. I don't know where to put it (pause for low humor). Forming a new perceptual class, which you

experience at a blow of the eye, is not that quick or easy.

At the program level, I propose, we run a model of the world that can, but doesn't necessarily, include simulations of all the lower levels of perception in the perceptual hierarchy. This program-type model can have as many levels as we can handle mentally (exactly as many). The simulation is carried out in terms of signals functioning as symbols representing perceptual classes, operated on by processes we could call, generically, "rules." Just as in electronic computers, it is possible for this program level to construct an emulation of itself (a direct analogy would be a debugging program that actually emulates the registers and operations of the hardware). The problem of infinite regression in such self-emulations has long been settled: it doesn't happen.

I am offering this proposal because I see in it the possibility of stating and correcting the difference between you and me in our approach to language, and to behavior in general. I think it offers a model within which everything you and I have been saying can remain mostly true by showing that apparent conflicts result from differences in the level of organization to which we attribute various phenomena of language.

==========

Date: Wed, 27 Nov 1991 13:12:21 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: Thanksgiving fare

Thank you, Bill, for that long and delicious-looking post.
I will savor it over the holiday weekend.

Best of all to everyone

Bruce

==========

Date: Wed, 27 Nov 1991 12:57:38 EST
From: "CHARLES W. TUCKER" <N050024@UNIVSCVM.BITNET>
Subject: Recent posts on levels, language, epistemology and related topics

[FROM CHUCK TUCKER 911127.1358]

TO: Bill, Bruce, Martin, Judd, Wayne, Maler, Rick and a few others I can't think of at the moment.

The recent posts on the topics mentioned above have been excellent. I have begun to take them off the file and transform then into WS files to edit them for use in my course next Spring. It is one experience to find a book that someone has labored over for several months (Peckham said that he wrote his book ART AND PORNOGRAPHY is 39 days) or for many years that is cogent, intelligent, relevant and useful BUT to find such writing that you know is done in several hours or at the most several days means, to me, that I am in the company of some very special people who one rarely finds in several lifetimes. I compare these posts favorably with the works of my most repected colleagues: Dewey, Mead, Bentley, Pierce, Rorty, Goffman, Bateson, Stone and several of my friends on CSGNET and not on the list. So, since this is the season set aside to be thankful, I want to thank all of you for your fine work and to all of my other colleagues on CGSNET I thank you for making this electronic experience so useful.

I only have one little request: STOP BEING SO USEFUL SO I CAN CATCH MY BREATH.

Best regards to all,

Chuck

==========

Date: Wed, 27 Nov 1991 13:39:53 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Nonlinearity and dimensionality

[Martin Taylor 911127 13:00]
(Rick Marken 911126)
>
>Martin Taylor (911126 17:30) asks a great question about how the PCT model
>can solve a problem that requires maving away from a goal in order to
>get closer.
>
>I think you will get a couple more questions (at least) about the problem:
>for example, are you assuming that A can see B through the wall. Also,
>since reorganization is not to be the basis of getting to B (a good
>requirement -- with reorganization my little e. coli can solve it in no time
>with a very simply perceptual/output model) you must say something
>more about what A "knows" about getting to B.
>
Assume that A and B are simply current and reference positions in a 2-D field,
and that the total configuration (where A and B are, and where wall exists)
is accessible to the sensors.

My question isn't whether a program can be developed to solve this problem,
but whether ANY hierarchic system can solve it if the system is constrained
to be built with elemental control systems that have (a) scalar (unidimensional)
values for percept and reference, and (b) monotonic relationship between
output and (error = percept - reference).

Actually, I'm assuming the answer is Yes, such a hierarchy can be built,
because the situation is so common in everyday life, and Bill is no fool.
Bill insists on the scalar variable criterion in the ECS, and I infer from
the discussions that the error-output relation is assumed to be monotonic
(indeed, if it were not, the ECS would automatically find unidimensional
problems it could not solve because it would get caught in local minima).
So, PCT has been dealing very nicely with complex problems for a long time
under these constraints, which means that it should be able to solve what
I see as a minimalist version of the more general problem of planning and
choice in such a hierarchy. But I couldn't myself solve the problem. I need
vectorial ECSs rather than scalars.

Martin Taylor

==========

Date: Tue, 26 Nov 1991 21:15:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: pROGRAM LEVEL

[From Bill Powers (911126.0800)]

Martin Taylor (911125) --

>I think all the levels are executed through a common method, though the
>implementation may be quite different at the different levels.

I haven't received your reprints yet, so I'm agog. What is the common
method?

Bruce Nevin (911126) --

>The "partitions" proposed are entirely metatheoretical, just to verify
>that all the bases I know about are being covered and that they are
>being covered in a coherent way. Sometimes it helps to rearrange
>familiar furniture, to see it fresh.

Scratch paper. I understand now. Objection withdrawn.

In this post you've brought up a number of excellent and substantive issues, and what's best, I think I more or less understand what you're talking about. This gives me some hope that we will get somewhere with this before much longer. You have spurred me into unwonted paroxysms of thought.

>Even in the unreduced, informationally complete and explicit base forms
>disclosed by an operator grammar, words do not map in a simple way onto
>nonverbal perceptions.

I agree. When I read or listen to someone, there's a great deal going on in my mind -- in imagination, I suppose -- and it isn't just an echo of the words. If there's a description, I'm constructing a picture based on the words, trying to fill in what the words don't say, to make the scene more complete or to interpolate the action, make sense of it. Words don't give a very complete picture of anything; the meaning is mostly constructed in imagination. "I noticed Harry at the station yesterday" doesn't mention that it's December, or that you had to drive a car through the snow to get there, or that you only "notice" someone at the station when you're there for some other reason, going or coming, or that there's a puff of steam when you say "That's Harry" (if you did). Or that you're a woman and Harry is your ex-husband who has sworn to kill you.

This is one reason that I don't find abstract discussions very interesting, at least above a certain degree of abstraction. As I listen to the words, all that comes to mind is more words. It's like listening to someone describing how a bridge hand was played. You can get into it, and sense the game that's going on, the strategies, the logic. But it's still all empty, unconnected to life. I want things to be connected to sight and sound and touch and feeling, to motion and action, even to relationships (and categories if there are such things) and logic and principle and system -- but to all of them, not just the upper levels.

The greatest weakness of linguistics is that it treats the language-game as though it's played in an isolation chamber under formal rules with nothing but internal consistency to worry about. It all seems to go on at one level. I hope that you and Martin and jbjg and other linguists on this net are going to do something about that eventually.

You said it yourself:

> ... language is conventionalized, nonverbal perceptions are not.

Yes. Language consists of controllable perceptions, and when we're concerned with language itself rather than what it means, the perceptions are controlled in formal and conventionalized ways that don't necessarily mesh with the rest of experience. So even if language begins as simple naming and denoting in relation to the nonverbal world, it can turn into a game played for its own sake, like mathematics. Like mathematics, language can be misused for this same reason. If the game of linguistics

is played for its own sake, if it is played IN language and ABOUT language, it leaves the rest of experience to tag along behind like a barrage balloon on a string.

A hint of where I'm going: conventions are program-level stuff.

==========

All this came from a morning and afternoon spent writing various responses to your post. A lot of words. I spent a lot of time puzzling over your Heretical diagram:

Heretical PCT

```
                        +-------------+
10 System         <-->  | Associative |
 9 Principle      <-->  | Memory,     |
 8 Program        <-->  | including   |
 7 Sequence       <-->  | Categories  |
 6 Relationship   <-->  |             |
 5 Event          <-->  |             |
 4 Transition     <-->  |             |
 3 Configuration  <-->  |             |
 2 Sensation      <-->  |             |
 1 Intensity      <-->  |             |
                        +-------------+
```

I was going to point out that this diagram refers to a set of categories, including the category called "associative memory." It was only after I gave up on this for a few hours that a different version of this diagram popped into my head, after a few lurches in other directions:

```
 ------------------------------------------------
|            Program level model                 |
|------------------------------------------------|
|                                                |
| Heretical PCT                                  |
|                         +-------------+        |
|  10 System         <--> | Associative |        |
|   9 Principle      <--> | Memory,     |        |
|   8 Program        <--> | including   |        |
|   7 Sequence       <--> | Categories  |        |
|   6 Relationship   <--> |             |        |
|   5 Event          <--> |             |        |
|   4 Transition     <--> |             |        |
|   3 Configuration  <--> |             |        |
|   2 Sensation      <--> |             |        |
|   1 Intensity      <--> |             |        |
|                         +-------------+        |
------------------------------------------------
```

This is what I meant in the preamble when I said that "It all seems to go on at one level." The above model (or rather the logical and formal operations implied by the picture of the model) shows an ordered sequence of names of categories of perception, interacting in some way with "associative memory, including categories." So the "categories" included in associative memory are not the categories referred to by the list of names on the left.

You ventured: "What if categorization is distributed throughout the hierarchy, just as associative memory is, and just as you say language

is?"

Well, it is, but not in the sense you probably meant. The operation by which a set of intensity signals is combined to produce a weighted sum that we experience as a sensation is not the categorizing process, but weighted summation. But *we perceive all sensations as belonging to a recognizeable single category: sensations.* In the same way, we can recognize a world of thinglike configurations, a world of events, a world of relationships. Each world is a recognizeable category, and within each of these worlds, we can perceive smaller categories. A nonverbal category-perceiving level lets us know that elements of each lower level belong together. That's because we have put them together, and we have done that because the result is something coherently controllable.

The point I'm working up to is that at the program level, we also have something called categories. They result from a logical process of categorization, quite different from the process of perceptual categorization. Logical categorization requires some basis on which to deduce that a set of perceptions satisfies the definition of a category.

It's at the logical, not the perceptual, level of categorization that "quadrilaterals" exist. A logical quadrilateral is "a plane figure having four sides and four angles." The terms plane, figure, four, side, and angle are names of categories that are conceivably perceivable: you can look at something and see if it is a figure, if it seems plane, if there are sides, if the number is four, and if there are angles. But you can't look at anything and see that it is a "quadrilateral." To see that a quadrilateral exists, you must tally the sides and angles, note that the figure is plane, and compare the observations with the definition. This might be plainer if I had said "duodecahedron."

Consider the class of relationships called "between." In the array ABC, B is between A and C. This is also true of the arrays AB   C and A BC. We can see instantly that B is between A and C in any of these cases, without having to reason about it.

At the logical level, B is between A and C if all three are on a straight line, B and C are to the right of A, and B is closer to A than is C: that is, the distance B - A is less than the distance C - A. Notice that in order to decide logically whether to class this relationship of B to A and C as "between," it's necessary to examine the evidence and see if it fits the definition. This is not the same process by which we recognize the "betweenness" of B in looking at ABC. It is done by symbolic reasoning, not by wetware.

This, by the way, is why graphics programs are so terrifically slow. They have to work out things by logic that the human brain sees at a blow of the eye: whether a point is inside or outside a closed figure, for example.

So, unfortunately, we use the word "category" for two quite different things: one is a matter of direct, if learned, perception: the other is a matter of symbolic reasoning.

An interlude:

>* The perception "triangle" can be used to refer to any of the other
> perceptions /\ that CW2 specifies.
>   /__\
>

>But wait a minute! Why does the linkup with other members amount to
>reference in the case of CW2 but not in the case of CV1 or CV2? Why the
>privileged status so that "triangle" can be used to refer to a triangle
>and not vice versa?

I didn't say that "triangle" refers to /_\. What I said was that either "triangle" or /_\ will give rise to the same category perception, not that they give rise to each other. Either one, therefore, will serve for satisfying the demand that CW2 be perceived.

It may be that hearing "triangle" results immediately in imagining /_\, and that seeing /_\ results immediately in imagining "triangle," but that is association, not categorization. It is not necessarily two-way. It does not amount to an assertion that these perceptions are equivalent. Seeing a car and thinking "accident" doesn't mean equivalence: it means only association. I think that's an entirely different phenomenon. Real, but different.

In trying to reconcile my notation, you conclude

>* The relationship perception "triangle" is a word in the object
> language rather than a metalanguage word (CV2).
>
>* The relationship perception "noun" is a word in the metalanguage
> rather than being an object-language word (CV3).

But this assumes that at the category level as I conceive it there are such things as metalanguage elements. When "noun" and "word" are collected into one category, this means only that the same category signal will be generated on hearing or reading either word. Only if we ask WHY such a category might have been developed do we raise the question of metalanguage. And that gets us to the levels that use categories: the levels of reason, in which we can lump sequences and programs.

Using a metalanguage is similar to going, in a dictionary, from the definition to the word. Rather than denoting the entire set of words designating a "person, place, thing, state, or quality" (and enumerating the names of members of those categories), we can say "noun." Whenever we use this substitution, we are actually manipulating all the members of this metacategory at once: in principle, if we've done this correctly, any statement involving "noun" could have the entire set of noun-words substituted for "noun" in the same assertion. This is the logic of a metalanguage, isn't it? In making an assertion about an entire set, we substitute the name of the set for the set: "Let 'noun' = {..........}.

When we say "plural noun," we are specifying two metacategories at once: the set of all words for a "person, place, thing, state, or quality" and the set of all words for multiple instances of things. Only plurals of nouns (and not plural pronouns) will satisfy both metacategories. (I get my definitions from dictionaries, not linguists, so pipe down).

However, we certainly do not ever actually make these substitutions. One result is that using language and metalanguage, we can say things about metaterms that are not in fact true of all instances of such terms. It's easy to overlook cases like the plural noun in "Siamese twins." It isn't true that plural nouns refer to separable things. Most of them, yes. But you can't separate trousers with scissors. Metalanguage is subject to the same defects as any form of generalization.

And these defects are caused by deviations of logic from experience. This brings me back to my main point.

I'm suggesting that we consider seriously the possibility that all of linguistics (or whatever part you say) is concerned with the program level. I was mistaken in accusing you of occupying the category level. Our differences concerning categories are the difference between thinking of perceptual categories and thinking of logical categories. What I think of and have diagrammed as category operations are in the realm of direct experience. It is the program level that, through subsituting names of classes for perceptual classes containing all kinds of perceptions, and substituting names of collections of words for perceptual classes containing words, creates metaterms and manipulates them according to rules appropriate to them.

I agree with your criticisms of my proposed naming process -- something better is needed than merely including the names in the classes. The kind of naming we need makes the formation of logical classes as easy as saying "Let x = any object with three legs." The elements of that statement are perceptual categories of varying sizes. But the substitution operation is the naming process. I don't know where to put it (pause for low humor). Forming a new perceptual class, which you experience at a blow of the eye, is not that quick or easy.

At the program level, I propose, we run a model of the world that can, but doesn't necessarily, include simulations of all the lower levels of perception in the perceptual hierarchy. This program-type model can have as many levels as we can handle mentally (exactly as many). The simulation is carried out in terms of signals functioning as symbols representing perceptual classes, operated on by processes we could call, generically, "rules." Just as in electronic computers, it is possible for this program level to construct an emulation of itself (a direct analogy would be a debugging program that actually emulates the registers and operations of the hardware). The problem of infinite regression in such self-emulations has long been settled: it doesn't happen.

I am offering this proposal because I see in it the possibility of stating and correcting the difference between you and me in our approach to language, and to behavior in general. I think it offers a model within which everything you and I have been saying can remain mostly true by showing that apparent conflicts result from differences in the level of organization to which we attribute various phenomena of language.

==========

Date: Wed, 27 Nov 1991 15:05:00 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: epistemology

[Martin Taylor 911127 15:00]
(Bruce Nevin 911127 0743)

>
>As we all know, "boss reality" doesn't really sit still for its picture
>to be taken. A model<2> or theory is possible only by categorization,
>subsumption with neglect, conventionalization, language. Only a
>model<1> with its continuous, live tracking can be veridical, and that
>only in a limited and local sense.
>

>Theories are models<2> of perceptions in our model<1>, which is a model
>or reflection of reality. All but the specific environmental
>feedback being tracked and controlled for is hidden from the model<1>.
>Potentially nothing is hidden from models<2> but their precision and
>accuracy are suspect.

Perhaps your model<1-2> distinction lies at the heart of our misunderstanding of each other's approach to language. I think I have been considering language in action -- the model <1> approach -- whereas you have been dealing with language as analyzed -- the model <2> approach, which has all the limitations you ascribe to it. Those limitations are why I keep saying things like: no satisfactory model of natural language has been developed in 2000 years of attempts, and there is no reason to believe that a symbolic description of language is possible. If you do categorize, subsume with neglect, conventionalize, and thus permit yourself to use language, then you can describe a skeletalized version of natural language quite well.

Do you agree that this model 1-2 distinction is the reason we seem to have been talking past each other?

It's really a different issue, but I would like to raise a flag to signal my objection to your claim that Model <1> can be veridical "only in a limited and local sense." It is true that mathematical theories of physical "reality" take us a lot further than intuitive physics in predicting the behaviour of the world, but it is not so clear that this is true for the less simple sciences. Physics is, after all, the only science so simple that the most intelligent humans have a reasonable hope of understanding some of it. But we do behave reasonably successfully in the much more complex world of nutritious and poisonous foods, friends and enemies, and so on, for which linguistically (e.g.mathematically) based models do a lousy job.

Martin

==========

Date: Wed, 27 Nov 1991 14:23:26 -0600
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Critiques of PCT

[from Gary Cziko 911127.1410]

I will be teaching a course next semester (if I can get any students to enroll) called "Psychological Theories Applied to Education" in which Perceptual Control Theory will be the major focus.

One thing I want to do is contrast PCT with other psychological theories and perspectives. And I suppose that a particularly interesting way to do this would be to look at critical reviews of PCT articles (and chapters or whatever) submitted by PCT types for publication. Rick Marken has indicated that he has a sizable collection of such reviews (I'd particularly like to see his reviews of the E. coli paper) and I suspect that Tom Bourbon (recent paper with Powers), Bill Powers (I'd love to see the BBS stuff) and Wayne Hershberger must have some "good" ones as well. (I'm expecting my own soon from my first attempt to publish a PCT article).

How could I get a hold of these for my class? Would people be willing to mail me copies? Perhaps it would interesting and fun if having such reviews to post his or her favorite reviews to CSGnet with some commentary. This could generate some interesting discussion of some topics central to PCT and might make for an interesting issue of _Closed Loop_. Whaddya say Bill, Rick, Tom, Wayne, et al.?--Gary

==========

Gary A. Cziko
University of Illinois
Educational Psychology
210 Education
1310 South 6th Street
Champaign, Illinois 61820-6990
USA

Telephone: (217) 333-4382
FAX: (217) 244-0538
Internet: g-cziko@uiuc.edu (1st choice)
Bitnet: cziko@uiucvmd (2nd choice)
N9MJZ

==========

==========

Date: Wed, 27 Nov 1991 14:30:24 PST
From: marken@AERO.ORG
Subject: Critiques of PCT

[From Rick Marken (911127)]

Gary Cziko (911127.1410) says:

>One thing I want to do is contrast PCT with other psychological theories
>and perspectives. And I suppose that a particularly interesting way to do
>this would be to look at critical reviews of PCT articles (and chapters or
>whatever) submitted by PCT types for publication. Rick Marken has
>indicated that he has a sizable collection of such reviews (I'd
>particularly like to see his reviews of the E. coli paper) and I suspect
>that Tom Bourbon (recent paper with Powers), Bill Powers (I'd love to see
>the BBS stuff) and Wayne Hershberger must have some "good" ones as well.
>(I'm expecting my own soon from my first attempt to publish a PCT article).

What a fun idea. I will go over my prodigious file of bad reviews (I think I might, indeed, have the dubious distinction of having more such reviews than any other PCTer) and mail you copies of the ones that seem the best -- especially the reviews of the 'Selection of consequences'(e coli) paper). I guess it's appropriate to do this on Thanksgiving -- It's a way for me to give thanks that I found PCT and didn't end up like the reviewers.

Hasty Bagels

Rick

****

Richard S. Marken
The Aerospace Corporation
Internet:marken@aerospace.aero.org
213 336-6214 (day)
213 474-0313 (evening)

USMail: 10459 Holman Ave
Los Angeles, CA 90024

==========

Date: Wed, 27 Nov 1991 21:18:27 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: John harp;psycho flux;scalar control

[From Bill Powers (911127.1000)]

Rick Marken (911125) --

I'm impressed all over again by the beauty and simplicity of your experiment with speed of perception at the different levels.

Bruce Nevin (911126) --

>   Mary plays flute and John [plays] harp.

There's a hint here about the mechanics of recognizing a sentence. I'll bet nobody would say this sentence:

  Mary plays flute and John who has been studying music from an early age but is progessing more slowly in developing proficiency harp.

When "plays" is heard the first time, the perception lingers, slowly fading away. If it's still sufficiently strong at the time it's needed again, it's perceived as already present in perception so there's no error to cause it to be said again. This is known as Powers' Persistence of Perception Postulate, or PPPP.

I'm going to wait for the rest until you've commented on my last post -- my spotty sampling times are getting me out of synch, so I'm always sending a new comment that pertains to the one before your last one. Too confusing.

Wayne Hershberger (911126) --

>I would say that the natural order immanent in the psychophysical
>flux is realized both in the form of perceptual stuff and in the
>form of conceptual stuff.

You have picked the first of my representations.

So does your model looks like this:

A:

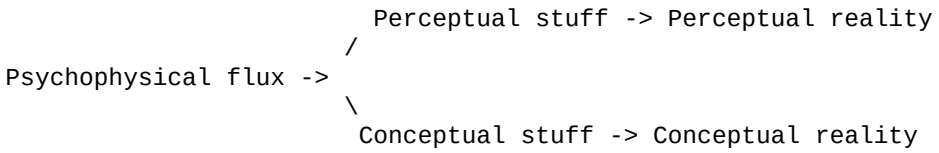

?

Or would you draw it this way:

B:

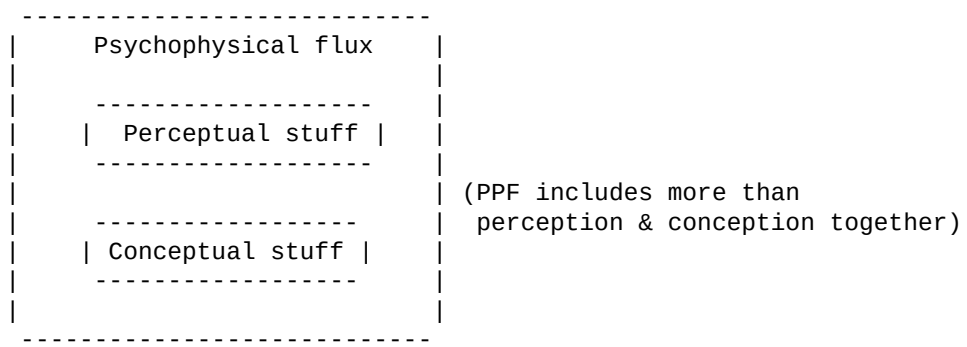

?

Or this way:

C:

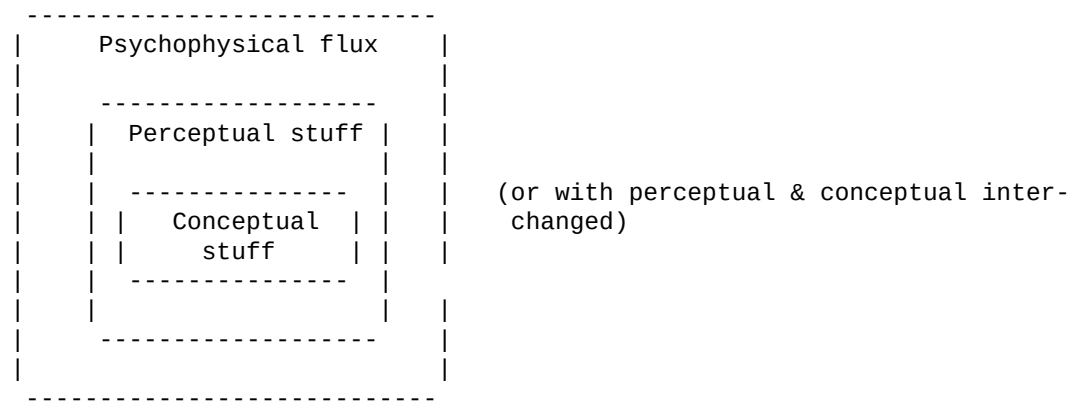

?
Or this way:

D:

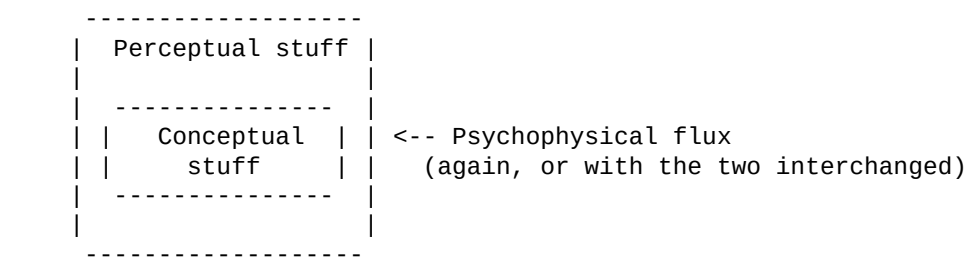

Or this way:

E:

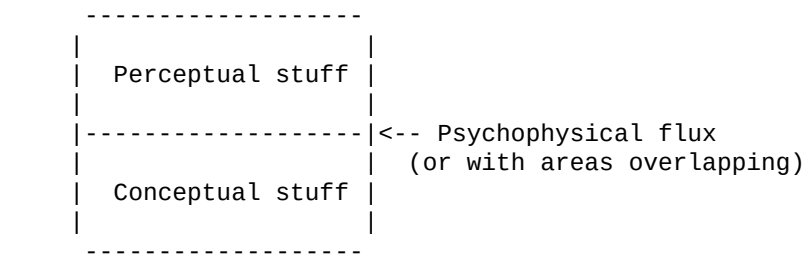

Or some other way?

Martin Taylor (911126) --

> What I would like to see is a discussion of a simple situation in which
> I am unable to see how a set of scalar ECSs can perform correctly.
> I think that a carefully considered explanation of this situation would
> go a long way to clearing up some of the discussions we have been
> having about language.
>
> Consider this picture (alpha-graphic pictures seem fashionable these
> days):
>
>
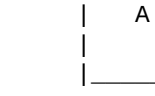

>   B

You ask how I would design A to reach B using only control systems that control one-dimensional perceptual signals.

As it happens, this problem is solved in the Gather (formerly Crowd) program I wrote for Clark McPhail and Chuck Tucker. Many setups are possible, but in the relevant one, a single moving individual seeks a goal-position and simultaneously avoids collisions with other stationary or moving people. When there are many stationary people on the field (randomly distributed), they often form barriers and pockets (which could easily be shaped and oriented as in the diagram above). If the moving individual encounters a barrier or pocket, the control systems that avoid collisions slow the speed and begin changing direction as required. The speed never falls below some minimum, so the person always keeps moving. If necessary, the moving person will backtrack, moving away from the goal, and find a way around the barrier or out of the pocket. Sometimes a closed path results, so the person wanders mindlessly around and around seeking a way out. The direction control systems, however, contain a small amount of random noise added to the error signals. The result is that the closed paths never repeat exactly, so it usually happens that a "trapped" individual hits on the way out (there is a way in, so there is a way out). This can take a long time, because there isn't much noise. There have been cases where I didn't feel like staying up all night to see if the individual would ever break out, so I guess there are truly inescapable traps. I suppose one could build in a "reorganizing system" so when the goal-seeking error persisted for a long enough time, the amount of random noise would begin increasing.

But a pocket like the one you drew above would be a cinch -- nobody would get stuck in that one. The "person" would just turn toward the open end and go around one side, moving therafter on a curve leading directly to the goal position at B without colliding with the wall's outer side.

General principles of operation:

The control systems can sense the "proximity" of objects on the playing field, and can sense the proximity of the goal and its direction relative to the direction of travel. When a direction-controlling system detects an error, it contributes to the CHANGE in direction of travel (the actual direction of travel is the time-integral of the total error). Thus error determines mainly the radius of curvature of the path. When a proximity-controlling system senses an error, it slows the rate of travel. This also shortens the radius of curvature.

Proximity is defined as an inverse-square function of distance to an object. The closer the object, the less distance there is and the more proximity is perceived (up to a maximum limit). This is analogous to perception of distance on the basis of retinal area of an image. An inverse-fourth-power function would also work, representing radar or sonar sensors.

Note that touch sensors would also work: they would simply provide a steeper "proximity function."

Avoidance of obstacles:

The avoidance system uses two sensors, one which detects the sum of all proximities to the left of the direction of travel, and another which detects the sum of all proximities to the right. This is a "vision"

system like a fly's, in which the sensors accept inputs from a wide angle. Each sensor's sensitivity drops as the angle points more than 90 degrees toward the rear. In the program, the distance to every object left of the direction of travel is computed, the proximity of each is looked up in a table, and all proximities are summed to provide the "PROXIMITY LEFT" perceptual signal. The same is done on the right. This is how you have to do things with a computer simulation: an actual sensor would simply detect total retinal area occupied by objects, in parallel. Objects are assumed, of course, to be all of the same size, although it wouldn't be unrealistic for the person to skirt larger objects at a greater distance.

So we have scalar Pl and Pr perceptual signals: proximity left, proximity right.

These avoidance systems seek an adjustable amount of proximity, usually set to somewhat less than the maximum. The left and right systems are one-way controllers: they react when the sensed proximity is greater than the reference proximity ("too close"), but when the sensed proximity falls below the reference proximity their error signals remain at zero. So if there is too much proximity either to the left or to the right, an appropriate error signal is generated. In the actual program I have collapsed these detailed relationships into a single comparator, plus a little logic. There are several ways to do this that work.

In each time-increment during the simulation, the error signals from the two proximity sensors are added to the angle (in laboratory space) of travel, with a factor to set sensitivity to error. As the person approaches a point-obstacle, the proximity will eventually rise past the reference proximity. If the left proximity is the one that rises (or it rises more than the right one --the fields of view can overlap), the error signal (via a sensitivity factor) will subtract from the angle of travel on each iteration, causing the path to veer toward the right until the proximity has fallen to or below the reference proximity. Similarly, an approach that makes right proximity rise causes a change of direction to the left. If the path happens to be aimed directly at the point-obstacle, no change of path would take place, so a small amount of random noise is added to the error signal on each iteration. This is sufficient to cause the path to start changing one way or the other and permit normal avoidance. Two such systems approaching each other sometimes act like people on a sidewalk, dithering back and forth.

A second avoidance control system perceives proximity, left plus right. When this proximity exceeds a reference level, the resulting (one-way) error signal subtracts from the velocity that is integrated on each iteration to produce motion. The result is that when the person nears the obstacle the velocity slows, which helps decrease the radius of curvature for collision courses.

With no goal-seeking in effect, the person equipped with only this avoidance system will wander around through the crowd and eventually go off the edge of the screen. It won't collide with anyone. It would escape from the trap above quite quickly, but wouldn't go to the goal at B.

Goal-seeking:

Two systems steer the person toward a goal-position on the screen. The direction control system is simple: it senses the angle between direction of travel and the direction of the line of sight to the goal. This goal direction signal is compared with the goal direction reference

signal (normally set to zero for a straight-in approach), and the resulting error signal adds to the angle of travel on each iteration (through a sensitivity factor). If the reference signal is set to a nonzero value, the approach to the goal will be along a spiral. In the program it is fixed at zero.

The second system is the speed control system. It compares the sensed proximity to the goal with the reference proximity (usually set to the maximum, meaning zero distance), and adds the resulting error signal to the velocity signal that is integrated to move the person. The scaling factor determines how fast the person will move at minimum proximity (maximum distance). As the person approaches the goal, the proximity error decreases, so the speed decreases -- at first slowly, then more and more rapidly as proximity to the goal increases. At the goal, the person stops.

Combined avoidance and goal-seeking behavior:

When far from the goal-position and any obstacle, the person will quickly find a path aimed directly at the goal. The goal-direction error signal will be zero, so the path will be a straight line. If this were the only control system, the person would run into the wall in the diagram above, and be physically stopped.

As an obstacle nears, however, the avoidance systems will detect significant proximity and the path will start to curve. At first, the goal direction control system will resist this curvature. But its error rises linearly with angle, while the avoidance systems detect a proximity that rises more and more steeply as the obstacle gets closer. At some distance from the obstacle, the avoidance error signal will have more effect on direction than does the goal-direction error signal, and the path will curve. The closer the person gets to the obstacle, the sharper the curvature of the path will be. Once the person is mostly past the obstacle, the avoidance error signal will fall back to zero and the goal-seeking system will turn the path once more toward the goal.

When the person passes between two obstacles, the path avoids both of them as the left and right systems both experience error. And because of the increased proximity, the person slows down and squeeeezes through. Very realistic when you get the parameters set just right.

So the result is that the person will find a way through the crowd and eventually get to the goal -- usually, quite efficiently. But this is a dumb system; it doesn't look for the shortest way to the goal, or "plan its path" at all. So it often overlooks shorter paths, because taking them might momentarily cause the threat of a collision (particularly when there are groups of obstacles all adding their proximities together). Modifications of the proximity detectors might be made to give groups of obstacles at the same distance a lesser effect, leading to better choices of paths, but I haven't experimented with that. The point was not to get optimum behavior, but to get reasonably intelligent-looking behavior out of the simplest possible set of control systems.

Generalities:

As presently set up, the goal-seeking and avoidance systems operate in conflict. By making either system very "tight" -- high loop gain -- you can negate the effect of the other. It would probably be better to make them hierarchical so one of them has its effects by altering the reference signal of the other. I haven't gone into that. Presumably,

Clark McPhail, who now has found a smart programmer willing to take over further developments of this model, will look into such possibilities eventually.

One important fact about this stimulated person is that there is no decision-making involved. There is nothing that says "Gee, this path looks pretty good, and so does that one -- which way shall I go?" There is no planning of trajectories. Of course it would be easy to watch the simulation in operation (knowing nothing about the underlying organization) and read into the behavior all sorts of decision-making that isn't going on, all sorts of look-ahead planning of the path that isn't happening. One point I tried to make with this model is that the decision-making interpretation of behavior is a model, and it's not the only model that can explain the same kind of behavior. If the concept of decision-making is needed as part of a model of behavior, it isn't needed to explain THIS kind of behavior. The control systems in this simulation don't deal in logical conditions: they are in continuous operation, handling analog variables.

Note that we get two-dimensional paths out of a collection of control systems each of which senses only a one-dimensional variable. If you include velocity, three dimensions are actually involved. But the control systems do not have to work in multiple dimensions. The relationships between the dimensions are imposed by the environment, in the link between output and input for each control system.

Best to all,

Bill P.

==========

Date: Thu, 28 Nov 1991 11:02:51 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: scalar control

[Martin Taylor 911128 10:30]

(Bill Powers 911127.1000)

I think Bill has answered my problem as to whether scalar control systems can handle the situation in which one has to increase the perceptual error in order to reach the reference. His solution is very neat, in that it involves the interaction of control systems that use polar coordinates (direction percepts) with those that use linear coordinates (distance percepts). But although it does satisfy the conditions I proposed, it doesn't feel quite right as a solution for what I had in mind. All the same, maybe it points the way.

Bill's solution involves proximity percepts that indicate imminent failure in a particular direction. It seems to me that the mechanism would find its path much as a blind man would, feeling its way along barriers but always attempting at the same time to get nearer the goal. If the barrier surface were complex, or if it went the wrong way initially (moving away from the opening in the barrier) solution should take a long time. It has much the same feeling as the solution I prohibited that used random reorganization (the e. coli solution, as Rick put it), though of course it is much more efficient and is not random. I'm not sure whether it would solve this modified problem:

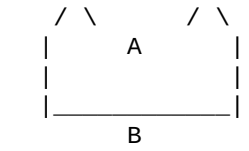

It seems to me that Bill's machine would climb the wall towards the opening, but when it came to the end of the reverse slope, it would go directly to the bottom wall, and would get caught in an endless loop.

The kind of solution I was looking for would result in the kind of behaviour we observe--a move upward, sideways, downward, and sideways to the target. There should be no initial move downward to the barrier, or even sideways. A sighted person would do this, and a blind person might make the loop once, but would do something different on arriving at the same place the second time.

Maybe (as often happens) I have failed to see the power of Powers' solution. I extended it in my mind to incorporate his machine working in imagined space, which should allow it to act as if it had planned a solution to the problem I initially posed, but then I could not see how even that amendment would provide a solution to the revised problem, since even the solution in imagined space would lead to an infinite loop (as I understand Bill's machine). It still seems to me that a machine that appeared to follow a purposeful track would need to have some two-dimensional perception in it somewhere. I'd like to be shown that I'm wrong.

I have a different, and also perhaps theoretically interesting difficulty with Bill's solution: How do the proximity detectors know whether the motion systems are not strong enough to push through (or over) the barrier?

Martin

Date: Fri, 29 Nov 1991 07:39:16 -0700
From: "William T. Powers" <powersd@TRAMP.COLORADO.EDU>
Subject: Thanks; Models<1> and Models<2>

[From Bill Powers (911128.0900)]

Yes, Chuck, it is the season for giving thanks: to you, to all the other contributors to CSGnet (especially its founder, Gary Cziko), and to all the silent listeners out there whose continued presence is an affirmation of their interest. I have the most reason of all to give thanks: without all of you, my work would have come to nothing.,

---

Bruce Nevin (911127) --

Epistemology: Models<1> and Models<2>.

Isn't the concept of models<1> part of models<2>?

What we know is the world of direct experience. If this is what Wayne means by the "psychophysical flux" then I would agree with him but would not use that term. To say either "psycho-" or "-physical" is to introduce description, characterization, and classification, which are added tg what we observe and go beyond what we observe. The term "psychophysical" introduces a model and a theory right at the beginning of the discussion,

which for me is too early. I want to explain the psychophysical flux, not assume it as a premise.

I do take the viewpoint of the HCT model in explaining the structure of subjective experience. I take the viewpoint of the physics-model in explaining the structure of the environment and brain. The purpose behind doing either one is to explain the way the world seems to me, as it is directly experienced. The ultimate criterion of truth for me is not any principle of philosophy, physics, neurology, or logic. It is simply whether these viewpoints, considered together, explain in an honest, testable, and self-consistent way what is puzzling to me about the world that I experience directly. I am not puzzled about its existence. I am puzzled about why it works as it does.

Why do I prefer honesty, testability, and self-consistency? Because I like such things. Everything seems to work better when such principles are accepted as constraints. Highly recommended. Best Buy. But not proven.

Model<1> is an assertion about the brain's built-in abilities to generate a perceptual world based on an external physical world. So it already contains an assertion about an independent physical world and a world of brain function. Model<2> extends model<1> to higher levels of brain function, by explicitly introducing modeling in terms of symbolic processes. But aren't these two models simply ways of classifying subsets in the general HCT model? Is either one more "real" than direct experience?

All of science, in my view, revolves around direct experience. We don't require models to explain each other: they are all required to explain what a human observer can experience and how human actions and spontaneous changes in the observed world affect human experience. A theory or a model must bring order into the relationship between actions and perceptions, where "perceptions" includes both unaided human observation and observations of the readings generated by instruments (there's no fundamental difference).

To me it is simply a fact that I don't experience anything but the surface of the world. I don't understand how anyone can claim that this is all there is. All you have to do is dig a hole, and you will see that the surface of the ground is held up by something else. That something else is hidden from the senses until you dig the hole. What holds IT up is hidden until you deepen the hole. I see no hope of ever seeing what holds it all up, at the center of the Earth: long before we could get there, our shovels would melt. I'm willing to entertain the possibility that there is really a nickel-iron core in the center of the Earth. I do not, however, confuse that possibility with an actual experience of the Earth's core, the only incontrovertible verification of the possibility. Another theory might claim that there is a black hole at the center. The universe is not expanding: we are shrinking into the event horizon of the black hole. How would we verify that?

Hierarchical control theory is verifiable to the extent that it predicts classes of perceptions that we can actually experience and control. Control theory is verifiable to the extent that it predicts relationships among actions and perceptions that we can actually experience. In neither case, however, can we verify intermediate processes required to make the model work but which do themselves correspond directly to aspects of direct experience. None of us, for example, can verify that these processes take place in a brain. That is conjecture. We will only know

that these processes take place in the brain when we can link each process to a perception or measurement of activities and relationships in a brain and show beyond doubt that affecting each process as measured affects direct experience exactly as predicted. On the way to doing this, our conceptions of the intermediate processes in the model will undoubtedly change, and radically. The only things that must not change are the correspondences between variables in the model and aspects of direct experience. They provide the anchor points in reality.

In PCT and HCT, certain identifiable aspects of direct experience are labelled "perceptions," and they correspond 1 to 1 with specific signals in the model (or they would if the model were complete). This does not change direct experience. It does change what we think about direct experience. We are led to think of all descriminable aspects of the experienced world as "perceptions," not just as givens. The perceptual signals in the brain model are linked theoretically through physical properties of neurons to other signals, and eventually to variables in the physical model of the world. The physical model deals primarily with variables and relationships that do not correspond to perceptual signals: a world beyond the senses. As predicted, we do NOT experience electrons, light waves or quanta, force fields or energy. The physical world becomes directly experiencable only at contact points established by meter readings of various sorts. What we experience is a meter reading, not the physical process that gives rise to it. Processes intermediate to those contact points and the physical variables on the other side of the meter remain conjectural and unverified. Therefore the two models together imply that what we perceive is not neccessarily in direct correspondence to the entities and relationships in the world proposed in the physics model. If we choose to use both models, the viewpoint we must take is that the world of experience is derived from or dependent on another world that is not experienced, just as the surface of the ground that we can see and touch is held up by deeper layers of unknown composition that remain invisible and intangible.

This is the only viewpoint I can see that is consistent with physical models, neurological or biological models, functional models of the brain, and direct experience. What we experience is not a model. Everything we say about experience is.

Best

Bill P.

==========

Date: Fri, 29 Nov 1991 21:55:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Avoidance/seeking behavior

[From Bill Powers (911129.1100)]

Martin Taylor (911127) --

>Bill's solution involves proximity percepts that indicate imminent
>failure in a particular direction.

That's one way to interpret the meaning of proximity perceptions but the model doesn't know anything about "imminent failure." It just keeps trying to control proximity and direction. If it gets stuck it gets stuck. This doesn't bother it. Do you want it to be bothered when it gets stuck? Tell me how it perceives stuckness. That could be controlled for,

too.

>It seems to me that the mechanism would find its path much as a blind
>man would, feeling its way along barriers but always attempting at the
>same time to get nearer the goal.

The proximity function is adjustable. The half-max-proximity distance (for the avoidance systems) could represent any distance on the playing field. If it were set to about a third of the width of the trap, and the reference proximity a little higher than that, the person would never come near the barrier. This is how a sighted person would move if physical contact were considered impolite or dangerous. If the proximity range were set large enough, the person would never get into the trap unless he/she/it were born there. We would then say that it avoids getting into traps too small to get out of again, which sounds pretty smart, but isn't.

> If the barrier surface were complex, or if it went the wrong way
>initially (moving away from the opening in the barrier) solution should
>take a long time.

If the barrier surface were complex enough, YOU couldn't get out of it. We're talking about the degree of intelligence that is to go into the model. I will match the model described against any one-celled organism you want to back. In fact, the initial conditions make very little difference in the behavior of the model as described. If it's pointed the wrong way it will just turn around and go the right way. The only random element in the model is whatever is needed to resolve head-on collisions, and it isn't much.

> I'm not sure whether it would solve this modified problem:
>

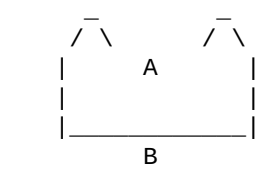

Yes, it would, unless you made the opening very small in comparison with the dimensions of the trap. Remember, the proximity reference level can be set so that the effective surface looks like this:

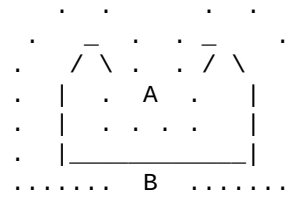

... so the person doesn't even gewt near traps small enough to be smoothed out by the perception of proximity to the surface.

The proximity function is actually a cardoid: the raw proximity to each point-obstacle is weighted as the cosine of half the angle from the direction of travel to the direction of each contributing obstacle. So it gives the most weight to obstacles in the direction of travel. With the

above arrangement, the person might have some trouble getting to the goal, because the goal position violates the proximity criterion as shown. A real person would have the same problem if the wall were red-hot and proximity were detected by temperature sensors.

>It seems to me that Bill's machine would climb the wall towards the >opening, but when it came to the end of the reverse slope, it would go >directly to the bottom wall, and would get caught in an endless loop.

Depends on how wide the avoidance range is, where the reference-proximity is set, and how sensitive the turn-generator is. The proximity of the bottom part of the trap, however, will always outweigh the proximities near the opening, where there is a span of zero proximity. I don't claim that endless loops are impossible, however. Just unlikely.

>The kind of solution I was looking for would result in the kind of >behaviour we observe--a move upward, sideways, downward, and sideways to >the target.

Sort of like a soldier doing close-order drill: about face; forward march; column right, column right, column right? I suggest that a real non-military person will move in continuous curves, as this model will do, skirting the obstacle but not moving in straight lines or turning at right angles.

>There should be no initial move downward to the barrier, or even >sideways.

If you put the sighted person on a bicycle facing downward, I wager that the initial move would be a small semicircle starting toward the barrier, turning to one side, and then straightening out in a path toward the opening. That is what this model will do. It always moves at some minimum velocity, and can't turn in place. It can, however, turn on a very small radius.

>A sighted person would do this, and a blind person might make the loop >once, but would do something different on arriving at the same place >the second time.

You're adding to the conditions: the model must now be able to take advantage of previous experience. In the simplest cases with the parameters adjusted reasonable well, this wouldn't improve performance any. How could you improve on turning toward the opening, going through it without colliding, then taking a path that is as close to the outside of the trap as collision-avoidance criteria permit?

By adding another level and sequence memory to the model, it could record a few of the positions along the way, then repeat the sequence of positions as intermediate goal positions, letting collision-avoidance systems take care of details (or changes in initial conditions). This would be useful in a maze. It wouldn't be useful in moving to a goal through a gathering of other people who are moving unpredictably around. My model can do this without any added intelligence.

>It still seems to me that a machine that appeared to follow a purposeful >track would need to have some two-dimensional perception in it >somewhere. I'd like to be shown that I'm wrong.

If the controlled variable is a particular path through 2-space, then two event-control systems, each controlling one of the dimensions, would be

required; each would produce a specific pattern of positions in one spatial dimension (whether x-y or polar). A relationship-control system would be needed to maintain a specific relationship between the two patterns of change. But that adds an unnecessary requirement if all you want to do is to escape the trap, avoid obstacles, and get to the goal. It seems to me that organisms very seldom insist on repeating a particular path. In fact I doubt that they ever repeat the same trajectories twice even when they're trying to, as in doing the school figures in competition skating. In most real cases, repeating the path would get you into trouble. The world changes.

A lot of people seem to suppose that because travel from one point to another always follows some path, the path must be preselected. I don't think that is actually a requirement -- what would the added advantage be if a reasonable path results without being specified in advance?

Would you like me to send you a copy of the program?

>I have a different, and also perhaps theoretically interesting
>difficulty with Bill's solution: How do the proximity detectors know
>whether the motion systems are not strong enough to push through (or
>over) the barrier?

[More conditions, eh? You tell me: on what basis could they "know" this? Define the perceptions and I'll tell you how to control them.]

Tell me at what distance the barrier will begin to give way. I can program the environment to make pieces of the barrier yield when that amount of proximity is reached. Then I can set the reference-proximity so turning begins a little closer than that. If the barrier doesn't give way in time to prevent a proximity error, avoidance will take place. Of course if you require that the reference proximity be set very close to the barrier to accomplish this, and forbid turning before a close approach to contact is made, then we're back to the case of staying close to the wall while looking for the opening. You can't have it both ways, unless you want to add a program level that first tries a close approach to the barrier for a while, using a short-range proximity detector, and then, after n failures, switches to the long-range proximity detector and lets it find the way out. If you want logically complicated behavior you have add logic to the model, too. That's why we have evolved a program level. Sometimes logic is required.

I could do this, but I'm not going to.

The point is not to complicate the model if you don't have to. If you have to, then complicate it by the least possible amount. Or start over: when my models get too complex I take that to mean I've started out wrong. Into the bit bucket with 'em. I have been known to discard a year's work. Not cheerfully.

Best

Bill P.