9112A CSGnet

Date: Tue Dec 17, 1991 5:05 pm EST
From: Gary Cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: CZIKO@vmd.cso.uiuc.edu

TO: * Hortideas Publishing / MCI ID: 497-2767
Subject: csg-l log9112a

==
Date: Sun, 1 Dec 1991 13:32:25 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Avoidance/seeking behavior

[Martin Taylor 911201 12:45]
(Bill Powers 911129.1100)

Forenote--I will be away from tomorrow lunch until late in the week, so don't expect an early response to a response to this.

Bill solved the problem I posed in which a barrier was interposed between a present position and a reference position for some object or tool. I felt that the solution did not have the generality that I was hoping, and modified the conditions a bit, in the hope that the desired generalization would emerge, and Bill solved that problem, too. But again I cannot find the generalization I seek. I'm not trying to be obtuse or to make things difficult, but here goes, again.

This time, please do not think of A and B being just places where a person might stand, and the wall as an object that can be detected by proximity. Think of a person holding a rod that touches a wall at A, the rod being spring-loaded so that it is easy to move while keeping it in contact with the wall. The person holding the rod wants to move the tip from A to B, but there is a barrier. The point is that the person can apprehend the whole shape of the situation, A, B, and the barrier, so that Bill's comment:

>If the barrier surface were complex enough, YOU couldn't get out of it.

would not apply. The barrier surface shape would be simply irrelevant if the person could see enough to avoid the whole region that contained the barrier. And Bill's other comment about the person holding a bicyclehaving to make a small circle and not being able to go straight out of the trap would also not apply.

What I am trying to get at is the question of whether unidimensional control systems can do all that Bill claims, by attempting to pose problems that seem to provide minimal conditions where 2-D systems might behave differently and more like people do. Whenever Bill solves one of these problems with a set of unidimensional systems, that strengthens my belief in his claims. But I am not yet convinced on this particular problem. I would like to see the solution lead naturally to a path that simply bypasses the barrier, without the need for tuning the comparative gain of control systems that one would think should act independently of one another. The present solution still makes me think of a blind man fumbling his way round the barrier until he finds the opening. Perhaps the sighted man also does that, but only in imagination.

The dynamics of the control systems include the constraints imposed by the external environment. The constraints are not unidimensional, and they impose relationships among the behaviours of the feedback systems whose controllers are unidimensional. In a lot of cases, Bill has shown how these relationships in the external part of the feedback loops simply do not matter. The control systems achieve their goals regardless. Looking at these systems from my (usually) geometric/topological point of view, I see the success of the scalar control systems as being permitted by the fact that the problem space is equivalent topologically to a multidimensional paraboloid. I am looking for a problem whose space intrinsically bifurcates, or in which there exists a quick solution that involves a segment in which the local motion is away from the final target, or in which

the control system can get caught in a local minimum of a complex solution space. Perhaps I am barking up the wrong tree, and what I should be looking for is not a problem involving an intrinsically multidimensional solution, but one in which non-local information is required for a solution. Non-local information means that the local gradient of approach to a solution could be overridden by information about the shape of the control function in other areas (Imagine being in a burning house with flames between you and the only door. If you go through the flames, you get burned, but after that you escape to the cool outdoors).
So, Bill, don't think I am being capricious or trying to catch you out by changing the conditions on a problem. If you solve a problem that I think poses a difficulty I can't get around within my understanding of PCT, that both increases my understanding and enhances the apparent power of the theory. You may not need that enhanced power, but science is a social endeavour, and there are others of us who do.

Martin Taylor ===
Date: Sun, 1 Dec 1991 15:44:38 EST
From: goldstein@SATURN.GLASSBORO.EDU
Subject: any comments?

To: CSGnet people
From: David Goldstein
Subject: method of levels, suicide
Date: 12/01/91

Recently Iposted on the topics: method of levels and suicide. So far, except for Bill and Rick, no one has made any comments. That made me wonder why. and so I ask: Why?

Thanks for your attention to this.
David Goldstein
===
Date: Sun, 1 Dec 1991 17:21:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: boss reality; scaler control loops

[From Wayne Hershberger--This may be a bit anachronistic because
I've not been able to read my mail since 911127; sorry fellas, I'm pedalling as fast as I can.]

Bruce Nevin (911127 0743)

>As we all know, "boss reality" doesn't really sit still for its >picture to be taken.

What!? You boggle my mind.

Your sentence implies what it denies: that is, although we can not picture it, "we all know 'boss reality....'"

Perhaps you meant to say that although we can picture it, we can not know boss reality. But, of course, such a transcendental reality as that smacks more of heaven than earth. The relationship between a hierarchical control mechanism and its environment is a much more mundane affair than picturing a transcendent reality. You imply in the following two paragraphs two different avenues of access.

>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.

I am not sympathetic with the first point. To say that the environment is hidden by all the proximal stimuli is to paraphrase the fellow who claimed not to be able to see the forest for all the trees (e.g., Gee officer, I couldn't see the fireplug; my eyeballs got in the way). Also, don't forget that the "intensity sensors" are spatially arrayed and sensitive to various forms of energy--over time. Further, transducers such as radar scopes vastly expand the range of our biological transducers. More trees to obscure our view?

However, I am favorably impressed with your second point, which is very similar to one I addressed last year--before you logged on to CSGnet. At that time, I observed that sensed efference affords a significant window to the world; that is, when an environmental variable is being controlled, sensed efference reflects the environmental disturbance (e.g., the weight of an object is proportional to the effort required to heft it). This principle provides a basis for the ideas of that other Taylor Martin frequently refers to.

Since that time, I have come across a delightfully lucid example from physics. Some physicists (Gerd K. Binnig & Heinrich Rohrer) won a share of a 1986 Nobel Prize by capitalizing on this principle in their design of the scanning tunnelling microscope, STM.

"The STM operates by passing an ultrafine tungsten needle over the surface of a sample to be studied. A low voltage is applied to the needle, creating a tiny electric potential between the tip of the needle and the atoms on the surface. Although the needle and the sample never touch in the classic sense, quantum fluctuations enable electrons to 'tunnel' through the intervening distance, hence the microscope's name.

"The current passing between surface and tip depends on the distance between them. A feedback mechanism continuously repositions the needle as it scans over the surface to maintain a constant voltage: the undulations of the needle are studied to reconstruct the sample's contours. (Scientific American, June 1990, p. 26)"

Martin Taylor (911126)

The Umweg (detour) problem you pose is one that some living control systems (e.g., chickens) CAN NOT handle. That implies to me that they DO comprise scaler control loops. Right? Perhaps what you are asking, is why can primates handle it. But can they? We touched upon this matter earlier when Judd described how one traps a monkey by placing a bit of fruit in a hollow log with a small knothole; when the monkey reaches through the knothole and grasps the piece of fruit he is trapped by his enlarged fist. "Why doesn't the monkey let go?" Joel asked. I don't know that he doesn't let go, but if he doesn't, that would be another example of scaler control loops. Right? Perhaps, what you are asking is why humans can handle it. But do they? In Chicago several years ago, a DC10 crashed when the flaps on the left wing retracted (because the hydraulic lines were severed) producing asymmetric lift, for which the ailerons could not compensate. The plane rolled slowly counterclockwise losing altitude and crashing on its left wingtip. Had the pilot used the ailerons to augment the flaps and quickly roll the aircraft counterclockwise 270 degrees, through an inverted to an upright attitude he might have bought himself enough time to have averted the disaster altogether, as Chuck Yeager once did when he was faced with a similar circumstance (see his autobiography). Apparently, the DC10 pilot did not have the "right stuff."

Right? Is that what you are asking, what is the "right stuff?"
Interesting question from a control theory perspective.

Gary Cziko (911127)

>critical reviews of PCT articles (and chapters or whatever)
>submitted by PCT types for publication.

I, for one, will see what I can dig up.

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: Mon, 2 Dec 1991 01:37:57 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From:
mmt@DRETOR.DCIEM.DND.CA
Subject: Re: boss reality; scaler control loops

[Martin Taylor 911202 01:30]
(Wayne Hershberger 911201 18:40)
>
>
>The Umweg (detour) problem you pose is one that some living control systems (e.g.,
>chickens) CAN NOT handle. That implies to me that they DO comprise scaler control
>loops. Right?

After Bill's responses, I'm not at all sure whether it is right or not. He has been satisfactorily solving my little problems using only scalar control systems.

But on chickens, I've been meaning to enter this little tidbit into the stew, from Scientific American several years ago. Chickens naturally peck at grains of wheat (or whatever it is they eat), and peck accurately. But if you put prism spectacles on them so that there is a deviation of a few degrees, they peck in the wrong place, perhaps a centimeter away from the grain. This
error never seems to decrease. So if they do have a control system for the peck, it seems not to be subject to a higher-level one that adjusts the reference for the peck system. I just thought I'd throw that in, because some time ago someone made a comment that the chicks would correct, as part of another argument. If anyone can look up Sci Am indexes and find the article (10-15 years ago is my vague memory), it might be helpful to get a proper reference.
Martin Taylor
==
Date: Sun, 1 Dec 1991 23:48:19 -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: 2-D contro; Straw men

[From Bill Powers (911201.1930)]

Martin Taylor (911201) --

I appreciate what you're doing, Martin. It's useful to consider what
added capacities would be needed as the environmental problems become more complex. This is roughly the kind of thinking I went through in formulating the levels. It may well be the story of evolution.

Maybe it would help if you considered that each kind of problem we can solve with the simple systems I propose is a problem that a more complex system doesn't have to solve. If you have a collision-avoidance system, higher-level systems don't have to consider obstacles of ordinary sorts: just specify the goal position and you'll get there, even if the obstacles shift around and your path shifts accordingly. This clearly simplifies the higher-level problems, so higher systems only have to decide where to go, not how to get there. And, of course, how to deal with extraordinary obstacles.

You comment

>What I am trying to get at is the question of whether unidimensional control systems can
>do all that Bill claims, by attempting to pose problems that seem to provide minimal
>conditions where 2-D systems might behave differently and more like people do.

I wonder if the problem isn't with the conception of a "2-D" system rather than with my models. Can you actually define a two-dimensional variable in a way that's not just a notational convenience (like a vector denoted with a single symbol, but actually consisting of two independent variables)? I ran into something like this in a discussion of some NASA "teleoperator systems" publications. The standard language was matrix algebra and matrices of differential equations, so of course the expressions for complex multidimensional control systems looked very compact. This made for brevity and avoided pages and pages of mathematics. But when it comes time to build such a system in hardware, the matrix notation is useless: you have to expand the matrices into their elementary operations and implement each one. This is true even if you simulate such systems in software: eventually, at the machinelanguage level, you have to multiply each element in a row by each element in a column and add the terms, do the same for the next row, and so on in all the tedious detail that motivated the invention of matrix notation in the first place.

When you say you want a system that will behave "more like people do," are you convinced that the systems I described would behave in some unrealistic way? Most people who see the program running think that the behavior is reasonably realistic, although not as intelligent as it might be. How is it that people behave that you imagine this model can't imitate? If you could pin that down, the result would be a possible improvement in the model. It might indicate what a useful next higher level of control might be.

I don't think your example of the collapsible pointer is specified well enough to imply a model. Where are you standing? Can you withdraw the pointer? How do you know if the pointer is too short to reach part of the obstacle? How do you detect when the pointer touches something? If you feel along a wall and encounter an inside corner, what degrees of freedom are allowed to get the pointer out of the corner? Can you see the tip of the pointer?

The question in designing control systems is always "what does this system need to know in order to act in the way I imagine?" To ask what it needs to know is to ask what perceptions must be controlled in order to lead to the observed behavior. The behavior itself -- the actual movements and trajectories -- should arise naturally out of the control of the proper variables in any normal environment that happens to exist; the system itself should arrive at movements and paths that look realistic, not through planning them but through controlling the right consequences of those actions. No realistic model can be based on planning of output, because such models can work only in a world that never changes. We should not be thinking in terms of how to duplicate behavior -- actions. We should be thinking of how a system can produce appropriate actions under unpredictable changes of conditions, resulting repeatedly in a preselected consequence. I think the way to get to control of complex consequences is to solve first the problem of controlling for simple consequences, which can then become the elements of more complex ones.

>I am looking for a problem whose space intrinsically bifurcates, or in >which there exists a quick solution that involves a segment in which the >local motion is away from the final target, or in which the control >system can get caught in a local minimum of a complex solution space.

All these things are possible, although I can't imagine an environment that "intrinsically bifurcates" and that a natural control system would be fit to occupy. I'm really interested only in the environments we actually encounter, and the most common ones at that.

It's an interesting question as to what the minimum control system would be to handle instances of each case. Each time you add new conditions that increase the difficulty of the problem, you introduce new requirements for perception, and often new capabilities for action that were not initially considered. This is a legitimate part of building from simple models to complex ones. But there is a danger of inventing environments that don't exist or that are too complex for any system with only human capabilities to master. If you want to baffle a human being's control systems, that's not too difficult a project. Just ask the human being to love his neighbor as himself.

You say

>The barrier surface shape would be simply irrelevant if the person could >see enough to avoid the whole region that contained the barrier.

But this is exactly my point. What variables would the person have to be able to "see?" What would be "enough?" What knowledge of the environment -- on the part of the system, not of an all-seeing observer -- would be enough to enable avoiding the whole region that contains the barrier? (Actually my model will do this if you make it very leery of collisions). And of course you also have to ask the converse, having raised the point in court: if you design the system to avoid the whole region containing the barrier, how could you then cause it to seek close proximity to the barrier to trace out its shape for other higher-level purposes? Would you have to redesign the system, or would it already contain adjustable parameters that could allow either kind of behavior (my model does). If solving one problem makes others, even simpler ones, impossible to solve, you haven't really got anywhere --or you may have to go somewhere you hadn't originally intended. The hierarchy of control was necessary to postulate because, having specified control systems for every muscle in the body, and having been convinced that they actually exist, I rendered it impossible for higher systems to use those same muscles directly for higher purposes. The only solution was to have higher systems use the first-order systems as parts of their output functions, and to use them by recommending inputs rather than commanding outputs. After that, one
thing led to another.

Simply saying "The point is that the person can apprehend the
whole shape of the situation ..." doesn't tell us what to model. What is "the whole shape of the situation?" Is this an actual perception you're talking about, or a vague impression that leaves most of the essential details to the imagination? The problem of producing a precise action on the basis of an indistinct perception and thus an uncertain error is, I think, inherently impossible to solve -- through logic or through evolution. It's a self-contradiction. The precision of control is exactly the precision of perception.

>Perhaps I am barking up the wrong tree, and what I should be looking for >is not a problem involving an intrinsically multidimensional solution, >but one in which non-local information is required for a solution.

This is more like it. If, to pick a simple alternative ennvironment, my control system could not see point B, but only a sign saying "Proceed backward out of this trap and go around it to point B," it would be helpless. It can't read, much less understand verbal instructions. A higher system could translate that instruction and convert it to a series of intermediate goal positions, the last one being point B. Bruce, how would I build that system?

This may bring in something like the 2-D considerations you want. Suppose that there are perceptual maps existing in the brain, that preserve angles and adjacency relations (if not the strict Euclidian geometry we attribute to the world). Specifying a reference position in such a map could amount to activating a point in this map. Such a point of activation would, simply by existing, specify two dimensions (or even three, or more) at once. If this were a visual map, then the problem of control might be that of bringing the visual field into coincidence with this map, centered at the activated point. True, the actual control processes would have to involve one system for each dimension in which independent variation can occur. But the higher system might identify the reference-position in some way that did not involve spatial considerations -- for example, it might be concerned with a characteristic other than its location of whatever is in that position.

But that characteristic, from the higher-level standpoint: would it turn out to be a collection of unidimensional variables also? I'm not against 2-D or n-D perceptions. I just don't see how to make error signals from them that will lead to actions in each required degree of freedom.

David Goldstein (911201) --

I'm not going to broadcast advice about suicide, pro or con. Your approach is on the right track. I leave this field to the clinicians who have to deal with this problem with real clients. If you like I will tell you about my own flirtations with this solution when I was young. But control theory can't provide one general approach that fits all.

Wayne Hershberger (911130 or so -- deleted the date) --

Your comments to Bruce:

>Perhaps you meant to say that although we can picture it, we can not >know boss reality. But, of course, such a transcendental reality as >that smacks more of heaven than earth.

Is that "of course" an argument against the proposition, or a bit of innuendo associating Bruce with a proposition that you DO know how to refute? You must have a better reason than that for rejecting the possibility of a boss reality. Are you arguing against the uses of imagination?

>I'm pedalling as fast as I can.
I'll let Bruce defend himself, and wait for your reply to my last comments, when your bicycle gets to them.
===
Date: Mon, 2 Dec 1991 08:32:19 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: "Active learning and control"

----BEGINNING OF FORWARDED MESSAGES----
Received: from LABS-N.BBN.COM by CCB.BBN.COM ; 29 Nov 91 13:54:23 EST
Received: from KARIBA.BBN.COM by LABS-N.BBN.COM id aa20486; 29 Nov 91 13:51 EST Received: by KARIBA.BBN.COM id ad03129; 29 Nov 91 13:30 EST
To: machine-learning@BBN.COM, neural-people@BBN.COM
Subject: Active Learning and Control Workshop
From: aboulang@BBN.COM
Sender: aboulang@BBN.COM
Reply-to: aboulanger@BBN.COM
Date: Fri, 29 Nov 91 13:35:09 EST
Source-Info: From (or Sender) name not authenticated.

From: pablo@cs.washington.edu (David Cohn)
Newsgroups: comp.ai.neural-nets
Subject: Active Learning and Control Workshop
Date: 25 Nov 91 18:58:19 GMT
Organization: Computer Science & Engineering, U. of Washington, Seattle

Below is the final schedule for the Active Learning and Control workshop scheduled for the post-conference workshop in Vail, Colorado on Dec. 6-7th, after the Denver Conference on Neural Information Processing Systems.

```
----------------------------------------------------------------------------
|##########################################################################|
|**************************************************************************| |*        *|
|*        NIPS Workshop on Active Learning and Control      *|
|*            *|
|*        Co-Chairs: David Cohn, University of Washington CS&E     *|
|*                   Donald Sofge, MIT Artificial Intelligence Laboratory *|
|*            *|
|*     An "active" learning system is one that is not merely a passive    *|
|*     observer of its environment, but also plays an active role in      *|
|*     determining its inputs.  This definition includes classification   *|
|*     networks that query for values in "interesting" parts of the domain, *|
|*     learning systems that actively "explore" their environments, and   *|
|*     adaptive controllers that learn how to produce control outputs to  *|
|*     achieve a goal.     *|
|*            *|
|*     Common facets of these problems include building world models in   *|
|*     complex domains, exploring a domain safely and efficiently, and    *|
|*     planning future actions based on one's world model.  *|
|*            *|
|*     Our main focus in this workshop will be to address key unsolved    *|
|*     problems which, once solved, may promote the application of active
*|
|*     learning to real-world control systems and other problem domains.  *|
|*     Our hopes are that research into unsolved problems in one field may  *|
|*     draw insight from research in other fields.    *|
|*            *|
|**************************************************************************|
|##########################################################################| ------------
----------------------------------------------------------------
```

Day 1, Morning   Active Learning

Intro: David Cohn, Univ. of Washington CS&E
                 Definition of Active Learning, theory from the computer
                    science perspective, simple query learning.
       Tom Dietterich, Oregon State Univ.
                   Data selection using response-surface methodology.
       Andrew Moore, MIT AI Lab
                 Intelligent experimentation with memory-based learning.

Discussion: The Ruff-Dietterich hypothesis: How important is it to simply do experiments
                vs. how important is it to do *good* experiments? Are the results from
                simple theory applicable to real problems? Perspectives on building
                accurate world models
                    Where do we go from here (i.e. how can these ideas be applied)?

Day 1, Afternoon   Learning Control

Intro: Michael Jordan, MIT Brain & Cognitive Sciences

Action Search with Forward Modelling

Kevin Markey, CU Boulder
Reinforcement Learning Benchmarks for Phonology Acquisition

Michael Jordan, MIT Brain & Cognitive Sciences
Multiple Network Modelling

Discussion: Reinforcement vs. error-driven learning
On-line vs. off-line training (e.g. real-time constraints, local minima?, noise effects, convergence time, when

to turn learning on/off)

Global vs. local approximation methods

Ties to Optimal Control Theory (e.g. how much of this is mathematically sound? what methods or lessons may be applied from Optimal Control?)

Various types of reinforcement Clues from neurobiology
Where do we go from here?

--Day 2, Morning

Active Exploration

Intro: Sebastian Thrun, Carnegie-Mellon University
Exploration in Reinforcement Learning and Adaptive Control Juergen Schmidhuber, Univ. of Colorado, Boulder
Active exploration in stationary and non-stationary domains Michael Littman, Bellcore
Perpetual Exploration

Discussion: Cost of exploration (exploring while performing system optimization)
Danger avoidance (or failure avoidance) in actively exploring control
(what if the cost of failure is too high?)

The use of uncertainty models in exploration
Tree search, LRTA*, other approaches
On-line exploration, use of on-line forecasting, and hybrid approaches
How may exploration methods be applied to existing learning control systems?

Day 2, Afternoon   Planning

Intro: Rich Sutton, GTE
Dynamic Programming, Planning vs. Reacting

Satinder Singh, UMass Amherst
Solving Multiple Sequential Tasks Using a Hierarchy of Variable Temporal Resolution Models
Chris Atkeson, MIT AI Lab

Strategy Formation

Discussion: The wide world of planning: review of different types of planning currently used (e.g. sequential look-ahead, back-propagation through models or time, dynamic programming).

Uncertainty and noise: how can they be accommodated in planning? How may these techniques be effectively combined in connectionist based models?

Conclusions of Workshop

Important Notice to Participants:

Please confirm that the Titles of your talks as given above are accurate and that other information (name, affiliation, etc.) is correct. Also, please verify that you will be available during the time slot for which we have you scheduled. Notify David Cohn (pablo@cs.washington.edu) or Don Sofge (sofge@ai.mit.edu)

of any exceptions.
(post bulletin board notice to Cohn/Sofge after conference starts)

An overhead projector will be available for making presentations. Please notify either Cohn or Sofge for any additional equipment you would like (e.g. vcr, monitor, slide projector, etc.) and we will try to accomodate your needs.

The workshop format will be 20 minute presentations, each followed by a 10 minute question/discussion period, thus allowing 30 minutes for each speaker. However, to encourage discussion, these times will not be strictly enforced and thus will serve as a guide to keep things moving.

A block of time has been set aside at the end of each session to allow for further discussion of the presentation topics, or in order to tackle the "Discussion Topics" listed above. We do not anticipate having time to address all of the topics listed; therefore the topics discussed will be selected based upon audience interest. If you would like to suggest addional topics for discussion, please send them to Sofge or Cohn, and we will put them on the list. As indicated, these topics will be addressed as time permits.

----END OF FORWARDED MESSAGES----
==
Date: Mon, 2 Dec 1991 10:14:44 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: Saturday's categories

[From: Bruce Nevin 911130 -- Gloucester]

(Bill Powers (911126.0800)) --

I spent some time yesterday formulating responses. Lots of words into the bit bucket. "Poit" back. I will put more effort today into isolating essential issues.

* Perceptions<L> (words) don't correlate with perceptions<W> ("the world") in a simple or direct way. This is a more fundamental issue than the observation that words evoke the

allusive richness of associative memory.

* Categorization is promiscuous in a way that other perceptual processes are not: we can and do categorize perceptions of any level.

    * There is a close identification between categorization and verbalization.

* I proposed an alternative view, in which the categorizing required for sequence- and program-level ECSs is done by the input devices of those ECSs.

    * The basis for this proposal was an idea that input devices at every level perform a categorization of lower-level perceptions.

* This helps to explain some of the apparent differences between categorizations of perceptions at lower and at higher levels: the conventionalization that is in language carries over into categories that are learned with the assistance of "recipes for recognition" rehearsed in language.

* Sequences and programs may take as perceptual input categorizations of perceptions from any level whatsoever. Either sequences and programs can accept categories input from everywhere in the hierarchy, or ECSs of the category level can accept input from every level of the hierarchy.

    * Analogy and metaphor is central to all of this, but we have no good account yet of how analogy and metaphor work in the control hierarchy.

    * Analogy (metaphor), categorization, and conventions (norms) seem to be mutually interdependent.

    * Conventions (norms) are just reference signals. They differ from other reference signals only in the kind of environmental feedback on which we base them.

I'll take these ideas up in turn.

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

Perceptions<L> (words) don't correlate with perceptions<W> ("the world") in a simple or direct way. By this observation, I don't just mean that words evoke the allusive richness of associative memory, as you have rightly indicated. The looseness of correspondence between language and non-language is much more fundamental.

We can find examples where words seem to correlate with perceptions or at least with categorizations of perceptions, especially with the primitive arguments (requiring no argument words for one to enter it into the construction of a sentence) such as "dog" and first-order operators on primitive arguments only, such as "jump". Harris has shown that the correspondence is best for science sublanguages, but becomes loose and inconsistent for the language in general.

At first glance, the word types established by their operator/argument dependence-on-dependence seem to correlate with levels in the perceptual hierarchy. You suggested this in your review of Harris's _Language and Information_. But even this degree of correspondence is not consistent, and breaks down.

For example, you might think of "time" as a relatively high-level perception, and "father" or "dad" as a relatively immediate, concrete perception. But "time" is a primitive argument, and "father" is an operator (in "A is the father of B" it requires two primitive arguments A and B). By reduction operations in the grammar, we frequently use "father" in the stead of its A argument: "John's father" meaning "he who is John's father," "Father" or "dad" reduced from "(he who is) my father" or "(you who are) my father," and so on. But "father" refers to a relationship of roles in a social system involving highly conventionalized prerogatives and obligations that differ enormously from one culture to another. It is not clear to me that the relationship level of the perceptual hierarchy applies here. We could have an interesting discussion of kinship terminology in various cultures.

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

Categorization is promiscuous in a way that other perceptual processes are not: we can and do categorize perceptions of any level. Even on the lowest level of intensities and efforts, where we are limited to primitive scalar judgements, we can say "that was a 6 on a scale of 1 to 10."

Your counterargument has been "when you are operating from level X you see everything in terms of level-X perceptions." Perhaps it is true that I can "configurationalize" intensities or sensations or programs in the same way that I can categorize them. Perhaps this is what is going on behind the kinesthetic/visual imagery of thought that I and others have reported.

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

This fits with the fact that I can talk about any sorts of perceptions in language, in the same breath as categorizing them (so to speak). There is a close identification between categorization and verbalization.

You have proposed that we associate a perception<L> with a perception<W> (represented in this discussion by e.g. "line" and _____, respectively) by virtue of either of these perceptual signals meeting the input requirement for a higher-level ECS. You said that this level/type of ECS was the category level. It is difficult to see how one could posit a closer identification between categorization and verbalization.

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

In the current standard theory, category ECSs just below sequence and program level provide input to sequence- and program-level ECSs. I proposed an alternative, in which the categorizing required for sequence- and program-level ECSs is done by the input devices of those ECSs. The effect of this is to demote the category level to the status of input devices at sequence level and program level. Although I think the original proposal has inadequacies that I am still working out as I write this, some questions that it raised remain.

The immediate question is whether categorization is operationally separate, on a distinct level in the hierarchy of perceptual control, or merely analytically separate because of the convenient synchronization of many categories with words that we use as category-names. Are category-level perceptions used for perceptual control, other than providing input to sequence/program perceptions and correlating with language? What disturbance may one apply, such that people respond by controlling for category perception, while not controlling for words, sequences, or programs?

Your observation still applies, by the way, about experiencing everything in terms of the level one has adopted for one's perspective. If one views things from the program level, everything looks like input to the program level, i.e. categories. It doesn't make any difference whether the categories are from input devices on the program level or from ECSs of a category level.

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

The basis for this proposal was an idea that input devices at every level perform a categorization of lower-level perceptions. This categorization is not identical with the weighted summation of perceptual signals, but rather an indirect product of it.
Let me try to explain what I mean, starting from two definitional facts about categories: different instances of a category are perceived as repetitions of that category, and there is a distinction of logical type between instance and category, commonly called a type-token distinction. A constellation of inputs abc is an instance (token) of a perception B if their weighted sum in the input device for the ECS controlling for B matches the reference signal for that ECS. A different constellation of inputs a'b'c' is a different instance (token) if their weighted sum also matches (i.e. is the same). Perception B is then a category (type) over instances (tokens) abc and a'b'c'. In general, different constellations of inputs are repetitions of a perception if the weighted summation of each matches the reference signal of the ECS for that perception, and any given instance (token) of the perception is a weighted summation by which two constellations of inputs from level n-1 are perceived as repetitions of the same perception on level n. We may speak of the level n category (type) as "referring to" the set or class of all such constellations of inputs (tokens) from level n-1 whose weighted summation matches the reference signal of the level n ECS, the extension of the category; or we may speak of the category (type) as "referring to" the single such level n-1 constellation (token) of inputs whose weighted summation currently matches the reference signal of the level n ECS, the referent properly speaking.

On this view, the most elementary categories are those by which we judge that two level-1 perceptions<W> are of equal intensity. Let us not be overly impressed by the lack of unitary vocabulary for intensity perceptions. That is merely a matter of creating shared vocabulary. We can freely use more elaborate constructions like "that was a 6 on a scale of 1 to 10," as I suggested above. And we can probably find vocabulary for perceptions<W> of this order in, for example, the sublanguage of wine-tasting.

Some turn-of-the-century phoneticians could judge the intensity of formants, apparently reliably since their accounts jibe well with our instrumental findings. This is a skill that improves with practice. It may depend as well upon innate differences, as is shown by the frequently made anecdotal observation that you have to have a good musical ear to be a good phonetician. These phoneticians used terms like "sharp." This terminology was developed further by Trubetzkoy, then Jakobsen, then Halle, as acoustic phonetic features such as acute, grave, compact, diffuse, and strident.

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

This helps to explain some of the apparent differences between categorizations of perceptions at lower and at higher levels.

* At higher levels, category-perception seems to involve programmatic "recipes" for recognition, less so at lower levels.
* At higher levels, categories seem to be conventional, less so at lower levels.

These are related. The conventionalization that is in language carries over into categories that are learned with the assistance of "recipes for recognition" rehearsed in language. Let me elaborate on this a bit.

About higher-level categories, you argue:

> 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."

I must question your major premise. People working as counters at the U.S. Mint count (or at least formerly counted) 100 bills at a time by riffling through stacks of new bills rrrrft! rrrrft! rrrrft! rrrrft! rrrrft! like a dealer in Vegas with a deck of cards. Every rrrrft! is 100 count, reliably. Such skilled perceptual control is surely a matter of practice. With practice, one can recognize quadrilaterals at a glance, with no computation of sides and angles, or a duodecahedron for that matter, just as we recognize a screwdriver at a glance, and just as a chess master recognizes positions on the board.

When we are learning a new category, we often have a recipe of description and instruction to guide us. Is there a figure? Does it seem plane? Are there sides, are they all straight, and is their number four? Are there angles, and is their number three? Then it's a triangle. But with practice we no longer need the instructions, and furthermore they get in the way. I believe it is possible to make all of the configurations we have discussed equally direct of apprehension, and that none of them depends upon the program level after sufficient practice.

With practice we can recognize types of programs or program constructions at a glance (written out in some programming"language" representation). Is it possible to recognize our own program perceptions by introspection? Well, yes, linguistics provides examples. We can recognize ourselves making program-level decisions and we can distinguish one type of decision from others, with practice. We can even observe utterances of others (socialproduct outputs) and make good guesses as to the programperceptions involved as they recast a sentence in mid stream.

(Judd has given some examples from his child's speech.) To do so, of course, one must be at the level of linguistic principles, informed by system concepts (a theory of language).

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

You have said that everything looks like configurations when you are looking from the configuration level. Then because configurations are input to the sequence and program levels and because of my concern with categories, you have concluded that I am identifying with the program and sequence levels.

In this context, you have observed that a program-level way of dealing with lower-level perceptions such as relationships ("between") is too slow for practical control, citing this as the main reason graphical programs on computers are so slow. Nonetheless, it is clear that programs can accept as input even perceptual signals from the intensity level. On a scale of 1 to 10, if intensity > 5 do A, if intensity < 5 do B, if intensity = 5 (or if you can't decide) wait.

What I am saying from the Principle level is that programs and sequences may take as perceptual input categorizations from any level whatsoever. We can even have programmatic perceptions of system concepts (human cultures) such as Benedict's ranking as to synergy or Bateson's analysis of cybernetic structures in human cultures (negative feedback countering a tendency to "schismogenesis" or positive feedback, for example, or the steadystate culture of Bali contrasted with our schismogenic "runaway" culture). This illustrates another aspect of the "promiscuity of the categorization process."

Whence this promiscuity of categorization? Either sequences and programs can accept categories input from everywhere in the hierarchy, or ECSs of the category level (which provide inputs to sequences and programs) can accept input from every level of the hierarchy.

If there is a category level that "preprocesses" input for sequences, then the input device for a sequence ECS has a loop for recognizing each of a sequence of category signals, in order. It produces a perceptual signal after receiving a sequence of signals that it recognizes in proper order. The loop mechanism imposes no implicit categorization of different category sequences that are the same for it (i.e. from which it produces the demanded input signal). For each category specified in the loop, either the demanded category signal is presented from the category level, or it is not. No weighted summations. No type-token distinction between the category level and the sequence level.

The type-token distinction is characteristic of the input functions at every level--unless there is an autonomous category level. This is scarcely ominous, but it does seem peculiar.

On the other hand, if there is no category level, then the input device for a sequence ECS performs a series of weighted summations in the course of the temporal summation. The loop mechanism for summation over time might involve a set of input devices, arranged in temporal sequence by their connections in the loop. Each successive input device computes a weighted sum of incoming perceptual signals. Because various constellations of input signals may satisfy each part of the sequence-detector, the categorization (type-token distinction) attributed previously to the category level occurs here.

I have been wondering how the distinction between the two views, with autonomous category level and with integrated categorization, might be tested. It is at least conceivable that a given lowerlevel input signal might be part of one element in one token of the sequence and part of another element in another token of the same sequence. Since this possibility is I think not available in the category-level formulation, evidence for it would support the alternative proposal, though the converse is not true.

If the categorization required for sequence perceptions and program perceptions are performed by the input functions for these perceptions, then there is a type-token distinction at every interface between levels in the hierarchy, and categorization is pervasive at every level.

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

I think analogy and metaphor is central to all of this. for example, it appears to me that we apprehend relationships between sequences, between programs, between principles, and between system concepts by making analogies to relationships between lower, more "concrete" perceptions. My own introspection about abstract kinesthetic and visual imagery as I recently reported tells me this. Others have reported similar mental processes (Kekule, Bohr, and Einstein come to mind as famous examples). We call this "thinking."

But how do analogy and metaphor work in the perceptual control hierarchy? From the perspective of some higher-level perception, two perceptions are partly alike, sufficiently alike that their differences may be usefully overlooked. How can two perceptions (two perceptual signals, two neural currents) A and B be "partially alike"? Perhaps the higher-level ECS would have to be able somehow to look down two levels, into the constellation of inputs summed as A and that summed as B, and see what they had in common. But even that would not do, since a variety of different constellations of perceptual inputs would meet the demand for perceiving an A, or a B.

Let's assume the perspective of that aforementioned higher-level perception X, a perspective in which the two perceptions A and B are sufficiently alike that X may usefully overlook their differences. Suppose X is a sequence or program perception in which either A or B may serve as the nth term. Then A is analogous to B for purposes of perception X. Now I think we're beginning to get somewhere.

Suppose X is a program for knitting or weaving. Configuration, relationship, and sequence perceptions appear to be involved in this program. Suppose another program Y is a program for contradancing (country-dancing). The relationship and sequence perceptions in weaving program X are in part like the relationship and sequence perceptions in dancing program Y. The very great differences between the configurations two levels under X and those two levels under Y are invisible at the program level. Perhaps it is just this abstraction that underlies analogy and metaphor.

The programs are analogous because some of the terms in them are alike. On the basis of this analogy, if it were durable enough and useful enough, one might derive principles and even system concepts--explanatory stories, theories. One might even teach the programs as derivatives of the system concepts and principles. Exactly this appears to have been done in more than one culture.

So it is that a social product, a resultant of people carrying out individual perceptual control over time in cooperative social groups, can come to be the basis for the creation of analogies by those same people, or their children, analogies and metaphors by which they may set reference signals that govern their behavior.

A categorization of programs may also be involved here. A myth that explains both a weaving pattern and a dance pattern as reenactments of the annual springtime rescue of a Persephone figure from the underworld and the restoration of warm weather, say, places the weaving and dance programs in one category with respect to a principle of renewal.

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

Analogy (metaphor) and categorization seem to be mutually interdependent in at least two ways. One is that the A and B which are interchangeable from the perspective of X are categorized as alike for purposes of X. The other is that the social processes that establish conventional analogies and metaphors for a group also establish conventional categories.

You said "conventions are program-level stuff." On the one hand I must object to this on grounds similar to my objection to placing all categories on a single level providing input to sequences and programs: conventions or norms are much more pervasive than this dictum would suggest. For example, you would have to conclude that phonemes are program-level stuff. We may indeed come to that if we get around to looking at what is going on in phonology, but I don't think that is a conclusion you intended.

On the other hand, I have to concur that conventions (norms) and categories have much in common, together with analogy, and are much used at the sequence and program levels. The *statement* of a convention, or of an analogy, or of a category, is a program sort of thing, of a sort similar to your recipe for recognizing quadrilaterals. Perhaps that is how we try to teach them, and it may be even how we try to learn them, though practice to the point of ready recognition is the real key to learning.

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

Conventions (norms) are just reference signals. They differ from other reference signals only in the kind of environmental feedback on which we base them. For threading a needle, success is determined by a remembered physical relationship of thread and needle, the tip of the thread passing through the eye, without which a motivating program (sewing a button on) cannot go forward. For saying the word "thread," success is determined by a remembered social product, the utterances of other people.

A reference signal at any level of the hierarchy might be subject to social convention. The only acoustic difference between the way "we" say the phoneme /t/ and the way certain others say it is in the intensity of the burst in the neighborhood of F3 and in some transitions to adjacent vowels. Articulatorily, they pronounce it with the tongue tip farther forward, against the back of the teeth, and with more intensity of effort in the closure. This is the stuff of two different dialectal norms of English pronunciation.

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

Events are like sequences except that they are brief and so familiar as to be immediately recognized. Perhaps their input mechanisms are nearly identical, their other differences being due to their relative positions in the hierarchy. Could it be that spatial relationships have this same correlation with configurations, and temporal relationships with transitions? It seems plausible that existing structures may be replicated, in evolution, in ontogeny, and in learning--replicated, then the replicas attached at a higher level and modified.

This possibility suggests a direct physical basis for at least some sorts of analogy. It may be that homologous ECSs on different levels maintain neural connections to support the drawing of parallels between levels. This would facilitate the practice of a new input-constellation at a higher and slower level, later facility coming with transfer to a homologous ECS at a lower, faster level. Perhaps the unused capacity of the brain is at least in part made up of replicas of existing ECSs ready to be connected.
Perhaps some are running redundantly in parallel, but available in case of damage or ready to be reconnected elsewhere.
-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-
This is enormous fun, and might even lead eventually to something useful to somebody, but I'm going to have to give it up for a while. I have to focus on exams coming up the 18th, and a new writing project at BBN with a deadline in January, not to mention the usual holiday craziness. Doubtless many will heave a sigh of relief if I give it a rest and shut up for a while, but have been too polite to say so. Forgive us our obsessions, etc.

Bruce Nevin
bn@bbn.com

==

Date: Mon, 2 Dec 1991 11:33:03 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: language in use

[From: Bruce Nevin (911202 1111)]

(Martin Taylor 911127 15:00)

>my objection to your claim that Model <1> can be veridical "only in a limited >and local sense."

By the end of that little note on the epistemology discussion I had come to the conclusion that model<1> is just the model constructed at level 1 of the perceptual hierarchy, and that erstwhile model<2> or theoretical models should for consistency be called model<11> or so (depending on how you count). Hence the conclusion that model<1> can be veridical only in a limited and local sense. I could have gone back and rewritten the whole from the changed perspective I had reached by the end, but that iterative process has to stop someplace.

Physics and the unformalized theories by which . . .

>we do behave reasonably successfully in the much more complex world of >nutritious and poisonous foods, friends and enemies, and so on

. . . alike depend on the entire hierarchy and are stateable (a fortiori)
only at the highest levels of the hierarchy. It is the dependency on
the entire hierarchy, I take it, that is the reason . . .

>linguistically (e.g.mathematically) based models do a lousy job.

Relevant here are the comments re recipes for rehearsal vs learning by practice in this morning's post (from Saturday).

>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.

The unfortunate red herring of models<1-2> aside, I understand that you
want to account for human communication in all its complexity.

I think it is appropriate to focus on the structure in language as a
social product, since people set their internal reference signals for conformity to that.
It is also appropriate to investigate communication of interpersonal
relationship/attitude, personal relationship/attitude to the linguistic information
transmitted, and so on, as you wish to.

This communication uses various means such as gesture and expressive intonation in the course of using language to transmit linguistic information. However, these means are not in language because they are continuous rather than discrete and because they do not participate in regular combinations with the discrete elements of language. They can be imitated but not repeated, and they can be used with any combination of words or with nonverbal vocalizations, indifferently. The question intonation of English is a grammatical element. A querulous tone of voice is not.
Nonverbal communication that accompanies linguistic information is dependent on the latter in a way that the converse is not so: If one says "Are you going?" with question intonation (represented here by the discrete element "?") the linguistic information is "I ask you whether you are going or not." If the speaker adds a querulous tone of voice, the tone of voice is a nonverbal communication about the linguistic information and about that to which it refers: about the relationship with the recipient (made explicit in the base form by performative "I ask you" here, or by "I tell you" for assertions, or "I request that you should" or similar for the imperative, and so on) and about the recipient's going. For the converse, the speaker could explicitly state in an expressionless, flat tone, "I ask you whether you are going or not, and I am alarmed at the prospect of being alone, which evokes feelings of abandonment and worthlessness in me, and I see by your even considering leaving that you do not value my feelings and am deeply hurt" etc., etc., or "It's about time you left!" or explications of any number of other occasions for querulousness (see R.D. Laing's _Knots_ for suggestive examples), but with scarcely the same communicative effect as the expressive intonation conveys. Making underlying issues explicit might be good therapy, but it is not normal human communication, such as you wish to account for.
Expressive communication has more to do with music and art than with language. Socially established conventions are involved, but not those of language.
A satisfactory model of human communication must include and refer to a satisfactory model of natural language, but it must also include and refer to much that is not in language.
> 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.

Your words "skeletalized" and "quite well" suggest to me that you are imagining a trivialized theory of language that excludes much that I understand to be included. Linguistic information is the content about which all observers would agree, made explicit by undoing reductions according to grammatical regularities in the language. Communication in addition to this concerns attitude and relationship to matters to which the linguistic information refers and to the information itself, and is something about which observers will seldom reach unanimity, largely because there is nothing corresponding to the grammatical regularities of language to provide means of making those attitudes and relationships explicit. The best shot available is to use the method of levels with the communicants, so that they can use language to represent their relationships and attitudes for you as linguistic information. Good luck.

If I am leaving something out that you intend by "language in use" that is neither linguistic information nor communication in the larger sense I have indicated, please help me out.

Bruce
bn@bbn.com
==
Date: Mon, 2 Dec 1991 13:39:52 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: portrait photography

[From: Bruce Nevin (911202 1239)]

(Wayne Hershberger Sun, 1 Dec 1991 17:21:00 CST) --

Sorry to butt in where angels fear, etc. I was supposing (out loud) why Bill might speak of aspects of reality being hidden. If Bill assumes the point of view of a perceptual-hierarchy model, and we assume a perspective supposedly outside of both that model and that which it is modelling, then we see that the only contact that a perceptual-hierarchy model has with "boss reality" is proximal stimulation of intensity sensors.

What might lie beyond that, accessible or potentially accessible (directly or in a further mediated i.e. inferred way) by way of proximal stimulation of intensity sensors may be reflected or imaged or modelled in the connections, input devices, and neural signals on up the hierarchy from those initial input devices and effectors.

Is the fidelity of that reflection or image or model verifiable? We postulate that coherent, successful behavior (however we define that) as an outcome of ongoing perceptual control constitutes a demonstration of fidelity. But the existence of conflict and reorganization must then be a demonstration of less than full fidelity. Since everything is connected to everything else, I suppose it might be argued that the "representation" immanent in the control hierarchy is complete--the universe in a grain of sand. But completeness in the same sense must be accorded the control hierarchy of a turkey.

All of which is only to say: there are grounds for assurance that the world of forces and impacts is there, but not for assurance that one knows everything going on in it. This is different from saying that some knowledge of it is in principle inaccessible. I know of no basis for either affirming or denying that.

>>As we all know, "boss reality" doesn't really sit still for its >>picture to be taken.
>
>What!? You boggle my mind.
>
>Your sentence implies what it denies: that is, although we can
>not picture it, "we all know 'boss reality....'"

Try this paraphrase: as we all know, our pictures of "boss reality" are imperfect. (Our pictures: our snapshots, portraits, models, theories.) We know this by internal inconsistencies (conflict), and the very provision of means for revision (reorganization) in the model itself indicates that coevolutionary mutual adaptation is an aspect of that which we are modelling. A moving target indeed.

I think this formulation is not ambiguous so as to allow the pernicious interpretation entailing that we "know `boss reality'," an interpretation that I did not intend in the original formulation. It relies only on our own perceptions, and on the assumption that these reflect reality, etc., as above.

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

I wonder if it would be useful to consider Bateson's distinction between pleroma and creatura, the old Gnostic terminology by way of Jung, in place of the mind/body dichotomy that is the usual starting place for
epistemology. From this perspective, the perceptual control hierarchy in a living control system is part of a continuum of cybernetic feedback loops extending throughout "boss reality." The fact that a control hierarchy is more strictly organized than other parts of this cybernetic soup is an important distinction as regards the control (behavioral) aspects of perception, but does not bear so strongly on the receptive (observational) aspects of perception.

A sense-intensity receptor is a difference detector, as I understand it. A perceptual signal is then news of a difference as it enters the control hierarchy by way of a receptor from some other part of pleroma, and as it passes up the hierarchy being transformed into other differences that make other differences in turn. The combining of signals to make a signal of a different type is unique to control hierarchies, I suppose.

I don't know quite where this goes, and haven't time to try to chase it just now, but I suspect that others here are more familiar with the turf than I am. I would have to reread the Korzybski map/territory lecture and _Mind and Nature_, and I presume the posthumous book with his daughter Catherine _Where Angels Fear_.

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

I had come across the STM description too, probably in _New Scientist_, which circulates to me from the BBN library, and enjoyed the parallel with CT, though I read it quickly and superficially. Thanks for making the parallel more explicit.

Bruce
bn@bbn.com
===
Date: Mon, 2 Dec 1991 13:43:03 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: 2-D contro; Straw men

[Martin Taylor 911202 13:30]
(Bill Powers 911201.1930

>I don't think your example of the collapsible pointer is specified well enough to imply a >model. Where are you standing? Can you withdraw the >pointer? How do you know if the >pointer is too short to reach part of the obstacle? How do you detect when the pointer >touches something? If you feel along a wall and encounter an inside corner, what >degrees of freedom are allowed to get the pointer out of the corner? Can you see the >tip of the pointer?

I'm sorry that I can't put in words clearly the example that I am trying to use. First I used a person standing at A trying to get to B but being inhibited by a barrier. This was solved by a barrier-tracing technique, which didn't take account of what I assume to be the critical point that the controller can see A, B, and all of the barrier and the surrounding area. So I tried to make it clear by changing the situation to a person trying to slide a rod tip from A to B without it leaving the wall, because the person clearly could see the whole set of possible paths, knows the rod cannot move through the barrier, etc. I tried to set up the simplest possible situation in which a visual representation of the whole would lead one not to expect barrier-following behaviour.

But words do not seem to be a good way for me to describe such a full-knowledge situation.

When I come back on Thursday, I'll try to respond to the rest of Bill's interesting posting. I get the impression that I am following clumsily in a path he trod long ago, but which I see only dimly.

(Bruce Nevin 911202 1111)
If you restrict "language" to the situation-independent set of structures agreed by students of the communicative habits of some social group, then of course they are describable by and in language (at least in the language agreed by those students). There ceases to be an issue, and there almost ceases to be a connection between language and communicative behaviour.

I (a) don't see why you dismiss by fiat behaviour that can be described on a continuum, and (b) don't see why you describe intonation patterns as belonging to the dismissed behaviours when you do not so dismiss allophonic variation.

Again, I'm holding your posting pending my return.

Martin Taylor ==
Date: Mon, 2 Dec 1991 13:51:00 EST
From: Hugh Petrie <PROHUGH@UBVMS.BITNET>
Subject: PCT reviews

Along with Gary Cziko, I, too, will be teaching a course on PCT next semester and would appreciate receiving copies of any critiques, reviews, etc. Thanks.

Hugh G. Petrie, Dean 716-636-2491 (Office)
Graduate School of Education 716-636-2479 (FAX) 367 Baldy Hall
PROHUGH@UBVMS.BITNET State University of New York at Buffalo
Buffalo, NY 14260
==
Date: Mon, 2 Dec 1991 11:11:46 PST
From: marken@AERO.ORG
Subject: Reviews/Reinforcement

[From Rick Marken (911202)]

Gary Cziko -- Well, I looked through the reviews of my old papers this weekend; pretty depressing. There are many unpleasant ones and it's tough to follow them without the versions of the papers at each stage of the review (which I have not always saved). But I am sending you (by US Mail) the reviews of the "e coli" paper that was originally submitted to Science. I am sending the original mauscript, the reviews of that manuscript, my reply to the reviews and the reviews of the updataed version of the paper (which is essentially the same as the paper that was eventually printed in Psych Reports as "Selection of consequences" -- which I am not sending; you'll have to get it in the library -- Psych Reports, 1985, v 56, pp 379- 383).

For the sake of the net, I will just mention that the paper was about an experiment in which subject moved a dot to a target destination of their choice and kept it there. The dot moves in a straight line at a constant rate until the space bar is pressed at which point the dot moves off in a new dirction that is selected at random. The direction of movement of
the dot after a press, being a consequence of behavior, is a reinforcer. Since consequences are random, no systematic behavior is expected. Yet, subjects always press the bar in just the right way so that the dot moves to the target. Reinforcement theory is rejected -- discriminative stimulus theory is also rejected since what constitutes a "good" or "bad" stimulus (direction of dot movement) is determined by the subject.

The reviewers said:
1) control theory and reinforcement theory are really the same thing
2) you are testing an obsolete version of reinforcement theory (the "straw man" and "beating a dead horse" argument)
3) the results are due to discriminative stimuli -- not reinforcment
4) the consequence of bar pressing was not really random (the "it's bound

to be better after a press" argument)
5) the results are the result of intermittant reinforcment
6) the notion of inner purpose is nonsense
7) even if there were control relative to an inner reference, control theory can't anticipate what that reference level will be any better than reinforcement theory

Surprisingly, not one reviewer said "You're right; reinforcement theory is wrong. Organisms control their input. Beautifully demonstrated in the experiment. It's now time to abandon this unfortunate illusion and start studying what variables organisms control, how they control them and why".

I have learned a lot since writing that paper. Bill Powers and I wrote a supplement where we showed how the control model is able to control even when the consequences of action are random. We had a hell of a time getting that one published -- especially when our earlier efforts were aimed at showing that we COULD NOT build a working model of this random walk control based
on reinforcement theory -- any version. I would send you (Gary) the reviews we got on that paper if Bill says it's OK. I would especially like to send you a copy of Bill's replies to the reviewers. Again, if that's OK with Bill.

Reinforcement is a notion so central to conventional psychology that it's just not going to be abandoned without a fight -- I doubt that it will disappear within our lifetimes. But it is fun to watch the contortions
of the true believers.

Evidence that reinforcement is not even mortally wounded just came over
the net today:

Bruce Nevin posted an announcement about an "Active Learning and Control Workshop" which includes the following, rather remarkable statement:

> An "active" learning system is one that is not merely a passive
> observer of its environment, but also plays an active role in
> determining its inputs.

In a conference with this description of "active learning" we get papers with titles like this:
> Reinforcement Learning Benchmarks for Phonology Acquisition
and
> Exploration in Reinforcement Learning and Adaptive Control
Maybe we can get into a little argument about reinforcement theory over the net. I seem to recall such a discussion last year. Want to try it again?

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: Mon, 2 Dec 1991 12:53:53 PST
From: marken@AERO.ORG
Subject: PCT reviews,levels/suicide

[From Rick Marken (911202b)]

Hugh Petrie writes:

>Along with Gary Cziko, I, too, will be teaching a course on PCT next semester >and would appreciate receiving copies of any critiques, reviews, etc. Thanks.

Gary Cziko -- could you send a copy of my review package to Hugh when you
get it in the mail -- thanks.

David Goldstein writes:

>Recently I posted on the topics: method of levels and suicide.
>So far, except for Bill and Rick, no one has made any comments.
>That made me wonder why. and so I ask: Why?

I think it's because this concept of levels is VERY difficult. I still don't feel like I have a real hold on it -- except in terms of my own ability to experience a couple of different levels in my little demo. I need to learn and practice getting a person to go "up a level" in a conversation. I plan to work on this when I get some free time. I think this must be done as a "practicum" -- your discussion of the process in your post helps but there is no substitute for just doing it.
I don't know why there were no comments on suicide. Maybe this is a difficult topic for people (like the religion topic of ages past). I don't have that much experience with suicide (well, I gues you only get 1 personal experience
with it). I don't know if I've seriously considered it; though I have considered it once or twice. I met a couple people who claimed to be considering it. I was appalled (sp?) by the way clinicans delt with the problem (electro shock, imprisonment, etc -- very repressive).
I don't know what control theory has to contribute other than the general claim that it is unlikely that anyone would ever consider suicide if they were "in control" of all the variables important to
them. Those variables are likely to include "philosophical ones" -- which
I think are the system level variables.
Keep asking about it and maybe someone out there will say something.
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: Mon, 2 Dec 1991 20:23:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: boss reality

[From Wayne Hershberger]

Bruce Nevin (911202)
A cybernetic perspective is certainly very appropriate. In fact, that is exactly my point. The environment is an integral component of cybernetic systems, and Bill's model is no exception. To speak of the environment as being outside Bill's model makes no cybernetic sense to me. It is OK sometimes to linguistically "zero" the environmental part of Bill's model (he certainly has greater proprietary claims on the internal hierarchical part) just so long as we don't forget that the loops are closed through an environment. Cybernetically, the environment is part of the epistemic system.

>Is the fidelity of that reflection or image or model verifiable?

If I understand what you are saying, "verification" can not possibly entail a demonstration of any correspondence between the "model" and what you are calling "boss reality" (there is no one to bring the boss). So, it seems to me that boss reality is really a gold brick: a charming fellow who is nowhere to be found just when you need him.

Warm regards, Wayne

==

Date: Tue, 3 Dec 1991 17:01:03 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Mental dithering

[from Gary Cziko 911203.1650]

Wayne Hershberger (911201) replies to Martin Taylor (911126):

>The Umweg (detour) problem you pose is one that some living
>control systems (e.g., chickens) CAN NOT handle. That implies to
>me that they DO comprise scaler control loops. Right?

It's a remarkable coincidence that all this Umweg stuff appeared just as I was reading Wolfgang Koehler's _The mentality of the apes_ where he used the test for chimps, dogs, children, and chickens.

It turns out that chickens could handle the Umweg problem if the detour was not too great AND if their random movements brought them to a position where they could see around the obstacle (he actually said that some of the unsuccessful chickens were not random enough!). So this is quite similar to the dithering in Powers's Crowd/Gather program. This contrasted with most of the chimps who apparently did the dithering in their head and then through "insight" (a word which I find quite misleading) figured out the solution and then quickly acted on it.

Koehler makes a big deal out of the contrast between the random movements the chickens and the "insightful" behavior of the chimps. But what I see is that the chimps can mentally dither without moving. Thought trials instead of motor trials. Certainly more convenient to dither this way, but is it really at a qualitatively different from the chickens? (Don Campbell would call it a vicarious blind variation and selective retention process.)--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: Wed, 4 Dec 1991 08:07:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: psychophysical flux

[From Wayne Hershberger]

Re: Bill Powers (911127)

Bill, if I were to draw Venn diagrams, I think I would want to
label them like this, meaning that the natural order immanent in the psychophysical flux is realized both perceptual and conceptually. Further, there is more natural order in the psychophysical flux than is currently dreamt of in our philosophies, meaning only that the subset boundaries are not fixed.

```
 -------------------------------
|     Immanent order      |
|           |
|  --------------------------  | |  | Perceptual realization |  | |  --------------------
------  | |         |
```

```
|   ------------------------     |
|  | Conceptual realization |  | |   ------------------------      |
|            |
                                   ------------------------------
```

Beyond this, I am reluctant to go, because it seems that I would then be doing what I claim we should not be doing: confusing control theory with cosmology.

However, I admit that the expression psychophysical flux does reflect my control theory perspective. When I think of "the psyche" I tend to think "reference values" and when I think of "the physical" I tend to think "disturbances." Each of these is an input to the canonical control loop giving the loop a psychical and physical pole. These poles are as inseparable as the poles of a magnet, making the canonical loop (incorporating the two inputs) a psychophysical whole.

The canonical control loop may be partitioned into separate arcs by a mechanism-environment interface, but the location of this interface is an accident of nature and does not separate matter from mind. The loop itself is NOT psychophysical in the sense of comprising a mental arc plus a material arc separated by receptors and effectors.

That is what I think--I think.

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: Wed, 4 Dec 1991 09:03:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Language modeling; model modeling

[From Bill Powers (911203.1700)]

Bruce Nevin (911130) --

There's something I'm trying to say about linguistics that's just not getting across. I suppose we'll eventually get it worked out, but I'm running out of ways of saying it. There's a kind of analysis, or maybe it's a concept of modeling, that seems to me to lie underneath the level at which you're presenting your arguments. I'm at a disadvantage here because I can't articulate this kind of analysis very well, whereas you are very good at arguing from the base of linguistics. It seems to me that you're taking for granted and using for other purposes the very capacities of perception and control that I'm trying to understand (at the same time that I have to use them, of course, in communicating).

Here may be an example that will get us a little closer. You refer to

>primitive arguments (requiring no argument words for one to enter it >into the construction of a sentence) such as "dog" and first-order >operators on primitive arguments only, such as "jump".

Why are these "primitive arguments?" How do you tell that "dog" is a primitive argument, while "jump" is a first-order operator? How can you tell that a primitive argument "requires" no further argument words?

You go on:

>For example, you might think of "time" as a relatively high-level >perception, and "father" or "dad" as a relatively immediate, concrete >perception. But "time" is a

primitive argument, and "father" is an >operator (in "A is the father of B" it requires two primitive >arguments A and B).

Again, how can you tell that "time" is primitive and "father" is an operator? I'm not asking argumentatively. Presumably, there is some way that you can look at words and see that some are arguments, some simple
operators, some operators on operators, and so on. What is it you have to be able to notice that tells you the difference?

And how can you tell that "father" is NOT an operator when the portrait photographer says "Father on the left, Mother on the right, please"?

When you say

>But "father" refers to a relationship of roles in a social system >involving highly conventionalized prerogatives and obligations that >differ enormously from one culture to another,

I wonder how "father" can "refer" to something that varies enormously from one culture to another. To say it refers to something implies to me that you can extract what that something is from all those variations. Is there an essence of fatherness? Or could it be just that this word is used to refer to quite different perceptions, and thus has no general power to "refer" to anything? And what's the difference when "father" is used just to indicate one person rather than another? Does it still "refer" to all those complicated things?

I am interested in the phenomenon we call "referring," but not just in terms of piling up examples. I want to know what it is we have to do to a word in order to make it "refer" to something else. What is the underlying process?

But even that gets me off the track. I'm trying to get at the fact that in order to know the things you know about language, you have to employ
capacities of perception that go beyond language. Those capacities are what we can build a model of (maybe), once they're identified. There's more to this than just following out substitution or zeroing rules. We need something that can follow rules in ANY context. When we have such a model, it will be able to do language, among other things. The capacities can't be identified within the boundaries of language, or that is my claim. You can't say in language what it is about a certain color that leads you to call it "orange." You can describe the coincidence of sound and sensation, but the question is still, so what? So they occur together: how does that make anything happen?

Another thing I'm having difficulty with (this seems to be Difficult Week) is your usage of category perceptions on everything. All you're doing is telling me that you can categorize every level of transformation in the model as "categorizing." Is this supposed to convince me that categorizing isn't a level of perception?

I think I haven't explained clearly enough what I mean when I say that the viewpoint from a given level is invisible. I mean that it's invisible as a viewpoint. When you look at the world from the category level, you don't notice that you're categorizing; you simply see the world you're looking at as if it, IT, is full of categories, independently of you. You see the word for a category as BEING the category.

My suspicion that I haven't got the idea across is strengthened when you say:

>If one views things from the program level, everything looks like input >to the program level, i.e. categories.

From the program level, everything looks the way programs interpret things, but it doesn't seem that you're interpreting them. You look at a category in which all the elements have something in common, and you say "Good, that's logical." You look at a category in which the elements consist of images and smells of dogs and the word "dog", and you say "That's illogical," because the rule is that categories ought to be made of things that have something in common. The logic and the illogic seem to be in the things you're talking about, not in you. You're not saying "I'm reasoning about these categories and their elements on the basis of certain premises and arriving at a conclusions about their logicality." To realize that this is what's happening, you'd have to be looking at least from the principle level, which would then itself be invisible to you. You'd just see that the reasoning is "consistent with" the rules. The consistency would be right there in the rules; it wouldn't seem to be a principle, an opinion of yours.

When you look at the transformations that occur at each level of perception, categorizing as you go, you naturally see the same thing occurring at every level. Intensities categorize stimuli; sensations categorize intensities; configurations categorize sensations. Each one is an example of different equivalent sets of inputs being reduced to a single output. Yet if you consider what the neurons at each of those levels are probably computing, you realize that the first level is responding to physical energy, the second level is computing weighted sums of intensity signals, and the third level is probably extracting some sort of invariant from the sensations. Clearly, these are very different sorts of computations, and "categorizing" would not be sufficient to accomplish any of them. And to put it the other way, if the operations are all alike, how is it that they lead to qualitatively different kinds of perceptions that are also ELEMENTS of categories?

What's really happening, I claim, is that you're classifying these different transformations on the basis of a superficial shared characteristic, the reduction in number from multiple inputs to a single output signal. Classifyinbg by shared characteristics is one typical form of categorization. You can name this category: you call it
"categorizing." This may or may not be the way categories are really computed. You're reporting, in other words, on the sense of categoriness that you get when examining these lower-level processes (or to be more precise, verbal descriptions of them). This has nothing to do with what is actually going on at those lower levels, which is not categorizing. THEY are not categorizing. YOU are.

From the category level, you identify words with the experiences they stand for: there's no difference. Thus you can categorize names of programs, names of principles, names of system concepts, just as if those names were the things they indicate. This gives the impression of categorizing things at a higher level than categories. But I don't think that is possible. You are only categorizing names.

True program perception has nothing to do with categories or things or names. It's a wordless comprehension of the form of a network of contingencies. True principle perception is not only wordless, it is beyond logic and contingency: system concept perception is even further removed. The only way to appreciate what is happening at these higher levels, the only way to become conscious of such perceptions separated from the unitary field of experience, is to be quiet and experience them. They can't be put into words. The best you can do is arrange for concrete examples of them at lower levels, concrete perceptions and descriptions in which can be seen, if you adopt a perceptual mode of the right level, the program or principle or system concept in question. Even writing out the steps of a program does not provide perception of a program, unless you have the capacity to see in those steps a network of contingencies, with all branches simultaneously present (regardless of which branches are actually taken). You see not only the path that was taken, but the path not taken, and why it was not taken. All in silence.

The higher levels are always present and in effect during word manipulations. But they aren't obvious unless you manage to see them from above, instead of only experiencing their effects from below.

There's more to say, particularly about social "norms" and what I think of as the unconscious reification of system concepts, but there may be an indirect way of getting at this. In your comments to Wayne Hershberger on the boss reality, you gave me an inspiration. I'll take off from this inspiration and direct some comments to Wayne about modeling; I hope you will read over his shoulder and be translating into terms appropriate to our discussion
. --
Martin Taylor, I'll wait until you get back to take up our thread again.

Wayne Hershberger (911202) --

Bruce Nevin has reiterated the basis in the control-system model for entertaining the concept of a boss reality. Your response basically says that if it's impossible to find the correspondence between the boss reality and perception, why bother with the concept?

You open that comment with

>The environment is an integral component of cybernetic systems, and >Bill's model is no exception. To speak of the environment as being >outside Bill's model makes no cybernetic sense to me.

Your comment and Bruce's finally, maybe, perhaps, have joggled me into the right point of view for explaining my recalcitrance and possibly bringing our mysterious controversy to an end.

Yes, in my model there is always an environment and a behaving system. Neither makes sense without the other. I have always taken both into account. So follow me as I outline a chain of reasoning, and see if there is any point where you detect a weak link.

****

We're being modelers now. Imagine a sheet of paper on which we draw two boxes, an Environment on the left and an Organism on the right. We don't need to model the environment; that has already been done better than we could do by physics, chemistry, and if you want to include raw meat, anatomy and neuroanatomy. We can put physical variables into that Environment together with all the laws that express relationships among them.

What we're trying to model is the organism part. So we draw two arrows: once from the environment to the organism representing effects the environment has on the sensors of the organism, and one representing effect the output devices of the organism have on the environment. We are sitting up here with a good view of the paper, so we can see what is in the environment and what we're putting into the organism.

The challenge is to build a model of the organism so it will interact with the environment exactly as the real organism does. This means that basically we can give the model no help other than to provide it with the functions and interconnections that will, by their operation, generate some sort of behavior. When we guess wrong, we find that the functions and interconnections do SOMETHING, but it bears no resemblance to real behavior. We just keep fiddling with the model until it behaves correctly. This leads us to a hierarchy of control systems and so on.

If this model is to be complete, however, it has to reproduce not just behavior, but experience. In other words, the physical environment over on the left has to appear to this model just as it does to us. If we see intensities, the model has to see intensities. Simple receptors excited by various forms of physical energy will do for that. If we distinguish sensations in which different intensities are interchangeable, the model must do so. No problem: weighted sums seem to make sensation perceptions depend on physical variables as they should.

As we go higher, the problems become tougher, but we know what we're working toward. We want the model to contain signals representing configurations, transitions, events, and the rest, because we can see the world in such terms. We can't just tell the model about such things, of course; it has to contain the equipment that will, all by itself, derived such perceptions from its inputs. At the moment we're pretty far from being able to do that, but we can at least draw boxes into the model showing where we will put the machinery for deriving the signals once we know what it is. As we know what the signals have to correspond to in our own experience, we can label them: "event perception," "relationship perception," "category," etc., corresponding to our subjective analyses of private experience. The model has to have those same private experiences. It has to have ALL the private experiences that we can discriminate into "natural kinds." That includes thought and reasoning.

If we now want to go FAR beyond where we are in the process of building this model, we may want to ask about epistemology. From our perch above this sheet of paper, we can see both the physical variables in the environment and the perceptual signals inside the organism model, the model of the person. It's perfectly clear that the perceptual signals are derived in systematic ways from energy fluxes connecting the physical variables to the sensors. As we fill in the boxes, we come to understand the details of that correspondence: just how an object in the environment, through the properties of light and optical devices, and
through the photoneural receptors, comes to give rise to signals indicating its size, its distance, its shape, its orientation, and so on.

But now we come to the crux of the problem. We want to let the model figure out what there is externally to it that corresponds to its perceptual signals. For example, the object it is looking at is actually a hologram, and all that actually exists in the environment is a set of wavefronts of light that don't actually originate at the surface of an object. How does the model go about checking into the reality of the object? WE have no problem; we can see exactly what is going on. But how can the model figure it out without us to whisper in its ear? The model doesn't necessarily understand holograms (this has to be a model of any person).

One way is for the model to extend a limb to bring its visual image into the same region of visual space as the apparent object. If no contact is felt, the object could be considered intangible (that being what intangible means). But is it an intangible object in that position, or is there no object at all? Is this sxome kind of plasma object, or a less familiar trick of nature?

Solving this problem would clearly require a lot of sophistication and experience on the part of the model. It would have to compare what one set of sensors reports with what another set reports. It would have to form hypotheses and test them by performing appropriate acts. In the end, it would probably narrow the possibilities down to a small set, and on the basis of preference or niceness or some general principle, pick one of them as the answer.

Would it pick the same answer we would give from our omniscient point of view? Possibly, possibly not. In truth the model would have to know everything we know about the environment, and interpret its information exactly as we interpret it, and know what operations take place inside its own perceptual functions (which are not represented in the signals) to arrive at exactly the correct conclusion about what corresponds to any of its perceptual signals.

There is one thing we can be certain that this model can't do. It can't rise out of the plane of the paper and peer across at the environment model to see what is going on there. We have given it no abilities that would allow it to see the environment except through the raw sensitivity to energy at its input sensors. The line separating it from the environment is a barrier that can be crossed only at the most primitive level, by physical energy.

So for this model, as we have constructed it, WE can know for certain how its perceptual signals correspond to what is happening in the environment model, but IT can't know for certain. All it can do is entertain possibilities. One of those possibilities might be absolutely correct. But it can't know which one, if any.

So that is the epistemology of the model. Now what about our own?

If this is indeed a model of a human being, if we've got everything right, then it is a model of the observer, of ourselves. It is a model of us sitting up here and looking down at a sheet of paper on which there are diagrams of an environment and of a nervous system. The model has eyes and limbs; they are models of our eyes and limbs. The model has sensors and neural signals which are supposed to represent our own sensors and neural signals. The model, if it were looking at a sheet of paper with diagrams on it, would know of those things only in the form of neural signals inside itself. As the model can't rise out of the plane of the paper to see what is really in the other diagram, the diagram of the environment, so we can't rise in a fourth dimension out of our brains, to peer at whatever it is that is causing our neural signals. As the model can't sense the internal workings of its perceptual functions, and use that information to deduce what is causing any given perception, so we can't deduce the transformations that lie between the environment and our perceptions.

The model might conclude correctly that it doesn't have access to an authoritative picture of the environment model; it could reach this conclusion simply by noticing that several plausible alternative interpretations exist. On that basis, it might decide that there is no point in guessing about a boss diagram that it now realises it can never experience directly. It might decide that all it can do is compare one perception with another, and take that as the beginning and end of reality. The boss diagram is an unneccsary frill, a religious superstition; it is to laugh.

Of course we, sitting up here, would laugh at that, knowing what a mistake it is. There really is a diagram of the environment there, and it really does have a particular state, and the model hasn't been so far off the track as to be completely hopeless. At least it could survive in its interactions with the environment on the basis of what it has deduced. What it thinks it is controlling is at least equivalent, in the necessary ways, to what it is actually controlling. It may have omitted a conformal transformation or two here or there, but because it omits the same transformations from perception of its own actions, the two mistakes cancel for all practical purposes. And if it gives up now, assuming that all there is to be known exists already in the perceptual world it has constructed for itself, it's going to miss most of the fun.

And what of us? We sit up here, experiencing our own perceptions, and debating whether or not they are connected to a physical world, and if so what kind of physical world. If we believe what the model of the person seems to imply, then we are in the same fix it is in: we experience our perceptual signals, but there is nobody sitting in a higher place still who can tell us what the environment diagram really looks like. We have to figure it out on our own, each in an individual private world.

So that's where my epistemology comes from. It comes from trying to think of a model that behaves and experiences like a person, and is built the way a person is built with sensors and a nervous system and effectors. the final step, to my personal epistemology, is simply an application of the model to myself. The model contains my best understanding of how the nervous system on the right, and the environment on the left, work and interact with each other. If I now don this model and imagine that I am experiencing the world from inside it, I transform my understanding of the physical world that seems to surround me. I realise that a very plausible thing to say about it would be: it's all perception.

But it is not implausible to add " ... of something else."

Best to all,

Bill P. ==
Date: Wed, 4 Dec 1991 11:47:07 MST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Ed Ford <ATEDF@ASUACAD.BITNET>
Subject: suicide

from Ed Ford (911204.1150)

As I began to reflect on David Goldstein's request for comments on suicide, I realized I hadn't been on the net for over a month.

I've dealt with suicidal people for many years and control theory seems to give a much clearer understanding, at least for me, as to what is going on. I think the key is to look at the reference signals or various goals, the things a suicidal person wants. Conflict is driven by what we want, our goals, as we compare them to where we are in relation to achieving those goals. As we increase the number of our incompatible or unattainable goals, as we perceive things as getting worse or we imagine or perceive our world falling apart or ourselves as less competent in handling our conflicts, and we continue to reorganize with no successful outcome, our belief in our own ability to reduce our conflict, find at least some satisfaction, begins to diminish. Suicide becomes an option, an option that allows for no more pain, of relief from the struggles in life.

I think that the key to helping others is to teach them how to rebuild that belief in self, to learn to choose internal goals over which they have some control and to develop the skill to satisfy those achievable goals. Helping a person restore that loss of belief in self demands someone who the suicidal person perceives as not only caring, but who believes that he/she (the suicidal person) can make it. The therapist has to teach the client how to move toward goals important to the client but which are possible to achieve. Early this year I had a client (who had been involved in drugs, with a highly crazy family ((his real mother had committed suicide and his step-mother was an alcoholic, and in and out of a mental hospitals))) who shot himself as I was talking with his wife by phone during an ongoing crisis in their marriage. He just couldn't handle the fact she didn't want to stay married to him (something over which he had no control but was trying to control for anyhow). Fortunately, he was a poor shot, and, after some counseling, he is back home in Ohio, surrounded by family and friends. He went back to where his extended family lived and people cared about him. The last I talked with him, he's was doing well.

I've found as people begin to understand how they're designed as a living control system, that they have control over their own goals, priorities, standards, and choices, and how they perceive things, they can be more easily taught how to manage themselves in a more efficient and satisfying way. It really comes down to learning the skills of operating your own system and not trying to control another system.

Ed Ford   ATEDF@ASUVM.INRE.ASU.EDU
10209 N. 56th St., Scottsdale, Arizona 85253   Ph.602 991-4860
==
Date: Wed, 4 Dec 1991 14:00:57 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: why arguments, etc.

[From: Bruce Nevin (911206 1210)]

(Bill Powers (911203.1700)) --

>Why are these "primitive arguments?" How do you tell that "dog" is a >primitive argument, while "jump" is a first-order operator? How can you >tell that a primitive argument "requires" no further argument words?

>Again, how can you tell that "time" is primitive and "father" is an >operator? I'm not asking argumentatively. Presumably, there is some way >that you can look at words and see that some are arguments, some simple >operators, some operators on operators, and so on. What is it you have to >be able to notice that tells you the difference?

All of this comes out of naturalistic investigation of language as social product--as utterances that we may compare with one another.

After you identify what is repetition and what is not, and establish some way of representing the different morphemes of the language so as to keep them distinct (phonemes, orthography, in principle perceptual
signals corresponding to something like phonemes or semisyllables), most of the work concerns what can cooccur with what (distributional analysis).
By this analysis, the linguist groups morphemes into classes and subclasses. When we say class A can occur in the environment before class B or between B and C, etc, we say that any morpheme in the given class can occur in sentences next to (many of the) morphemes in the other classes. One can then make generalizations about sets of utterances by describing the sets as sequences of morpheme classes.
This is all analysis of observables, behavioral outputs. Perhaps I need to motivate that. The motivation is to find out what the social conventions are. The social conventions are internalized by all the members of a speech community, with a high degree of unanimity, and they control their language perceptions with reference to those shared conventions. Therefore examination of behavioral outputs can disclose the social conventions or norms. Even better, we can examine our own imagined behavioral outputs,

and interrogate others about theirs. "Can you say `the acid was washed in polypeptides?'"
We find a lot of sentences that satisfy ntuples of morpheme classes ("sentence forms") such as the following:

Sentence Form	Examples
T N1 t V T N2	The dog ate the cupcake.
	The insight will illuminate the theory.
it t be T N1 wh-R t V T N2	It was the dog who ate the cupcake.
	It is the insight that will illuminate the theory.
T N t be P T N	The chair is in the corner.

. . .

The morpheme classes are variables over sets of morphemes. Some constants appear in these formulae among the variables, for example "be" and "wh-" in those shown above. (T is definite article, t is tense, R is pronoun, the other symbols I think are self explanatory.) Each sentence form is a set of sentences that have the same form or structure at this level of analysis.
Then we find pairs of sentence forms, such that for each morpheme class of one there is a corresponding morpheme class in the other sentence form (though not necessarily vice versa). At this point we move from morpheme classes back toward the individual words. For some of these pairs of sentence-forms, when we find a satisfier in one, we find a corresponding satisfier in the other. Thus:

```
T   N1  t V T   N2  it t be T   N1  wh-R t V T   N2
The dog ate the cupcake.   It  was the dog who  ate the cupcake.
```

In addition, when language users say that a satisfier of one is nonsensical, or peculiar, or jocular, or restricted to certain contexts, they make the same judgement about the corresponding satisfier of the other:

```
T   N1      t V T   N2     it t be T   N1      wh-R   t V T   N2
The cupcake ate the dog.   It  was the cupcake which  ate the dog.
The vacuum  thought the doctor. It  was the vacuum  which thought the doctor.
```

(The difference between who and which is already understood at a prior level of analysis. "`The vacuum!' thought the doctor" is a member of a different sentence form.) This is what Harris talks about as acceptability or likelihood. It gives a diagnostic or criterion for establishing a set of mappings from the set of sentences into itself (transformations).

By further analysis of the network of transformations, we can break them out into elementary sentence-differences. These turn out to be of four kinds:

Transpositions	I saw John	-- John I saw
Changes of morpheme shape	in a quick manner	-- quickly
Zeroing	Mary left and I left	-- Mary and I left
Increments	She left	-- I thought she left

Then it turns out that the increments can be partitioned from the rest, so that the increments are words entering successively into the construction of a sentence, and the others are reductions that take place at the time each word enters. It is this that gives the partition into an informationally complete report sublanguage (with much redundancy, no paraphrase, and many sentences that are unspeakably awkward and unconventional) plus a set of reduced sentences that provide paraphrases and the rest of the language.

You can tell that "dog" is a primitive argument because its entry into the construction of a sentence does not depend on the prior entry of any other words. You can tell that "jump" is an operator because its entry into the construction of a sentence *does* depend on the prior entry of another word, one primitive argument. Similarly for "time" and "father."

>And how can you tell that "father" is NOT an operator when the portrait >photographer says "Father on the left, Mother on the right, please"?

There is another sentence from which we can derive this:

The one who is the father (of the child) on the left . . .

The motivation for this is to have just one "father" morpheme rather than two, one in one class and one in another, but with the same meaning. "Father" in its seeming noun role can be derived easily from a source in which it is an operator, as above. "Father" in its predicate role, as in the following, cannot be derived from a source in which it is a primitive argument:

John is the father of Mary

(This "is" is not an operator, it is a carrier for the tense morphology with more stative operators, including prepositions, adjectives, and relational nouns like "father.")
>>But "father" refers to a relationship of roles in a social system >>involving highly conventionalized prerogatives and obligations that >>differ enormously from one culture to another,
>I wonder how "father" can "refer" to something that varies enormously
>from one culture to another. To say it refers to something implies to me >that you can extract what that something is from all those variations.

I'm sorry, I joggled two levels of abstraction there. The reference to cultural variation was only to emphasize that the perceptions to which our word "father" refers (with which it is somehow associated in the perceptual hierarchy) are cultural, normative, based on social convention, and at a fairly high level in the perceptual hierarchy, rather than at the level that we perceive "directly" as physical objects. "Father" refers to a socially defined role relationship. When we translate a word from some other language as "father" we take the biological part of that relationship as basic (even though it is less basic than the social prerogatives and obligations in our practice).
The prerogatives, obligations, etc. associated with that word in the other language are typically quite different, for example, much of the
"fathering" may be done by the mother's brother, while the "father" does the "fathering" for all of the children of his sisters. (I believe this is one historical basis for "problems" of absentee fathers in Black American communities.)

>Is there an essence of fatherness? Or could it be just that this word is >used to refer to quite different perceptions, and thus has no general >power to "refer" to anything? And what's the difference when "father" is >used just to indicate one person rather than another? Does it still >"refer" to all those complicated things?

Perhaps "referent" is the wrong term, since this is a generic reference to "fathers" rather than to a specific referent "my father" or "(my) Dad." The meaning of the word is certainly not limited to the latter. Yes, I think it has both aspects of meaning.

>I am interested in the phenomenon we call "referring," but not just in
>terms of piling up examples. I want to know what it is we have to do to a >word in order to make it "refer" to something else. What is the >underlying process?

Operator-argument dependencies among words correlate with dependencies among nonverbal perceptions. The information structure or "report structure" in language (or in "things we know" in language) correlates with networks of dependencies among nonverbal perceptions, including our expectations, such that we take the structure in the former (to which we can give utterance in various discourses, and which we have learned in part from hearing and participating in various discourses) to be an account of structure in the latter. Primitive arguments correlate for
the most part pretty well with identifiable perceptions--they're mostly concrete nouns. Many first-order operators, which require only primitive arguments, are also largely fairly straightforward. Other words are less easy to identify as individual words in a straightforward way with individual perceptions. "Thought" is an abstract noun, but the sentence "I thought that John left" is referrable to perceptions, and it is as a generalization of many such sentences that we get some sense of the reference of "think". Does that mean there is an elementary control system for a nonverbal perception to which "think" is correlated in the same way that "triangle" seems to be correlated with a configuration perception? I doubt it.

>in order to know the things you know about language, you have to employ >capacities of perception that go beyond language. Those capacities are >what we can build a model of (maybe), once they're identified. . . .
>We need something that can follow rules in ANY context. When we have such a >model, it will be able to do language, among other things. The capacities >can't be identified within the boundaries of language, or that is my >claim. You can't say in language what it is about a certain color that >leads you to call it "orange." You can describe the coincidence of sound >and sensation, but the question is still, so what? So they occur >together: how does that make anything happen?

Perhaps the outline of the process of ferreting out the structure in language that I sketched above can indicate what you're looking for here? The capacities of perception required to do linguistics are capacities of perception that we can and do use for all sorts of other things, yes. There is nothing unique about language in that regard.

We can talk about those capacities in terms of a child learning a language. However, that account alone does not suffice to answer the last question. Language already exists as a highly evolved tool when the child comes to learn it. The child doesn't invent its structure or the loose correlation of that structure with nonverbal perceptions. All that comes as a given, a social inheritance, and what the child learns is how to use it as others in the speech community do.

The rest will have to wait. I've already overrun my available bandwidth.

Bruce
bn@bbn.com
==
Date: Wed, 4 Dec 1991 14:20:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: model modeling

[From Wayne Hershberger]

Bill Powers (911203)

How in the world do you manage to maintain all your various CSG dialogues, virtually on a daily basis!?

Having said that, let me give you something more to occupy your time. I've just now read your recent post. Although I've not got time now to address your comments in detail, a thought has occurred to me that you might agree would prove helpful. It seems to me that the little man model may be said to know where the target is, while knowing nothing either of computers or of yourself, his creator. That is, it seems to me that if you were to rewrite your essay, putting the model in a computer where it can function as a simulation, the epistemological implications might appear more clear cut.

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: Wed, 4 Dec 1991 23:53:00 GMT
From: "Apeiron, Kurt Christensen,PAS" APEIRON.TECH@APPLELINK.APPLE.COM
Subject: Unsubscribe

unsubscribe "Kurt K. Christensen" <APEIRON.TECH@APPLELINK.APPLE.COM>

Dear all,

I find the discussions interesting, but you guys are far too prolific for me to keep up in my current situation.

Later...

Kurt ==
Date: Thu, 5 Dec 1991 13:24:20 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: OK, answering myself

[From: Bruce Nevin (911207 1222)]

(Bill Powers (911203.1700)) --

First, an observation about the precis of linguistic analysis that I sent yesterday. It involves nothing special or unique to language. It could plausibly be that it is paralleled pretty closely in the process by which a child comes to know a language. (Martin's reference to literature describing neural nets "learning" word classes, etc., supports this. Yes, if you could send me the paper I would appreciate that, Martin.)

The main difference is that the child learning the language is also concurrently learning the correspondences of constructions in language with other perceptions.

The historical reason linguistic analysis does not exploit these correspondences, other than for judgements of acceptability, is that it has not had as a context a science of how people control these other perceptions.

The fact that so much can be determined without broader reliance on our knowledge of meanings is an indication of how very much is conventional in language. It is the conventional aspect that emerges most easily and clearly in the analysis: so clearly that we can even use it to account for an important and rich aspect of meaning. And indeed, people can use the conventionalized structure of language to transmit this aspect of meaning precisely because it is conventionalized and socially available, rather than dependent upon private experience of perceptions. And this is why language provides indications where to focus attention in one's own experience, rather than providing any replication or image of the author's experience itself.

I've made a copy of the last two chapters of Harris's _A Theory of Language and Information_ which I will mail to you. You may be interested in his discussion of such matters as reference.
-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=-=+=>the viewpoint from a given level is invisible.
This is what I mean when I say "ask a fish about water."
>When you look at the world from the category level, you
>don't notice that you're categorizing; you simply see the world you're >looking at as if it, IT, is full of categories, independently of you.

> You see the word for a category as BEING the category.

I have broken this in two parts because I do not know that the indentification between words and categories that they name is as intimate you say.

Maybe it is only that not all words are names of categories.

Someone protested a while back that there are categories for which we have no words. Are there words that do not label categories?

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

>True program perception has nothing to do with categories or things or >names. It's a wordless comprehension of the form of a network of >contingencies.

The inputs to an input device at the sequence level are category perceptions, as I understand it. Sequence perception has nothing to do with categories, it's a wordless comprehension of the form of a succession. The fact that categories are terms in the sequence is made invisible by the input device. Similarly for the program level: the fact
that sequences and I guess categories too are terms in the network of contingencies is made invisible by the input device that constitutes a network of contingencies out of the input signals.

All I was suggesting was: what if the input device for a sequence-level ECS does the constituting into categories as well as the constituting into sequences? What are the reasons for supposing an autonomous category level?

It sounds like a reason for you is your identification of categories with category names (words), as for example when you say

>From the category level, you identify words with the experiences they >stand for: there's no difference.

(One should broaden this to symbolization in general, i.e. association
of a category with some nonverbal perception as its index or label.) I have questioned the intimacy of the association of categories with words, but it might work for most primitive argument words and first-order operators.

But you have agreed with my questioning a proposed input function that sends up the same perceptual signal when presented with either the word "line" or the configuration signal from a line-detector ECS. Do you no longer agree in questioning that proposal?

Another reason for supposing an autonomous category level might be the process of learning that starts with categories artificially constituted by rehearsing some verbal or symbolic recipe or formula (count the number of sides, etc) and concludes with immediate recognition of the category without computation at higher levels of the hierarchy. (Response-delay experiments presumably underwrite this distinction.) There must be something that does this latter recognition, below the sequence and program levels where the recipes live.

It could be that such rapid-recognition categories are recognized by the input devices of the sequence-detectors and program-detectors that "use" them. We are never called upon to recognize them but for some purpose.

This discloses another reason: parsimony. Suppose I gain skill in immediate recognition of icosahedrons, regular and irregular. This was in context of some higher purposes, even if only carrying out some researcher's bizarre experiment and getting some compensation for it. Absence of a category level would predict that it would take me as long to learn the same category in some different context requiring different sequence and program perceptions, and that my recognition of the category might be better or differently skilful in one context than in the other. It does not and, so far as I know, it is not. In general, it is implausible that each sequence and program detector that "uses" a given category has its own separate copy of the means for recognizing that category.

So I guess I've talked myself into a category level. These are answers of the sort I was looking for. (Maybe it's time for the old "Oh is that all you were saying" response.)
===
Date: Thu, 5 Dec 1991 13:12:42 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Neural Hierarchy

[from Gary Cziko 911205.1300]

To Bill Powers, Joel Lubin, any other neurally sophisticated CSGnetters:

I attended a lecture yesterday by Gerald Westheim (I forget from which

California university; there are too many) who spoke about visual hyperacuity. He showed lots of data of how subjects can make very fine visual discriminations which are actually much finer than the "grain" provided by the density of retinas. Very interesting.

He didn't offer much in terms of explanations of how this was achieved, but did mention that as you move from the geniculate nuclei to the visual cortex you get an increase in the number of participating neurons by a factor about 100 (relationship is about one to one between retina and geniculate nucleus).

Now, all the diagrms I've seen of the PCT hierarchy (as well as Marken's spreadsheet model) CONVERGE as you move up levels, sort of like a classic Christmass tree shape with systems concept sitting on top as the bright star). Indeed, this seems to be a basic characteristic of a hierarchy (as opposed to a heterarchy, whatever that is). Can someone explain to me why there instead appears to be this DIVERGENCE as you move up into the visual system? How can the Christmas tree be upside down in HCT?--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: Thu, 5 Dec 1991 15:54:27 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>

Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From:
Joseph Michael Lubin <jmlubin@PHOENIX.PRINCETON.EDU>

[from Joe Lubin 911205.1400]

[to Gary Cziko (911205.1300)]

> I attended a lecture yesterday by Gerald Westheim (I forget from which
> California university; there are too many) who spoke about visual
> hyperacuity. He showed lots of data of how subjects can make very fine
> visual discriminations which are actually much finer than the "grain"
> provided by the density of retinas. Very interesting.

> He didn't offer much in terms of explanations of how this was achieved, but
> did mention that as you move from the geniculate nuclei to the visual
> cortex you get an increase in the number of participating neurons by a
> factor about 100 (relationship is about one to one between retina and
> geniculate nucleus).

From de Valois and de Valois' _Spatial Vision_ (Oxford Press, 1988), p. 173:

Of the many types of hyperacuity, two of particular interest are Vernier acuity and stereoacuity. In Vernier acuity a subject is asked to judge whether or not two line segments, end to end with a gap between them, would, if extended, form a continuous line. One can do almost as well with a variant of the task consisting of just two dots, which are judged to be lined up vertically or not. A lateral offset between the lines or dots as small as 2 to 5'' of arc has been found to be detectable, as against 1' of arc between then finest resolvable bars in a classical acuity test (Westheimer & McKee, 1977). The same threshold is found in stereoacuity, in which the precision with which an observer can align two lines in depth is measured in binocular vision....

The similarity of the terms acuity and hyperacuity is unfortunate, because it suggests that similar processes are involved -- which is not the case, as pointed out by

Westheimer (1979). In the classical acuity tests the fundamental nature of the task is to discriminate between two different figures -- to tell whether one or two lines are present. In terms of the CSF [contrast sensitivity function], acuity is related to the highest spatial frequency detectable. On the other hand, the hyperacuity tests measure the observer's ability to determine the LOCATION of the relevant stimulus characteristics. Hyperacuity could thus be related to the determination of the PHASE of the sine wave components involved (although Westheimer, 1977, deems it unlikely).

In principle, localization, or determination of phase can be carried out to any degree of precision, limited only by noise, given a sufficient amount of information. This is not true for acuity. Hyperacuity is not limited by the spacing of the spatial smaples in the same way as is two-point resolution. As discussed in Chapter 2, the spatial separation of foveal receptors would by itself determine the high spatial frequency limit of the CSF and thus acuity (even if optical factors did not). However, the receptor seapration does not set the limit for the localization of some feature, or the determination of the phases of detectable (lower) frequencies in the pattern, which is the basic capability erquired for so-called hyperacuity. A simple linear model of visual processing would predict that hyperacuity would be far superior, by a factor of ten or more, to two-point or grating acuity limits (Geisler, 1984). One thus need not postulate some complex, higher-order visual processing capability to account for fine hyperacuity thresholds.

I have not read the (Westheimer, 1977) paper which deems the direct use of phase unlikely, but this would be my thought also. I don't believe phase information would be sufficiently accurately preserved for "large" visual objects like lines by the time the signals arrive in visual cortex where the work is done. My feeling is that while this type of mechanisms could work in the retina, the phase resolution of the cells in visual area 1 (V1; primary visual cortex) is not sufficient. I may be wrong however. I have built a mechanism, called the Boundary Contour System (BCS) by Grossberg et al. which maps very nicely into the functional and cyto-architecture of the first few stages of visual cortex, and which constructs emergent representations of visual objects based on the statistics of supporting and disconfirming visual splotches. It is here where phase information might be CREATED and used for hyperaccurate determinations (if necessary).

Using the phase information is probably not necessary, however, because, as you mentioned, the magnification (in terms of number of cells devoted to a small patch on the cornea) of the cortical representation to the retinal representation is high. Simple cells in V1 use high and low frequency information to represent macroscopic location, orientation and phase. There are many of these cells for each small patch of retina (espacially in the fovea).

> Now, all the diagrms I've seen of the PCT hierarchy (as well as Marken's
> spreadsheet model) CONVERGE as you move up levels, sort of like a classic
> Christmass tree shape with systems concept sitting on top as the bright
> star). Indeed, this seems to be a basic characteristic of a hierarchy (as
> opposed to a heterarchy, whatever that is). Can someone explain to me why
> there instead appears to be this DIVERGENCE as you move up into the visual
> system? How can the Christmas tree be upside down in HCT?--Gary

The real point is why doesn't the PCT hierarchy doesn't look like the visual hierarchy. If you consider the fact that our sensory transduction apparatus can be, should be, and is quite small, and then realize that small as that receptor array is, it is extracting a tremendous amout of data. Much of the data is, however implicit. (Explicit data would be intensity of neural responses; implicit data would be psychophysical color at position (x,y) on the cornea.) This implicit data is unfolded by a variety of circuits in cortex. So we can view the visual system as consisting, longitudinally, of two divisions:

(i) A divergent processor of the visual stimuli which extract features along a variety of dimensions, and

(ii) a (semi-)convergent processor which creates increasingly abstract (and holistic) representations.

The convergence (of the object recognition aspects of the visual system) probably terminates in the inferior temporal cortex, where all cells have very large receptive fields which include the fovea. Such cells are thought to respond maximally to stimuli like faces and hands.

So, if you take the leaves of the tree to be in visual area 1 or 2, then there is a convergence, and Christmas is saved, at least for this year.

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

===
Date: Thu, 5 Dec 1991 17:28:15 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: mmt@DRETOR.DCIEM.DND.CA
Subject: Re: Neural Hierarchy

[Martin Taylor 911205 17:00]
(Gary Cziko 911205.1300)
>
>
>I attended a lecture yesterday by Gerald Westheim (I forget from which >California university; there are too many) who spoke about visual >hyperacuity. He showed lots of data of how subjects can make very fine >visual discriminations which are actually much finer than the "grain" >provided by the density of retinas. Very interesting.
>
Gary asks how this can be done. Actually, it is an example of something that is so common as almost to be called "ubiquitous" in the nervous system. The key to understanding it is not to see the individual sensors, such as the retinal cones, as resolution cells that are isolated from
each other and from their own recent history. For example, I imagine that Wertheim talked about vernier acuity (the ability to detect small offsets in a linear edge) or the ability to detect a fine wire, each of which has a limiting resolution on the order of 1 second or arc, whereas
the retinal cone array has a cell size of around 1 minute of arc. In either case, it is the population of affected retinal cones that provides the information, not the individual cones. Colour is another example. We
can discriminate so finely that some people misinterpret the data to say that we can duscriminate milions of different colours. But we have only three different types of cone (labelled red, green, and blue, although the red and the green differ by only a very small amount in the frequency
to which each is maximally sensitive). And there are VERY few blue cones, by comparison to the others, almost none in the central fovea.

Colour makes perhaps the easiest illustrative example, which can be extended to other cases, such as pitch discrimination in hearing, tactile discrimination, ... What happens is that a particular colour patch excites quite a few cones, some of which are red, and some green. The ratio of outputs of the two types changes quite rapidly as the frequency of a spectrally pure colour passes through yellow (equal red and green). Each ratio signifies a particular colour (I am not forgetting all the important context effects, but I am ignoring them for this discussion).
So although there are only two different receptor types (forgetting the rare blue type, which affects the colour perceived in larger patches), very many different ratios among their outputs can be discriminated.

Now think about the spatial array of receptors in the retina. Suppose there is a black-white edge crossing some part of that array. As the
eye moves (and it does, at some 10s of Hz, like 50-70 Hz), the edge moves coherently back and forth over this array, allowing the detectors to react according to the proportion of time they see black or that they see white. The precision of locating the edge can thus be much better than the size of the retinal elements. But one can do better than this, by noting that the coherence is different if the eye motion is across the edge as compared to if it is along the edge. This difference allows the array to learn which receptors are aligned in which directions. Aligned receptors tend to be preferentially

connected (I don't know whether this happens in the retinal, or higher up in the cortex, where the Hubel and Wiesel line and linear motion detectors are found), so that they can work in concert to improve still further the accuracy of locating the black-white edge, if it happens to be linear. Then, if there is an offset in the edge, it can be detected with great precision, compared to the
size of a single cone.

This is a very quick and only crudely accurate description, but it should be enough to make it clear that the apparent crudeness of peripheral receptor systems is no indicator of the precision with which cooperating sets of them can provide information. I think the same kind of thing happens at much higher levels of abstraction, as well, such as word recognition in reading (see my "Convenient Viewing and Normal Reading", in Working Models of Human Perception, Academic Press, 1988). There, I quote Hinton, McClelland and Rummelhart (in volume 1 of PDP): "The intuitive idea that larger zones lead to sloppier representations is entirely wrong because distributed representations hold information much more efficiently than local ones." And, I may add, are more resistant to distortion and the effects of noise.

Martin Taylor ==
Date: Thu, 5 Dec 1991 17:41:16 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: mmt@DRETOR.DCIEM.DND.CA
Subject: visual phase

[Martin Taylor 911205 17:30]
(Joseph Lubin 911205.1400)

>
>I have not read the (Westheimer, 1977) paper which deems the >direct use of phase unlikely, but this would be my thought also. >I don't believe phase information would be sufficiently >accurately preserved for "large" visual objects like lines by the >time the signals arrive in visual cortex where the work is done.

Maybe so, but I was once (many years ago) shown a demonstration of the importance of visual phase in the appearance of large visual abojects.
The demonstration consisted of reconstructions of photographic images by inversions of Fourier transforms from which some information had been removed. Reconstructions from which the phase information is removed look like random noise, whereas reconstructions in which the amplitude was randomized look like quite good cartoons of the scene. So at least in a gross sense, the phase is more important than the amplitude information for object recognition. My informer, whose name
I am sorry I have forgotten (at the Swedish Defence Research laboratory
in Stockholm), said that about 75% of the Shannon information in photographic scenes is in the phase.

Martin Taylor ==
Date: Thu, 5 Dec 1991 17:20:54 CDT
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: Re: 2-D control; coordination

From Tom Bourbon [911205]
The conversation on 2-D control, and scalar control, has
drifted in and out of my world, during a series of major crashes and minor glitches with our computer. I am missing many posts, but I do have Bill Powers [911202] and Martin Taylor [911202].
For the past couple of years, I have been working on models
of interactions between hands, models and combinations of the two. One of the first ideas I explored was whether independent control loops could produce interactions, through their control of variables that are linked in the environment, that are coordinated. Many widely accepted theories of coordinated action rely on the

notion that coordination of different degrees of freedom (dimensions) in the environmant is achieved via a prior coordination, or articulation, of the control processes inside the organism.

In two simple studies, I showed that independent control systems can indeed produce what appears to an uninformed observer to be coordinated action. (Rick Marken has produced similar results. His work explored people using both hands to control two degrees of freedom in the environment.)

In the more complex of my studies, two control handles affect two cursors on a computer screen. Each handle affects both cursors, but it affects one cursor twice as much as the other cursor. Each cursor is also affected by an independent random disturbance. One person can easily control the positions of both cursors, relative to a stationary target -- in spite of the fact that, whatever the person does with each hand also disturbs the cursor controlled by the other hand. Following conventional reasoning, one might assume that the person established a plan, or schema, for coordinating the two hands, prior to controlling the two cursors, but that is not a necessary assumption.

In fact, two people can perform the task together, each using one of the handles. The results are indistinguishable from those produced by one person using two hands. That result suggests to me, as it did to Rick, that the coordination in such tasks occurs in the linkages between variables in the environment -- that any hand that uses one handle to control one environmental variable might be modeled as the output device for a single control system, controlling a single degree of freedom (a scalar variable).

PCT modeling of the task confirms that notion: two independent control models, each with a reference to control one of the cursors produce results that are nearly identical to those from two people, or from two hands on one person. Further, when a person and a model run concurrently, each controlling the position of a different cursor, the results again match the now-familiar pattern.

Whenever one system (person, model, hemisphere) moves its control device, it unavoidably disturbs the cursor controlled by the other system (whatever it might be), but the disturbed system merely eliminates that influence as part of the net disturbance acting on the controlled variable at that moment. There certainly is coordination in the actions of the systems, but the coordination comes through environmental couplings, not from the operation of a 2-D control system.

Recently, one of my thesis students (Wade Harman) defended his research, in which a person, using a mouse and a joystick, controlled 4-D relationships. An arrowhead changed in four "dimensions," each under the influence of its own independent random disturbance. Changes occurred in X and Y, in the angle of orientation of the tip of the arrowhead, and in the size of the figure ("depth"). With a bit of practice, all participants mastered the task. More interesting to us was the fact that the performance of people was duplicated by four independent control loops, each controlling one degree of freedom in the figure.

In Wade's task, the coordination of actions is obvious, but it need not imply a 4-D controller. I suspect that, contrary to conventional wisdom that coordination requires a prior linking of control processes inside the system, coordination is enabled by a process in which we functionally "carve ourselves up" into independent control loops, one each scalar perceptual signal.

I think Bill Powers was making a similar point when he described some of the inner workings of the "people" in the "Crowd" program. Now, I hope this goes out and that I will have an opportunity to follow the conversations.

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

==

Date: Thu, 5 Dec 1991 21:31:00 MST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: language, reality

From Bill Powers (911205.1000)]

Bruce Nevin (911204) --

I asked how you can tell "dog" is a primitive argument. You answered in much the way I thought you might after I realized that "you" should have been "anyone":

>All of this comes out of naturalistic investigation of language as >social product--as utterances that we may compare with one another.

>After you identify what is repetition and what is not, and establish >some way of representing the different morphemes of the language so as >to keep them distinct (phonemes, orthography, in principle perceptual >signals corresponding to something like phonemes or semisyllables), most >of the work concerns what can cooccur with what (distributional >analysis).

>You can tell that "dog" is a primitive argument because its entry into >the construction of a sentence does not depend on the prior entry of any >other words.

When I asked "how can you tell," I meant by the "you" the generic you. What I should have asked, less colloquially, was "How can your model of the linguistic system recognize that 'dog' is a primitive argument?"

My question is how the model can see that this dependence does not exist.

What you describe in this post is not language, but a linguist analyzing language. Doing this analysis requires identifying perceptual variables such as repetitions, morphemes, cooccurrance, and distributions.

>... the linguist groups morphemes into classes and subclasses.

The linguist is capable of grouping things into categories

>When we say class A can occur in the environment before class B or >between B and C, etc, we say that any morpheme in the given class can >occur in sentences next to (many of the) morphemes in the other classes.

The lingist can perceive the way in which classes occur, and from observing occurrances derive a stable sequence: any element of one class can occur in the same sequence with elements of (many of the) other classes. This is a sequence-level perception in the linguist.

>One can then make generalizations about sets of utterances by describing >the sets as sequences of morpheme classes.

So the linguist can perceive that some different sequential utterances obey the same rule.

You have illustrated the levels of perception quite nicely, but they are not in language. They are in the linguist.

The process you describe is similar to what goes on in an intelligent but naive onlooker watching a game of chess being played in silence. The onlooker identifies classes of pieces, and of moves that go with each piece, strictly by remembering past behaviors of these differently-shaped and -colored pieces. The taking of one player's pieces by moves of the opponent's pieces follows rules that can be deduced: if a legal move by piece A lands on piece B, B is removed from the board. At some point, the onlooker might feel a tremendous sense of insight:
"Hey, the horse could have bumped that little guy off the board!" At another point, the onlooker will see the same possibility and add, to himself, "but if he had done so, the one that looks like a watchtower would have bumped HIM off the board!" So the program-level concept of strategy appears in the behavior of the pieces.

Notice how the onlooker's attention is on the pieces and what they can do by way of moves or effects on other pieces. Gradually, the pieces take on properties that predict what they can do. Patterns of possibilities such as forks and castling become apparent, and with long enough study, the principles behind those patterns emerge from all these details. Formations of pieces take on strategic significance. The chessboard becomes a living thing, with Black and White surging here and there across the board leaving the casualties piling up on the sides.

This way of understanding chess, like the linguist's way of understanding language, is pure empiricism of the pre-modeling kind. It never occurs to the onlooker to ask, "Why do the horses always move in that funny pattern? What keeps them from jumping straight ahead, or farther?" The onlooker can imagine many things happening that never happen, even though it would be simple to reach out and move a horse one square, or ten squares, or to move a watchtower diagonally. But that isn't how watchtowers and horses move, so that's that. There is a social convention that says they move in these patterns and no others. At some point the onlooker will get the concept of "the game of chess."

If all the onlooker wants to do is enter into the game, this kind of understanding will suffice. But suppose the onlooker doesn't simply accept that horses move one way and watchtowers another. Suppose the onlooker expands the field of view to include the two human opponents in this game. This transforms the question. The question is now not "Why does the horse move that way?" but "Why do you move the horse that way?"

Chances are that even the expert player being asked this question will interpret it just as the onlooker had done before, and answer something like "Because moving it the other way would have exposed my Queen," or even more simply, at a lower level, "Because that's how knights move." But these are no longer the answers that the onlooker wants.

What the onlooker is trying to ask, haltingly and in considerable confusion because this is a question he hardly understands himself, is "What is it you do inside of yourself that results in the horse moving that way instead of some other way?" The onlooker has realized that the pieces could have had any shapes, the board any number of squares or other subdivisions, the moves any arbitrary constraints, the strategy any form, the game any principles; the same question would apply. The onlooker has realized that the rules and conventions are not in the pieces, but in the players.

The game is only an externally-visible consequence of what is going on in the players. There are no chess-pieces inside the players, there is no board. What is inside the players is something far less tangible. It is some set of operations that results in the apparent properties of the pieces, the apparent rules of the game.

So the onlooker is really asking questions like , "What is it that can state a rule and act so that the consequences, when described, turn out to fit the rule? What is it that can apprehend two different objects and treat them as if they were the same?"

The behavior of the chess pieces illustrates what results when capacities like these inside the players are used to perceive and affect the outside world. But what are these capacities? Is rule-following a generalization of a "fork" or a "castling?" Can the underlying processes be perceived just by examining the chess pieces and their behavior more closely, by combining and recombining the rules and strategies of chess?

No, because these are only consequences, not causes. The very following of a stateable rule is itself a symptom of a more basic process, which I describe as taking place at the program level in a brain. The very recognition of a "pawn" is evidence of configuration perception and naming or classification.

Even to begin trying to understand language, it's necessary to view language as only one example of more basic processes. I've tried to name these processes with 11 words. All these processes are used in language, but they are also used in chess and mathematics and hunting for prey. It is not language that makes these processes run, but just the opposite. I am just stating what we both agree to.

You say

>We find a lot of sentences that satisfy ntuples of morpheme
><classes ("sentence forms") such as the following:

```
>
>Sentence Form                    Examples
>=============                    ========
>T N1 t V T N2                    The dog ate the cupcake.
>                                The insight will illuminate the theory. >it t be T N1 wh-
R t V T N2      It was the dog who ate the cupcake.
>                                It is the insight that will illuminate
>                                  the theory.
>T N t be P T N                   The chair is in the corner.
```

My question is: why do we find such sentences? I don't doubt that we do, if you say we do. What I want is an explanation, a model that will, out of its own properties, generate sentences like these and not sentences of other forms. I would want the same thing no matter what the distribution
turned out to be. The observed distribution just reveals the rules of the game, the way the chess pieces can move. The rules could have had any other form. The answers to my questions are not to be found in any specific set of rules, but when found, will explain how ANY set of rules becomes effective.

You say

>You can tell that "dog" is a primitive argument because its entry into >the construction of a sentence does not depend on the prior entry of any >other words.

And I say that is describing how the chess pieces move, not modeling the system that moves them. I have nothing against becoming an expert chess player or an expert linguist, but I am a modeler.

Wayne Hershberger (911204) --

>How in the world do you manage to maintain all your various CSG >dialogues, virtually on a daily basis!?

I'm retired. Can't you tell?

>... it seems to me that if you were to rewrite your essay, putting the >model in a computer where it can function as a simulation, the >epistemological implications might appear more clear cut.

Actually that is the route I took to my present position, only it came from building real systems more than from simulations. In the late 1950s, for example, Bob Clark and I built an "isodose tracer" that used an analogue computer as a control system to make a tiny radiation probe move along curves of constant radiation intensity in the beam of a Cobalt-60 treatment machine (VA Research Hospital in Chicago). In the early stages we got some strange curves, because the long stem that held the probe turned out to be radiation-sensitive. The control system was keeping what it assumed to be sensed radiation at the probe tip constant, but it couldn't know where that radiation was being detected. The variable under control wasn't quite the one that was supposed to be under control.

I've mentioned the voltmeter effect before: the reading on the voltmeter is not the "true" voltage because the meter draws current. In my electronics ventures with radiation probes and photosensitive equipment, often incorporated into control systems, it was almost always necessary to correct the meter readings when measuring low-current high-voltage sources. Automatic control of such voltages required compensation so that the "real" voltage, not the measured voltage, was controlled.

The whole world of electronics is fraught with examples. A simple circuit board is, to the electroniker, largely imaginary. The surface appearance of the board has almost nothing to do with what is "really" going on. Every component carries in it mysterious properties like resistance, capacitance, inductance, and amplification that are never experienced directly (voltage is one example, but it doesn't feel like voltage. It feels like hell). Usually such things are known only after calculations based on the few contact points with direct experience. Yet when you assume that such things exist in some boss reality, as you must in order to make any sense of "correcting a meter reading," the result is the power to make things happen in highly predictable ways. You adjust a tuned circuit a little below resonance, so it will be exactly at resonance when you remove the capacitance of the probe you're using to measure the response. The true operation of a circuit is what you deduce would take place if you weren't measuring anything!

You've picked this diagram from the possibilities I suggested:

```
                             -------------------------------
|     Immanent order      |
|      |
|  --------------------------  | |  | Perceptual realization |  | |  --------------------
------  | |  |
|   ------------------------      |
|  | Conceptual realization |  | |   ------------------------     |
|      |
                             -------------------------------
```

"Immanent order" wouldn't be a bad term for "boss reality." From my viewpoint it has the nice implication that there can be order without our knowing what it is. But I have some more questions.

By choosing the diagram in which the immanent order extends beyond the boundaries of the realizations, you have agreed with me that there is more to know than meets the eye. By making the two realizations independent and non-overlapping, you have said that each has its own relationship to the immanent order independently of the other.

In this diagram, there's no connection shown between perception and conception, nor any indication of how these "realizations" might relate differently to the immanent order. You describe the figure as a Venn diagram. This implies that within the outer boundary there is some immanent order, and that it's simply marked off into regions, with the elements of the largest field being no different inside and outside the two "realizations."

As you didn't specify the difference inside and outside the realizations, two possibilities are:

1. A realization is simply a noticing of something that was always there, the noticing in no way altering what was always there but merely bringing it into the field of attention.

2. A realization is some transformation or projection of the immanent order, so that the realization is an invention or at least an expression of the nature of the system becoming acquainted with the immanent order.

In both cases there is an implicit relationship between a realization and the immanent order. In the first case the realization is completely passive; it is merely recognition. In the second case there is a difference between the realized and unrealized states of portions of the immanent order. Does either of these choices fit your conception?

I take it that the rationale for the term "immanent order" is that neither perception nor conception is random; that both reflect some orderliness that constrains them. Does this not imply some effect of the immanent order on the realizations?

Gary Cziko (911203) and Martin Taylor (various) --

Gary, your discussion of Koehler's chickens has helped me understand what Martin Taylor in going on about. Tell me, Martin, if this isn't it.

I think what Martin wants is a system that can look at the barrier, its own location, the goal location, and the enclosing space, and reason its way to the best path, or at least one good-enough path, before actually moving. This, rather than feeling its way along strictly in present time, dealing with only the local environment as it goes.

I can agree that there are cases in which this would be the appropriate way to get from A to B. As Gary suggested, there might be an internal model involved, in which moves are tried and rejected in imagination, until a successful set of moves is found. Then a series of immediate goal-positions would be selected, or perhaps a series of velocity vectors, to be executed from memory of the successful solution. The first goal position or direction could be directly away from point B.

This kind of behavior requires a much more complicated control hierarchy. So I wonder what would be the circumstances that would call for this much complexity. The obvious answer is, circumstances in which a model as simple as mine could not succeed. While my model can get from A to B under more conditions than one might at first imagine, it isn't hard to complicate the environment to the point where my model would fail. A creature as simple as my model would simply perish in that environment -organisms like my model would have found a different and simpler niche.

Do we have to choose, then, between my model and a more complex one as a model of organisms? I don't think so. The concept of a hierarchy of control allows us to have both in the same organism. My model will get an organism past obstacles to a goal-position without involving any higher systems, as long as the local topology is simple enough. But even if the relationship of the present position to a distant goal is topologically complex, it is not likely to be complex if the goal is very near -- say one step away. So a higher system could survey the situation and using perceptions of relationships, categories, sequences, and programs imagine a series of goal positions that would lead from the inside to the outside of a Klein bottle. If the surface of this bottle were studded with randomly-place bumps, many of which are invisible from the starting point, the planned path would still work, because it is executed as a series of short goal-seeking moves with the avoidance systems taking care of unforseen obstacles of simple kinds. Or to go to the continuum, the goal-position could be moved along the planned path, with the lower-order systems keeping the organism's position at the specified reference position all during the trip, save for minor detours around obstacles.
So we can have our cake and eat it too, with a hierarchy.

Best to all
Bill P.

==

9112B CSGnet

1 message in INBOX

No.	Posted	From	Subject	Size
6	Jan 04 18:22	Gary Cziko	log9112b	99999

Command: print 6
Date: Sat Jan 04, 1992 6:22 pm EST
From: Gary Cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: CZIKO@vmd.cso.uiuc.edu

TO: * Hortideas Publishing / MCI ID: 497-2767
Subject: log9112b

==

Date: Sun, 8 Dec 1991 11:06:00 MST
Reply-To: "Control Systems Group Network (CSGnet)" CSG-L@UIUCVMD.BITNET

Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Language, vision

[From Bill Powers (911208.0900)]

Bruce Nevin (911206) --

We're getting closer again. I did understand that you were showing how the linguist follows the same levels as the language learner. I've said before, although the point gets lost in my nit-picking, that the study of language is a window into the way the hierarchy works because whatever we do with words we can do with other perceptions. I'm trying to turn the spotlight on the linguist in order to separate formal methods of understanding the structure of language from the methods that a nonlinguist (without the data from statistical studies always at the ready) would have to use.

>The linguist in fact follows perceptions of meanings, but (for Harrisian >methodology, anyway) always makes sure that the conclusions could have >been reached, in principle, from just two tests: the pair test for >repetition (phonemic contrast) and tests that rankings of acceptability >are not reordered in a proposed mapping (transformation).

Can you explicate what the linguist must perceive in order to do these tests? What are the reference conditions, and what perceptions might result in errors?

>In a person who has learned the language, "dog" is a primitive argument >and "eat" is a first-order operator requiring two primitive arguments >because these two words meet the demand of a sequence detector and a >program detector that controls for word dependencies between just such >an operator and its arguments.

OK, I see. "Operator" and "argument" are formalizations like "input function" and "comparator." Could you express these terms so they could be applied to, say, tennis? I don't mean to a description of tennis, but to playing it. Behind these terms, which carry a linguistic slant, are processes that I assume can also be applied nonlinguistically to perceptions of sequences and programs in which the elements are not words. If you can come up with a way to describe what I'm after that would apply equally well to forming sentences and to playing tennis, we will be that much closer to figuring out what the sequence and program levels do in general, rather than in specific applications.

You say

>It does not matter what sentences you encounter, you find others that >are like them or partially like them.

"Likeness" depends on criteria for judging similarity, doesn't it? I'm thinking here of the difference between transforming one sentence into another with a similar linguistic structure (Why do you do that vs. I ask why you do that) and by paraphrasing (I wish you'd tell me why you have that bad habit). One transformation is based on purely linguistic criteria according to a theory of structure, while the other is considered similar because it leads to a similar meaning. I think that it would be difficult to tell, in many cases, whether two sentences are considered equivalent (by nonlinguists) because of their formal structures or because they produce the same perceptual meaning. My bias, as I'm sure you can tell, is toward the latter interpretation of "similarity," although I don't deny that it may be possible to find formal rules that explain some of the cases.

You're really going to have to handle both sides of this argument, because I don't know enough about this subject. As I see it, the basic question is whether sentences obey transformation rules, or whether they obey the requirements for conveying meaning. Simply finding rules that fit selected sentences doesn't show that this is the true explanation: it shows only that such rules can be found. If you adopt the idea that meaning is the central organizer of sentences, then "Why do you do that?" is NOT equivalent to "I ask why you do that." The second version emphasizes the asking while the first concentrates on the question. If you ask it the first way, and five seconds later the second way, you indicate that you insist on an answer and that you have some sort of right to do so. So these forms are not equivalent in discourse.

I guess my point is that while there are certainly social conventions for communication that we learn, these are like learning how to hold pencils or forks; other ways also will serve, as long as they don't create consternation. It's as though we are taught "here's A

way to speak that others will understand because they were taught the same way." But it isn't THE way to speak. It's just one way that works, to get started.
If there are other ways to structure sentences that lead to the same meanings, then clearly the rules don't capture the essence of language, and we ought to be suspicious about whether the apparent rules are ad-hoc or fundamental. I don't mean that I know the answer. But I'm seriously asking.

>The theory of language we are talking about accounts well for the >socially inherited structure that is in language.

When you say "it accounts well" do you mean that it explains every utterance by every person? Or do you mean that it fits more examples than it doesn't fit? If the former, you've got a good theory. If the latter, you have only a generalization
. --
Bruce Nevin and Joe Lubin (911206) --

You might also look up the Edwin Land theory of color vision, which requires only a long-wavelength detector and a short-wavelength detector. I was privileged to be at a biophysical society meeting in which Land demonstrated the effect. He prefaced his demonstrations by remarking that vision specialists had assured him that they knew all about the "faint suggestions" of color that could be obtained by combining monochromatic images. He had two slide projectors, one projecting a black-and-white slide with white light and the other projecting the same slide with red light. He showed each image by itself; first a red image, then a white monochrome image. He said "Watch for a faint suggestion of color," and turned both projectors on. The whole audience gasped. The screen showed a full-color picture that could have been taken with high-quality color film. It was astonishing.

Later he put an orangish-red filter over the white-light projector and left the deep-red filter on the other. The picture showed somewhat dim green and red fruits in a brown bowl with a BRIGHT YELLOW banana lying on top. Land explained that all the wavelengths of light in both images lay shortward of "yellow." Yet there was that banana.

Land's theory, as I remember it, entails normalization. The visual systems hypothesize, in effect, that the average color over any whole visual field is gray. Individual objects are seen in colors that deviate from this gray. So, within wide limits, the color of illumination of a scene does not affect the relative colors observed; they are seen as if the illumination were white light. Of course it's possible to fool the visual system, as Land's demonstration showed.

Joe, this is probably grist for your model of perception. It shows that categorial boundaries aren't fixed, unless you put the normalization at a lower level.

Best to all

Bill P.
===

Date: Sun, 8 Dec 1991 13:51:44 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Re: color, categorical perception, hyperacuity

[from Gary Cziko 911208.1330]

What fascinating stuff Lubin, Taylor, Powers, and Nevin (did I forget someone?) contributed as a consequence of my query concerning hyperacuity. Thanks so much for giving me so much to ponder.

Joe Lubin (911206.11000 said:

>The phenomenon of categorical perception (CP) applies to

>color perception. While the physical spectrum is linear, its >psychophysical partitions are not. If you have a base color say >at 500 nm wavelength pure light and a test1 color at 475 nm and a >test2 color at 525 nm, in this particular case, the wavelength >delta between the base and test1 (25 nm) would cross a perceptual >boundary (nonlinearity), while the delta between the base and >test2 (also 25 nm) would not. Base and test2 would both be >instances of psychophysical green, while test1 would be blue, a >totally different category.

>CP is basically a contraction in the psychophysical space within a >category, i.e. 500 nm and 525 nm appear closer together than they >ought to be based on the linearity of the spectrum, as well as a >separation of between category stimili in psychophysical space, >i.e. 500 nm (green) and 475 nm (blue) appear much further apart >than might be expected.

This is something that has always fascinated and puzzled me. Looking at a rainbow, I perceive a continuum of change from red into orange and orange into yellow, but red and yellow appear as two distinct categories and I have no sense of the underlying frequency continuum involved in color.

Now, while the pitch of sound of is also dependent on frequency, the continuum is much more apparent. Not only can I tell that D is higher than C, but that G is higher still. And I am also sensitive to sound frequency ratios with octaves (frequency doublings) sounding like the same note although I can still say which is higher, and I can transpose a song into
any key and still recognize it (the original "Gestalt" phenomenon).
Perhaps Joe Lubin or others can help me understand why these differences between color and pitch perception exist. I do have some ideas about this, but would like to save them until Joe and others have had a say.--Gary
P.S. Joe Lubin also said:
>A good friend of many on the net, Stevan Harnad, is a CP expert.
I do not doubt that Stevan Harnad is a good friend of yours and many on CSGnet and that he is a CP expert, but as far as I can tell (and I am the "listowner"), he is not on CSGnet. Perhaps an invitation and endorsement from you might encourage him to join
===

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: Sun, 8 Dec 1991 21:21:16 CST
From: URROBERT@ECNUXA.BITNET
Subject: replies to Gary & Hugh re reject letters

[from Dick Robertson]
I noted your requests for rejection letters from journals, and will send off a couple that I got from APA to you in the mail next week. Also I would like to request Gary to please send me {the instructions for subscribing again. I have a couple of students who would like to get on the list. Thanks, Dick.
===
Date: Mon, 9 Dec 1991 09:24:43 -0600
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: Brain experts

[from Joel Judd]

Now I Understand Dept.:

From yesterday's Chicago Tribune, a short blurb copied from the Seattle Times entitled "Human brain likened to a confederacy of dunces." Quoting and summarizing books by psychologists Robert Ornstein and William Calvin (_The Evolution of Consciousness_ and _Ascent of Mind_), the reporter provides some wondrous images of how human brains evolved and why we don't function well in present circumstances.

Ornstein is quoted as saying our present capabilities are "largely an accident" and we aren't prepared to deal with modern civilization. He also suggests "that educators spend less time on math skills that emphasize logic and more on probability, because that is how our brains really make judgments. When was the last time you used algebra? Instead people should be taught statistical theory and probability."

One of the main questions for both authors is put this way: "Why are we as smart as we are when related apes stayed and survived with far smaller brains?" Citing "environmental stress" as one of the current popular brain evolution theories, Ornstein suggests the need for cooling a larger brain led to an upright posture where cooling would be immediate, and the "strain of chasing game across the widening savannahs created new heat stresses,
which led to evolution of an improved circulatory system to cool the head." [Maybe this could be called the "Prestone Theory"]

The best part, though, is the last where comments on our abilities to deal with modern life are provided. Again, Ornstein is summarized as saying "we react more to change than to _persistent problems_; we don't comprehend numbers very well in thinking about budgets; our memories are faulty; our forecasts are colored too much by our _personal past_; we are blind to infrared light, deaf to high-pitched noise; we overestimate how many people share our beliefs; we are easily manipulated, easily distracted." (underlining added) "No matter how much we prize our individuality, almost all human beings act automatically in the same way," Ornstein writes.

To his credit, the reporter notes: "Ornstein notes that admitting this bleak picture is the first step in overcoming the primitive limitations of the brain and solving modern problems. But his book is weakest in prescribing exactly how thinking can or should alter, given our cranial 'simpletons.'"

Hmm, it is kind of depressing to ponder the future of humankind when all you do is react to the environment.
===
Date: Mon, 9 Dec 1991 19:00:41 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: 2-D control; coordination

[Martin Taylor 911209 18:40]
(Tom Bourbon 911205)

I know I said I would be back on Thursday, and I did manage to get off a couple of quick responses on that day, but they took down the machine with
my user file-system on it, so I could do no more until today. Now there are some 40 messages backed up, some with 300+ lines (thanks, Bill). When I can get to them, or back to the old threads, I don't know, but I'll try!

Tom says:
> For the past couple of years, I have been working on models
>of interactions between hands, models and combinations of the two.
>One of the first ideas I explored was whether independent control >loops could produce interactions, through their control of >variables that are linked in the environment, that are coordinated. and goes on to describe some such experiments, of which I was aware because he had kindly sent reports on them. But they do not seem to cover the situation that formed the basis of the discussion, which can be described abstractly as one in which if either of the obvious low-level control systems attempts to bring its percept closer to its reference, the system as a whole cannot meet its target, but if each moves away from the reference in a particular manner, both can then get back to their target (by a detour). The question was whether this problem could be solved by a system of one-dimensional control systems. Bill showed that it could, by introducing two more low-level control systems that could engage in a conflict withthe ones having to make a detour, in such a way that the problematic ones would eventually be brought into a condition in which the approach to their reference conditions would be monotonic. That solution struckme as being the way a blind man would approach the problem, by feeling his way along the barrier, rather than the way a sighted person would, by "appreciating" (don't ask me to define that, yet) that there exists a simple route--"reculer pour mieux

sauter". When I left, Bill was about to convince me that a system of unidimensional control systems would behave like a sighted person, if appropriately designed. Tom's coordinated control of multiple *monotonic* dimensions is neat, but irrelevant tothis problem (I think).

I am interested in seeing whether there exists a naturalistic but simply described situation in which unidimensional control systems will failbut natural organisms (including humans) will succeed. By a unidimensional control system I mean one that has a scalar value as reference and as percept, and that has an error signal based on the difference between
the reference and percept, a single-valued function of which is used as an output that provides (part of) the reference signal for lower-level unidimensional control systems. (I hope that conforms to the intentions of PCT; it is the way I understand it).

The feedback loop that constitutes the rest of the control system lies in the environment. In the environment there are all sorts of intrinsic
interactions among the dimensions reflected in the percepts of the different controllers (the elements that compute the differences between references and percepts). The question is whether the overall system behaviour can accommodate these interactions without any internal provision for one controller to have access to two or more percepts.

I have a persistent feeling that the issue involves bifurcations in the behaviour of the total control loop with slow changes in the environmental part of the loop. If so, then it has an impact on the program level at
the very least, and quite probably as low as the category level (which
I think involves a catastrophic bifurcation in the input-output relation
as a function of system gain).

Martin Taylor
==

Date: Mon, 9 Dec 1991 19:31:09 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA|EJeHn>
Subject: Re: language, reality

[Martin Taylor 911209 1930]
(Bill Powers 911205.1000)

>
>Gary, your discussion of Koehler's chickens has helped me understand what >Martin Taylor in going on about. Tell me, Martin, if this isn't it.
>
>I think what Martin wants is a system that can look at the barrier, its >own location, the goal location, and the enclosing space, and reason its >way to the best path, or at least one good-enough path, before actually >moving. This, rather than feeling its way along strictly in present time, >dealing with only the local environment as it goes.
>
I am hoping that "reason its way" is not the answer, for two reasons. First, the process seems to be simpler than that, finding what amounts
to geodesics in some space of effort. Second, I anticipate a later discussion as to whether it is possible to represent logic with scalar control systems, so without a resolution of that debate-yet-to-come, an appeal to "reason" would not satisfy my question as to whether a hierarchy of unidimensional control systems would work. If you take away "reason" and use the nicely vague word "see", I think you would have what I want. The solution might involve executing your machine in fast time in imagination, but I am not sure it would.

Martin Taylor
==
Date: Mon, 9 Dec 1991 19:43:50 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Hyperacuity; 4-D control

[Martin Taylor 911209 17:30]
(Bill Powers 911206.0745)
>

>
>However, your note that a wire that subtends 1 second of arc can be seen >implies that very small stimulations of retinal cells add up over the >length of the wire enough to yield a perception, implying that there is a >functional grouping of cells (somewhere) corresponding to long distances >across the retina (in every possible direction), so that all the retinal >signals together indicate presence of the wire. It would be most >interesting to know how different in angle TWO 1-sec wires (passing >across a common point) would have to be in order for two rather than one >to be seen. Since the retinal cells subtend about a minute of arc, one >would suppose that differences in angle, in order to be perceived, would >have to carry the ends of one wire at least 1 min of arc away from the >other wire. It's hard to imagine how preferential directions of >connection could discriminate angles more finely than that.

That's a lovely experimental idea. I don't suppose anyone has done it, though it's a long time since I was in that business. I do not think
that there is any requirement on the retinal connections to be linked in "straight" arrays over a long distance for this integration to work. I would be very
sUEHI(f§_%__"!§I§_:_M_)+
!_%__§I§9
§_J9_"§Q§
Q_ %1%Qe5~between a long straight wire 1.5 sec of arc in width and a slowly curved one of the same length and width.

(Having said that, the eye does seem to be very good at seeing geometric straightness. When I was a summer student, trying to attach straight bits of platinum wire to translucent plastic disks without leaving any orientable traces, my supervisor was studying the ability of the eye
to discriminate arcs of large radius of curvature (up centre from down,
for example). He gave up because he couldn't fabricate arcs of sufficiently large radius to provide a problem to the eye.

As a practical matter, 1 minute of arc is a bit larger than the best a good eye can accomplish for resolving two objects. That ability is more or
less what is called 20/20 vision. The best eyes can do almost twice as well, about 30-40 seconds of arc. It is eyes of that quality that can see the wires down to nearly 1 second of arc.

Martin Taylor
===
Date: Mon, 9 Dec 1991 20:02:31 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" CSG-L@UIUCVMD.BITNET
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: questions on color perception

[Martin Taylor 911209 19:50]
(Bruce Nevin 911206 09:12)
>
>This may bear on Whorfian experiments that say the distance between >colors predicts that some will be more basic and earlier evolved in >languages and cultures, some later evolved.
>
("This" being the interpolation of colours using only three receptor types).

No, I think not, because the actual discriminative sensitivity curves
seem to be quite culturally independent, peaking in the yellow, but not terribly dependent (order-of-magnitude) on spectral wavelength. The spectral discriminative sensitivity curve is probably due to the snesitivities and distributions of the receptor types.

The Whorfian aspect is that the language determines the boundaries among categories, and that colours within a category look more similar than
colours in different categories. You yourself forwarded an expansion on this by one of the original authors of the Mexican study, saying how strong and how surprising this effect is. The number-of-categories thing is different, I think, in that cultures with more categories more or less subdivide the categories of cultures that have fewer. That

"dark", for instance, is subdivided into (say) blue and red (I forget where the actual subdivisions are) if there is another splitting of
the categories does not say where the boundary will lie between blue
and red in terms of mixing ratio. That's where the Whorfian question arises, as I see it.

Martin Taylor
==

Date: Mon, 9 Dec 1991 20:13:43 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: color, categorical perception, hyperacuity

[Martin Taylor 911209 20:00]
(Joe Lubin 911206.1100)
>
>I think the theories of language-creates-psychophysical-boundaries >(at least for color perception) are no longer in vogue in CP >theory. Demonstrations of CP in human infants and ferrets caused >a reformulation.
>
I don't think the argument ever was that language *creates* psychophysical boundaries, in the sense that without language there are no psychophysical boundaries. We've known for a long time that little bunny rabbits respond in a categorical way to Voice Onset Time (VOT). If anyone made that claim I think they misrepresented Whorf and the mainstream of Whorfians. The claim is more that there is a kind of resonance between language and perception, each reinforcing the other. All the same, there is also a claim that language *can* create distinctions where alinguistic perception would see none (in some contexts, that's called reification). In colour perception, in particular, it seems clear that the perceptual experience is strongly affected by the language of the perceiver as well as by the physiology.

Colour was chosen for these experiments because it seemed the least likely place to find an effect, being so close to the physiology (apparently). Finding Whorfian effects there was quite a surprise to the investigators (as reported first-hand in this group a few weeks ago). Given that language can affect perception in an area of perception so closely constrained by physics and physiology, I think it would be most surprising if there were no such effects in areas further from the sensory periphery.

Martin Taylor
==
Date: Mon, 9 Dec 1991 20:25:49 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: linguistic and nonlinguistic intonations

[Martin Taylor 911209]
(Bruce Nevin 911206 1207)
>
>I'll wait until you revisit what puzzled you in my post of a week ago >Tuesday (Bruce Nevin 911202 1111) and integrate it with intervening >discussion Bill and I have had. I will include here some clarification >about expressive intonations not being in language, though we use them >for communication just as we use language for other aspects of >communication. In turn, I would like you to clarify what information is >conveyed by white noise (other than turning it on or off as a binary
>signal).
>
After a couple of hours of going through the backed-up mail, I was coming to the end of my endurance when I came to this posting. I won't try to remember the older problem yet (maybe after dinner), but I can give a quick answer to "In turn..."

White noise in a band limited channel can be completely specified by a series of samples from a fixed Gaussian distribution whose RMS value
is given by the power capacity of the channel. The sampling rate is
twice the bandwidth (any more frequent sampling provides no new information, but helps with the filtering if the noise is to be reproduced from the samples). Each sample specifies an amplitude selected from the Gaussian distribution, and it is specified with some precision. The choice from

the distribution provides an amount of information given by p log p, where p is the probability that a point from the original distribution is found in the distribution implied by the precision with which the sample is measured (that's imprecise language, but it is easily made precise in mathematical notation). The average amount of information provided per sample is greatest if the samples are independently selected from the Gaussian distribution (Shannon showed this). No signal can convey more information for a given power than can a white noise of that power.

Now, when I remember what brought up that question, I'll get back to the issues that were at hand.

Martin Taylor
==
Date: Mon, 9 Dec 1991 20:36:08 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: thanks, and a plaint

[Martin Taylor 911209 2030]
Bruce Nevin (911206 1259)

(A final gasp for this evening--)
>
>I was just thinking about the redundancy in our categorial perception of >the elements in language. When you spend hours transcribing tape >recordings, as I have done, you become aware of the variety of cues that >might suffice for recognizing a particular word.

I concur. It is extremely difficult to perceive just what people actually say, in transcribing tapes of natural dialogue. One regularizes to an extraordinary extent, and it is only by comparing one iteration of the transcription with the actual tape that you realize that the transcription missed an "uh" or a syllable repetition, or that the talker really didn't have an article before that noun... We do hear the intended speech, as our decoders interpret it using the redundancies of language. Bruce talks about all the cues to words, and I think the same applies to the higher level constructs of language.

Naturally, one gets better at not hearing what the speaker intended as one gets more practiced, but it's still difficult.

(Aside: I'll bet one could do it easily under hypnosis. Has anyone tried?)

Martin Taylor
==

Date: Tue, 10 Dec 1991 08:34:18 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: no information without redundancy

[From: Bruce Nevin 911210 0703)

(Martin Taylor 911209) --

>White noise in a band limited channel can be completely specified by a >series of samples from a fixed Gaussian distribution . . .

> Each sample specifies an amplitude selected from the Gaussian >distribution, and it is specified with some precision. The choice from >the distribution provides an amount of information given by p log p . . .

> The average amount of information provided per >sample is greatest if the samples are independently selected from the >Gaussian distribution (Shannon showed this). No signal can convey more >information for a given power than can a white noise of that power.

This describes a measure of quantity of information, not a characterization of information. A signal can convey information.
White noise cannot.

What makes a signal a signal rather than a noise is some specification of differences that make a difference. In a code, that specification is in a list or table external to the code itself that maps discriminable differences onto "meanings." The "meanings" are typically specified in natural language (<noise1>="dog," <noise2>="cat" . . .). It is the mapping itself that specifies the elements of the code.

In natural language, the specification of differences that make a difference (after learning to control phonemic contrast) is done by redundancies in the language.

In a code, information about what are elements and about the correspondence of those elements to "meanings" is given in the meta-code table of mappings. Redundancies in the table correlate with that information *content*--the information about what are elements, etc, just mentioned. In a code, information about what elements may be combined is given covertly in the background vernacular of natural language, on which the code-mapping relies for its "meanings."

In natural language, information about what are elements and about their combinability is given in redundancies in the language. These redundancies are recognizable in terms of the complementary and contrastive distribution of the elements relative to one another, specifiable in terms of a few classes of words and reductions of the phoneme sequences constituting the words. The classification is in terms of dependency on the dependency-class of previously said words (or perhaps "previously intended words," allowing for reduction).

Of course, the lowest-level elements, the phonemic contrasts, are given in something like a table-lookup form, but the only information or "meaning" here is the fact of contrast. Redundancies in phoneme sequences specify morpheme boundaries (reduction in next-successor count within a morpheme, return to a higher inventory at morpheme boundary). Redundancies in morpheme sequences specify morpheme classes. All of this is learnable as a socially inherited system or structure in principle without reference to meanings (nonverbal perceptions).
Perhaps this is the achievement of idiots who can produce words and nonsensical utterances without making the correspondence to nonverbal perceptions.

This summary may make explicit the covert reliance of codes on the structure in natural language:

```
      Elements        Combinability          Correspondence
                                of Elements      to Meanings
-----------+-----------------+---------------------+------------------+ Code         |
In external    | Specified by    | In external|
                    |  meta-code       | natural language    | meta-code  |
                    |  (lookup table) | "meanings" in table  | (lookup table)   |
                    |                 |        |      |
-----------+-----------------+---------------------+------------------+ Language      |
Specified by   | Specified by word      | Correspondence    |
                    |  redundancies   | classes (based on   | of constructions |
                    |  of lower-      | redundancies) plus  | to nonverbal      |
                    |  level elements | reductions plus -->  | perceptions      |
-----------+-----------------+---------------------+------------------+
```

Word classifications and many of the reductions of phonemic shape of words are learnable just from redundancies in language. However, the membership of all words in operator classes and the detailed working of the reduction system is learnable only with reference to meaning--the correlation of operator-argument dependency structures with nonverbal perceptions, perhaps a correlation with dependencies or expectancies (themselves perceptions, of course) held with respect to other nonverbal perceptions.

A code cannot be learned only by possessing the meta-code table. One must also know the language in which the "meanings" are expressed. Possessing the table, but not knowing the language, one might know <noise1>=<squiggle17> and so on, but one could not say which combinations of code elements constituted a meaningful message. Since all combinations are possible, no particular comibination bears any particular information.

Elements of a code are mapped onto "meanings" which are elements of a language. Constructions in language are mapped onto "meanings" which are nonverbal perceptions. The seeming parallel, which underlies the code analogy to language, is I believe misleading. Relations among nonverbal perceptions are not constrained and structured to the extent and in the conventional ways that relations among the elements of language are constrained and structured. If they were, then the mappings from language to nonverbal perceptions for different languages and indeed for different speakers of the same language would be much closer to unanimity than they are. And we would not have to work so hard to discover and refine scientific theories--the language we speak would incorporate a fully adequate account.

Bruce
bn@bbn.com

===

Date: Tue, 10 Dec 1991 08:20:56 CST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: URROBERT@ECNUXA.BITNET
Subject: Returned mail: grade control project to Bill Powers

Forwarded message:
From daemon Sun Dec 8 22:00:29 1991
Date: Sun, 8 Dec 1991 22:00:07 -0600
From: MAILER-DAEMON (Mail Delivery Subsystem)
Subject: Returned mail: Internal error
Message-Id: <199112090400.AA22299@uxa.ecn.bgu.edu>

To: urrobert

----- Transcript of session follows -----

btmail: Invalid address R@FLC@VAXF.bitnet

btmail: Letter saved in dead.letter
554 R@FLC@VAXF.BITNET... Internal error

----- Unsent message follows ----Received: by uxa.ecn.bgu.edu (5.65c/IDA-1.4.4) id AA22220; Sun, 8 Dec 1991 22:00:07 -0600
From: urrobert (Richard Robertson)
Message-Id: 199112090400.AA22220@uxa.ecn.bgu.edu
Subject: GRADE CONTROL PROJECT
To: PETERS, R@FLC@VAXF.BITNET
Date: Sun, 8 Dec 91 22:00:06 CST X-Mailer: ELM [version 2.3 PL11]
[From Dick Robertson] (9112.08)
TO: Bill Powers
Thanks for your call the other night. Here is my sketch of the revision of the grade control analysis that you worked out for me at the Durango meeting.

```
                                  R                Where  S equals the next computed
                           ____|____              test score derived from:
                           S-> |_______| -E    (1)    ko  = computed preparedness
                           |___       __|__                       estimate*@
                    |k(1)|      |k(o)|  (2)    E   = original RS-PS
                                  |   ______     | q(o)
                                  |<- |test|<-- |      (3)    qo  =k1(R-S)  {?qo=ko?}*
                                  D    k2=x5    (4)    S   =k2 qo  {?k2ko?}*
                                                       (5)    S   =k2(k1(R-S))=
                                                                   k1k2R-k1k2S   so,

                                                       (6)    S(n)= k(o) k(2) R(next)/
                                                                          1 + k(o) k(2)
```

It all seemed to make sense to me then, but when I got around to revising

the program for recomputing it, I ran into some snags.
*@ Recall that when we looked at the preliminary results that I had before you moved, you suggested computing the regression of the next-prepar.est. on the previous error.
I read that as substituting the est.PEnxt for the PE value given by the student, at k(o).
* You had k(o) and q(o) on your sketch, as I indicate here, but am I right in thinking they are alternative notations for the same thing and we only need k(o) as in your final equation?
Next, k2 is in there simply because I got the prep.est. in the form of question numbers, & k2 multiplies them by 5 to put them in percentages to be comparable with the other variables.

Your sketch didn't define or identify k(1). I put it into the INPUT box, but I don't see what value to assign it. I took S to be a computed test score, to be derived from your final equation, that could then be correlated with the student's actual test scores to see how well the equation simiuIates the performance. Is this right?

Recall that in the earlier version OP was the mean of sumK(o) where K(o) was Prepar.est./Error. We had concluded that the students in general didn't seem to be basing RSnxt on K(o). But there was that finding that MnK(o) did separate the students who took incompletes from those who finished the course. So I am keeping that measure in the program. But you thought that the regression of PE (prepar.est.) on the previous ER might better reflect how much the Error was affecting the next preparation. Then, if we plugged in the estimated PE (prepar.est.) in the K(o) column I should use equation (6), above, to compute S and correlate that with actual PS.

I had also made some other modifications. I encouraged students to change RS as seldom as possible this time, to think of it as thier goal for the course grade. And I redefined PS to be the Mean of all tests taken at each quiz point. I asked them to compute that each time they got a test graded so that they would be looking at the relation between their Ref Sc and their current average--as more analogous to a tracking task. Hope this is clear.

9112 CDE CSGnet

Date: Fri Jan 03, 1992 5:22 pm EST
From: Gary Cziko
EMS: INTERNET / MCI ID: 376-5414
MBX: CZIKO@vmd.cso.uiuc.edu

TO: * Hortideas Publishing / MCI ID: 497-2767
Subject: log9112c, d, e

==

Date: Sun, 15 Dec 1991 13:41:00 CST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: TJ0WAH1@NIU.BITNET
Subject: acuity; epistemology

[From Wayne Hershberger]

(Bill, Martin, Gary, Joe, et al.)
Regarding acuity:
There are a variety of operational definitions. Resolution acuity (detecting the orientation of lines comprising a grating--vertical vs. horizontal) yields a nominal normal of about 30 arc seconds (as Martin has implied), approximately the diameter of foveal cones. Detection acuity, detecting the presence of a black dot on a white field, yields a nominal normal of .5 arc seconds.
Although the latter is called acuity, or hyperacuity, it does not appear to involve a registration of size. Rather. it reflects Weber's ratio for the detection of a difference in luminance (about 1 to 2 %). That is, a dot subtending .5 arc seconds generates an extremely poor "image" on the retina amounting to reduced retinal

illumination (by 1-2%) across a relatively large area of the retina. I have read that WW II fighter pilots detected enemy aircraft by looking for such smudges in the sky.

Bill Powers (911203)
>Yes, in my model there is always an environment and a behaving >system. Neither makes sense without the other. I have always >taken both into account. So follow me as I outline a chain of >reasoning, and see if there is any point where you detect a weak >link.

Gee, Bill, that is like waving a red flag in front of a bull. How can I refuse. Here goes.

Although you say, above, that the environment is in your model, you soon speak as if it were external to it:

>But now we come to the crux of the problem. We want to let the >model figure out what there is externally to it that corresponds >to its perceptual signals. For example, the object it is looking >at is actually a hologram, and all that actually exists in the >environment is a set of wavefronts of light that don't actually >originate at the surface of an object.

The basic (perhaps only) problem I see is your confusing separate issues. In your elaborate analogy there is one dichotomy (creator-created) and one dyad (organism-environment). You confuse the two, shifting from one to the other as though they were one and the same. As I have said before, this confuses physiology with metaphysics.

For example, you say that:
>we can't rise in a fourth dimension out of our brains, to peer at >whatever it is that is causing our neural signals.

The first part of the remark, mentioning a fourth dimension, is alluding to the inability of the created dyad to assume the epistemic perspective of the creator (epistemology), whereas the last part is concerned with the relationship between the two parts of the created dyad (sensory physiology). Apples and Oranges.

Then you reverse course and switch from sensory physiology back to epistemology with the following remark, effectively by substituting the word "perceptions" for the words "neural signals" (moral: perceptions are not to be equated with perceptual signals).

>the transformations that lie between the environment and our >perceptions.

This last phrase is also part of a sentence that prompts a question I think might be constructive.

>As the model can't sense the internal workings of its perceptual >functions, and use that information to deduce what is causing any >given perception, so we can't deduce the transformations that lie >between the environment and our perceptions.

If a neural model could monitor its role in the perceptual process, could it deduce the nature of the transformations that lie between the neural model's signals and the neural model's environment?

(Bill Powers (911205)
>"Immanent order" wouldn't be a bad term for "boss reality." From >my viewpoint it has the nice implication that there can be order >without our knowing what it is.

Yes, exactly. It seems to me that the expression "immanent order" (or natural order, or what have you) would be a much better term for your purposes than "boss reality," for precisely the reason you mention. The word reality connotes a verifiability you are denying to "boss reality," making the expression an oxymoron.

>By choosing the diagram in which the immanent order extends >beyond the boundaries of the realizations, you have agreed with >me that there is more to know than meets the eye.

No. I think there is a point of agreement here, but not for the reasons you say. The aspect of the diagram that implies that there are some things which can be known but which do not directly meet the eye or the ear or the other sense organs (e.g., your example of voltage) are the elements which are both IN the set labeled conceptual realizations and NOT IN the set labeled perceptual realizations. In contrast, the elements that are in neither subset (neither type of realization) simply imply an immanent order which is not realized--either perceptually or conceptually. Whether this unrealized order is potentially realizable is something a static venn diagram doesn't capture.
But if one takes the view that at least some of the immanent order unrealized at present may be realized in the future, it is presumptuous to suppose that this realization CAN NOT be perceptual. Further, any immanent order which can not possibly be realized at any time in either way is simply not to be known; it does not mean that there is more to know than can be known. I readily admit that there can be more than what-can-be-known, but I can not agree that there is more to be known than what can be known, without contradicting myself. Nor can you. We are talking here about the limits of the epistemic process not the limits of a man--obviously, there is more to be known than any one man will ever know.

>By making the two realizations independent and non-overlapping, >you have said that each has its own relationship to the immanent >order independently of the other.

I would say that one is not a subset of the other, but their intersection is not nil, meaning that the two realizations are independent of each other. Your drawings did not seem to include this alternative, so I selected the one which I thought would "suggest" independence (actually non-overlapping subsets depict mutual exclusion, a form of dependence).

>As you didn't specify the difference inside and outside the >realizations, two possibilities are:
>1. A realization is simply a noticing of something that was >always there, the noticing in no way altering what was always >there but merely bringing it into the field of attention.
>2. A realization is some transformation or projection of the >immanent order, so that the realization is an invention or at >least an expression of the nature of the system becoming >acquainted with the immanent order.

Your two alternatives are not a matched set. They are not necessarily mutually exclusive.

First the latter:
The expression "system becoming acquainted with the immanent order" seems to me to suggest that the "system" transcends (stands apart from) the immanent order. That would insinuate a gratuitous wild card. For me, the system becoming acquainted with the immanent order must be part and parcel of the immanent order. Perhaps they are even coextensive. The system responsible for the two types of realizations is best characterized as an ecological system (i.e., an organism-environment dipole). If we attribute the "becoming acquainted" merely with the organism pole and the "immanent order merely with the environment pole we are being arbitrarily inconsistent. Therefore, if one is to be consistent, it seems to me that the realizations would inevitably be "an expression of the nature of the system becoming acquainted with the immanent order" because the system (the ecological dipole) becoming acquainted with the immanent order IS part and parcel of the immanent order.

Now the former:
However, does "self-acquaintance" rule out the possibility that acquaintance is simply a registration of what is "there"?

Fortunately, the question appears to be academic. If some aspects of the immanent order are hidden by the recursiveness of self-acquaintance, or whatever, I would say, so what? Call it Noumenon, and let the faithful worry about it, because, by definition, it is not to be known.

>I take it that the rationale for the term "immanent order" is >that neither perception nor conception is random; that both >reflect some orderliness that constrains them. Does this not >imply some effect of the immanent order on the realizations?

The relationship is not cause->effect, but, YES, realizations of both types reflect some orderliness that constrains them. That's what I think--I think.

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, 15 Dec 1991 15:10:35 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: goldstein@SATURN.GLASSBORO.EDU
Subject: PCT and suicide

To: Ed Ford, interested CSGnet people
From: David Goldstein
Subject: suicide and PCT
Date: 12/15/91

Let me see if I can elaborate on the difference between "random" and "learned" suicide which you said confused you. But before I do, let me ask you a question: Let us talk about a person who has verbalized the intent to suicide. Here we have person A who has a history of four suicide attempts. Over there is person B who has never attempted suicide before. As I said, both persons A and B have verbalized the intent to suicide to you. Let us imagine that you had one hospital bed available in the mental health unit of a local hospital. If if were your decision, who would you select to enter the hospital? (My decision would be person A.)

In your last post, you stated: I think the possibility of a successful suicide depends on the means to the end. Don't you think that person A, who has attempted suicide four times, has acquired better "means to the end" meaning that person A has more know how when it comes to suicide. Furthermore, in your last post you said: The more chronic or long term the error and the more a person has tried to rebuild his/her life, the more inclined a person would be to give up. Doesn't this sound more like person A than B?

Everybody knows what suicide is by the time they become adults. The "random" suicide person, person B above, is temporarily entertaining making this idea a goal. Since person B has never done it before, I assume that there will be some conflict about this. One part of B will want to do this, another part of B will not. I think that person A, the "learned" suicide attempter, has resolved this conflict in favor of suicide.

Let me summarize what I think we agree on so far:

(1) A person who is intent on suicide is not experiencing much success in life and has come to the conclusion that the future will be much the same. As a result, there are intense negative feelings/moods. Suicide seems to be a means of escaping the bad feelings/moods.

(2) The therapist instructs the person to calm down and gives him/her a task which is distracting and which is controllable. The therapist can try to arrange the person's environment to minimize the the experience of stress. This may involve some time off from work and other usual routines.

(3) When the person is calmed down, the therapist recognizes the person's situation but asks the person to commit some time to work things out. A very specific plan of action is drawn up which is monitored very closely.

(4) If the person makes the committment to live a little longer, the therapist starts to work on the problems of living which resulted in the crisis, my and your preference being along PCT lines.

(5) The person is hospitalized if steps 2 or 3 cannot be accomplished. It is in the decision to hospitalize that I think consideration should be given to a person's history of prior attempts.

(6) The person's social support network needs to be alerted to the situation and asked to cooperate.
===

Date: Sun, 15 Dec 1991 16:42:36 -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: Martin Taylor stuff

[From Bill Powers (911215.1700)]

Martin Taylor (911213) --

>Conscious perception seems to have one function (among probable others);
>it is there to remove the ambiguity of perception. (Now I go into >personal hypothesis mode)--This is because at a sufficiently high level of the control hierarchy the possible ambiguities are usually related to >incompatible actions/intentions and it is of negative value (in an >evolutionary sense) to be dealing with the irrelevant meanings.

If you'll buy into my hypothesis that consciousness directs reorganization to specific areas of the hierarchy, we will be in agreement. I would resist the idea that consciousness can play a systematic role in removal of ambiguity. If it can, then there is some function within the concept of consciousness that ought to be removed and put into the hierarchy itself -- consciousness would have to be endowed with knowledge of the particular hierarchy that has developed in an individual, and we would then have to explain where it got that knowledge.

>In the context of this group, it might be relevant to consider the >possibility that one function o consciousness might be to affect the >gain of control systems associated with relevant and irrelevant aspects >of percepts.

I would rather leave control of gain as a mode of hierarchical control to be added to the main mode now in the model. If you allow consciousness (awareness) itself to be structured, you're just building up a backlog of phenomena that require modeling, phenomena of consciousness.

> The gains of those associated with the aspects that relate to the >higher-level intentions might be raised, those dealing with irrelevant >aspects depressed.

You see? Now you have to explain how consciousness detects which subsystems relate to higher-level intentions, and which are irrelevant. This information has to be sensed somehow, and interpreted properly for the levels involved -- and suddenly you have a lot of machinery inside consciousness.

>The control of gain has been mentioned a few times in the year or so I >have participated here, but its effects have not really been discussed. >Could it be an important factor?

Yes, it could -- but I'm waiting for a specific problem to come up that can't be handled with the hierarchy as it stands. So far I have only one possible example. In the arm model, the effect of a sudden upward movement of the upper arm is to straighten the elbow angle if that angle is more than 90 degrees, but to collapse it if the angle is less than 90 degrees. The ideal way to handle this would be to have a cross-connection between the two control systems whose gain depends on the elbow angle. Also, the damping needed varies as the arm goes from flexed to extended, greatly altering its moment of inertia. It would be nice to have control of gain in the rate-of-change control system (phasic stretch reflex). But I'm putting off introducing those ad-hoc changes until I'm sure that there isn't some clever method, or some physiological fact, that will make it unnecessary. One physiological fact turned up just yesterday as I browsed through a kinesiology text I picked up at the Humane Society Bookstore for 25 cents. One segment of the biceps and also of the triceps anchors on the forearm and on the scapula, crossing both the shoulder joint and the elbow joint. So there is a very helpful mechanical interaction between the two control systems. And in the text there was a brief sentence saying that the effect of the biceps on the humerus segment is present only with the forearm extended. This could take care of the position-dependent effect I wanted.

Apropos of nothing, the C version of the arm model is finished and I'm now working on details of the dynamic stabilization. It should be ready for distribution before very much longer (January, probably).

RE: information

You say to Bruce

>I've followed a consistent definition of information that has formed the >core of my research work for some 36 years, and it is a bit >disconcerting to find that my understanding has now been defined out of >existence.

I think I side with you here: an outside observer, contrary to Bruce, can't know what interpretation is being applied to incoming messages, and so can't determine the information content just by looking at the message. If the message says "yes", the outside observer can't know how much semantic information that word carries without knowing the question that the listener is trying to get answered.

When I equated information to SNR I was simply citing the simplest form of information most closely related to the Shannon-Weaver idea. All my talk about channel capacity was meant only to show that information as an engineering concept has nothing to do with semantics. It seems to me that you are extending the concept into semantics, so that the "reduction of uncertainty" now encompasses uncertainty as to the meaning of a message, not simply uncertainty as to whether the message received is physically the same as the one that was sent. I have no objection to this extension.

I agree with you that white noise can be a signal. An example I thought up long ago was that of a field engineer trying to determine the ambient electromagnetic noise level. A broad-band reading of this noise level will be interfered with by a radio station broadcasting a symphony. So how do you characterize the signal-to-noise ratio in the noise-meter reading? Obviously, the patterned and organized music would go into the "noise" term, and the white noise into the "signal" term.

Perhaps your concept of information would be better correlated with the SSR instead of the SNR -- the signal-to-signal ratio, where one is the wanted signal and the other is unwanted.

>Both Bruce and Bill argue that the channel capacities are the same at >all levels of the hierarchy,

See later post. That was Bruce's interpretation. Actually it's pretty hard to define a "channel" in a massively parallel system. A high-level perception can be computed from lower-level signals arising at many different levels and through computations of highly different complexity. There really isn't "a" channel involved.

>Bill, do you think that ALL possible sequences of events form signals >that are appropriate and DIFFERENTLY usable by the sequence detector; >that EVERY set of intensities leads eventually to a different category >percept?

I don't know what you mean by "the" sequence detector. In my pandemoniumstyle model there is a separate sequence detector for every sequence that matters. All sequence detectors operate at the same time, independently, even when they are logically redundant. Same for category detectors: each detector is waiting for its inputs to indicate a member of the specific category it is specialized to perceive. Different category, different physical input function.

To answer the first question: when multiple events are being perceived, their signals are available in parallel to all sequence-detectors and to all category-detectors. The sequence-detectors will each respond only to a subset of the event signals, but the subsets are allowed to overlap so the same signal can be an element in more than one sequence. All sequences that we have learned to perceive and that are exemplified in the set of event signals are perceived at the same time. It's up to higher level systems to select among the resulting set of sequencesignals. More than one may be used at the same time.

We recognize only those sequences for which we have formed specific sequence-detectors: so the answer to the first question is NO. There are vastly more potentially-recognizeable sequences in a set of events than we or all of us together will ever learn to perceive.

The second question: it's the nature of a category perceiver to treat some set of input signals as equivalent: any input from that set will generate the category-signal. So large sets of intensity-signals will, rather than giving rise to different category signals, all give rise to the same category signal. Other sets of intensity signals (different arrangement of intensities) will give rise to a different category signal. And again, we recognize (in parallel) only those categories that we have learned to perceive, which is far, far, fewer than the number possible.

So the answer to the second question is also NO.

>The amount (and perhaps type) of redundancy occurring
>between levels of the hierarchy is, I suspect, critical for the >stability of the whole system, and of its ability to reorganize.

I don't disagree; I just don't understand. How can there be "redundancy" between levels that handle different types of variables? I have always understood redundancy to mean delivering the same information by two or more means, any one of which would have been sufficient.

This probably isn't a big problem, but it would be good to get our terminology coordinated.

RE: challenges to the model (a good thing):

>I sympathize, but there is a cycle here. I'm not trying to go out of my >way to find cases in which the model fails. I'm trying to find places >where one should think it ought to be applicable because the situation >is common and the behaviour (on the surface) obvious, and yet where an >easy application of the model does not seem to work properly.

OK, think of an example and we'll see what can be done. I'm not trying to keep you from throwing challenges at the CT model. The more the better. But I'm not going to play by rules that required me to model hypothetical environments. I have a hard enough time with the world as it is.

>My intuition still says that scalar control should not be expected to >work in an environment that has bifurcations,

You may well be right -- if such environments actually exist. If they don't, then living control systems don't have to work in them.

Best

Bill P. ==

Date: Sun, 15 Dec 1991 23:33:51 -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: Epistemology

[From Bill Powers (911215.1800)]

Wayne Hershberger (911215) --

We keep going around and around on the same points without getting anywhere. You keep saying that I am missing the distinction between modeling and metaphysics, and I keep saying that metaphysics is just one of the things a brain can do. Let's take it from the top.

RE: immanent order

>It seems to me that the expression "immanent order" (or natural order, >or what have you) would be a much better term for your purposes than >"boss reality," for precisely the reason you mention. The word reality >connotes a verifiability you are denying to "boss reality," making the >expression an oxymoron.

So in your book, "reality" is identical with "verifiable reality." It's not, in mine. I don't need to understand electricity to comprehend that touching certain objects is highly unpleasant. I can generate acts like touching objects, but I can't decide what their consequences will be. That is decided for me by something I don't sense and only partially conceptualize. I can choose whether to repeat a consequence or to avoid it, but I can't make an act have a different consequence. In that department, something else is boss.

Are you saying that I must realize in perception or conception the connection between an act and its consequence, and VERIFY the nature of that connection, before I can accept that there really is a connection? Or are you saying that it is sufficient to verify only that the consequence reliably follows the act, and never mind why? I would argue against the latter as being simply pre-Galilean empiricism, and reject it because it works so poorly in comparison to the method of modeling. The method of modeling posits an unseen reality mediating between act and consequence, and has most profitably interpreted nature in those terms. The assumption has repeatedly been vindicated. How could the purely empirical approach ever predict a new perception, and experimentally reveal the link explaining the surface appearance of a causal sequence?

Later in your post, you say

>... if one is to be consistent, it seems to me that the realizations >would inevitably be "an expression of the nature of the system becoming >acquainted with the immanent order" because the system (the ecological >dipole) becoming acquainted with the immanent order IS part and parcel >of the immanent order.
This would be consistent. It would also be an empty generalization, a true statement of which one can legitimately ask, "so what?" To say that all of knowledge is an expression of the immanent order (whatever that is) is meaningless: any statement that is true of everything is trivial. Even that statement and my response to it are part of the immanent order. I repeat: so what? Knowing that does not contribute to our understanding of any specific phenomenon -- in fact, it seems to discourage asking questions and conjecturing. All of our useful understanding comes from discriminating one part of the immanent order from other parts, and from realizing that different parts of it have characteristics of their own unlike the characteristics of other parts. It is out of these differentiations that all knowledge comes. From these differentiations, we come to realize that organisms and environments are NOT alike. We realize that some parts of organisms function differently from other parts. We realize that brains exist.
And ultimately we are faced with a paradox, the one you and I have been arguing about. We find by experimentation that the presence of certain signals in a brain is the sine qua non of perception. Remove those signals and you destroy, as far as the victim is concerned, a chunk of the immanent order. Yet you don't destroy it for anyone else. What other conclusion can we reach but that perception is absolutely contingent on those signals? That puts us, as perceiving entities, inside the brain. To deny that would be to

destroy the whole structure of perceptual and conceptual organization we have so painfully built up. That structure is at least as well worked out as any metaphysical argument in words, and a whole lot better tested experimentally.
I don't see that any philosophical conception, any combination of words, any exercise of pure reason, can be more persuasive than these simple observations. By simple and straightforward reasoning based on close attention to experiment and observation, we are led to conclude that the object of perception and thought is a world existing inside, not outside, a brain. We can see how this world of experience is related to what we conjecture to exist in a physical environment outside of us, but we can also see that the relationship is not a simple or direct one, nor is it wholly verifiable because of our peculiar circumstance of being inside the very system we model and by necessity having to perceive and think using its equipment.
Until you can come up with an equally persuasive set of observations and deductions that lead to a different conclusion, I will continue to be satisfied with my view of the relationship between consciousness and reality. Simply reiterating your point of view without revealing and justifying each step of the way that leads to it will not win me over. I understand that if I believed as you do, all would be explained. But I do not.
Yours unpersuadedly,
Bill P. ===

Date: Mon, 16 Dec 1991 17:40:06 +0100
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Francis Heylighen <fheyligh@VNET3.VUB.AC.BE>
Subject: Principia Cybernetica Symposium-CFP

CALL FOR PAPERS
****
* SYMPOSIUM: THE PRINCIPIA CYBERNETICA PROJECT * * computer-supported cooperative development *
* of an evolutionary-systemic philosophy *
****

as part of the

13th International Congress on Cybernetics
NAMUR (Belgium), August 24-28, 1992

Symposium Theme

The Principia Cybernetica Project (PCP) is a collaborative attempt to develop a complete and consistent cybernetic philosophy. Such a philosophical system should arise from a transdisciplinary unification and foundation of the domain of Systems Theory and Cybernetics. Similar to the metamathematical character of Whitehead and Russell's "Principia Mathematica", PCP is meta-cybernetical in that we intend to use cybernetic tools and methods to analyze and develop cybernetic theory.

These include the computer-based tools of hypertext, electronic mail, and knowledge structuring software. They are meant to support the process of collaborative theory-building by a variety of contributors, with different backgrounds and living in different parts of the world.

As its name implies, PCP will focus on the clarification of fundamental concepts and principles of the cybernetics and systems domain. Concepts include: Complexity, Information, System, Freedom, Control, Self-organization, Emergence, etc. Principles include the Laws of Requisite Variety, of Requisite Hierarchy, and of Regulatory Models.

The PCP philosophical system is seen as a clearly thought out and well-formulated, global "world view", integrating the different domains of knowledge and experience. It should provide an answer to the basic questions: "Who am I? Where do I come from? Where am I going to?". The PCP philosophy is systemic and evolutionary, based on the spontaneous emergence of higher levels of organization or control (metasystem transitions) through blind variation and natural selection. It includes:

a) a metaphysics, based on processes or actions as ontological primitives,

b) an epistemology, which understands knowledge as constructed by the subject, but undergoing selection by the environment;

c) an ethics, with survival and the continuance of the process of evolution as supreme values.

PCP is to be developed as a dynamic, multi-dimensional conceptual network. The basic architecture consists of nodes, containing expositions and definitions of concepts, connected by links, representing the associations that exist between the concepts. Both nodes and links can belong to different types, expressing different semantic and practical categories.

Philosophy and implementation of PCP are united by their common framework based on cybernetical and evolutionary principles: the computer-support system is intended to amplify the spontaneous development of knowledge which forms the main theme of the philosophy.

About the Symposium

After the succesful organization of a symposium on "Cybernetics and Human Values" at the 8th World Congress of Systems and Cybernetics (New York, June 1990), and of the "1st Workshop of the Principia Cybernetica Project" (Brussels, July 1991), the third official activity of the Principia Cybernetica Project will be a Symposium held at the 13th Int. Congress on Cybernetics.

The informal symposium will allow researchers potentially interested in contributing the Project to meet. The emphasis will be on discussion, rather than on formal presentation. Contributors are encouraged to read some of the available texts on the PCP in order to get acquainted with the main issues (Newsletter available on request from the Symposium Chairman).

Papers can be submitted on one or several of the following topics:

The Principia Cybernetica Project
Cybernetic Concepts and Principles
Evolutionary Philosophy
Knowledge Development
Computer-Support Systems for Collaborative Theory Building

About the Congress

The International Congresses on Cybernetics are organized triannually (since 1956) by the Intern. Association of Cybernetics (IAC), whose founding members include W.R. Ashby, S. Beer and G. Pask. The 13th Congress takes place in the "Institut d'Informatique, Facultes Universitaires Notre-Dame de la Paix, 21 rue Grandgagnage, B-5000 Namur, Belgium". The official congress languages are English and French.

Namur is a quiet little city on the confluence of the Meuse and Sambre rivers, at the foot of a hill supporting impressive medieval fortifications. The congress atmosphere is relaxed and informal, with a lot of small symposia going on in parallel in adjacent rooms. There will be a welcome cocktail, a congress dinner, and a meeting room available for coffee breaks. Participants are responsible for making their own hotel reservations, but, if necessary, student's rooms will be available.

Registration fee :
members of the IAC and authors of papers: 6000 BF (about $180)
other participants: 10000 BF (about $300) Young researchers under 30 years 2000 BF (about $60) (with certificate of their university)

The fee covers congress attendance, preprints and coffee-breaks.

Partial Congress Programme

The Congress will feature over 30 symposia, including the following: (CHAIRPERSON Subject)

ACALUGARITEI G. (Roumania)
Evolutions and Metaevolutions from the Point of View of the Invariants Associated to the Transformation Groups

BAHG C. (China)
Complex Systems and their Evolution

COLLOT F-C. (France)
Les notions de temps et d' e'volution en Cyberne'tique

FRANCOIS C. (Argentina)
Les syst`emes humains home'ostatiques ou e'mergents

HEYLIGHEN F. (Belgium)
The Principia Cybernetica Project : Computer-supported Cooperative Development of an Evolutionary-systemic Philosophy
JDANKO A. (Israel)
- Cybernetic Systems Approach to History
- Cybernetic Systems Interpretation of the Religious Idea : From the Primitive to the Monotheist

GASPARSKI W. (Poland)
Cybernetics and Human Behaviour

GELEPITHIS P.(United Kingdom)
Invariants of Cognitive Science : Scope, Limits, Implications

STEG D. (USA)
Determinacy and Indeterminacy in Complex Systems

VANDAMME F. (Belgium)
Cognitive Modelling for Knowledge and Information Technology : Manual and Automatic Tools

Submission of papers

People wishing to present a paper in the Principia Cybernetica symposium should quickly send the filled-in application form below, together with an abstract of max. 1 page, to the addresses of the Symposium chairman (Francis Heylighen) and of the Congress secretariat (IAC) below. Submissions or request to the chairman can be done by email only, but for the secretariat it is advisable to send an application in paper form. In principle, all applications should be received by December 31, 1991, but it may be possible to come in late. People wishing to present a paper in a different symposium can directly submit their abstract to the secretariat.

You will be notified about acceptance not later than 2 months after receipt of the abstract, and will receive instructions for the preparation of the final text. Final papers (max. 5 pages) should be ready by the end of the congress. The Proceedings will normally be published by the IAC about 1 year after the congress.

==Deadlines==
* for summaries (1 page max): December 31, 1991
* for paper submission: March 31, 1992
* for final texts (max 5 pages): August 28, 1992

....

For submissions of papers or further information about the Principia Cybernetica project, contact the symposium chairman:

==
Dr. Francis Heylighen
PO-PESP, Free Univ. Brussels, Pleinlaan 2, B-1050 Brussels, Belgium
Phone   +32 - 2 - 641 25 25   Email fheyligh@vnet3.vub.ac.be
Fax   +32 - 2 - 641 24 89   Telex 61051 VUBCO B
==

For congress registration or further information about the congress contact the secretariat:

=== International Association for Cybernetics
Palais des Expositions, Place Ryckmans, B-5000 Namur, Belgium
Phone   +32 - 81 - 73 52 09   Email cyb@info.fundp.ac.be Fax   +32 - 81 - 23 09 45
===
____ Application Form

Symposium "The Principia Cybernetica Project"
in the framework of the
13th Int. Congress on Cybernetics (Namur, 24-28 August, 1992)

Name : .... First name(s) :... Profession and titles:...................................... Institution: ....
.... Address : ....
....
....
Nationality:.. Phone : (office).................... (home)..................
Fax:.. E-mail :...

o I would like to receive more information about the Congress
o I would like to attend the Congress
o I would like to receive more information about the Principia Cybernetica Project (Newsletter)
o I submit a paper for presentation at the Symposium "The Principia Cybernetica Project" (abstract sent to the Symposium chairman AND to the congress secretariat)

Title of Paper :...
....

Date :   Signature:
____
==

Date: Mon, 16 Dec 1991 13:28:34 PST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: marken@AERO.ORG
Subject: Epistemology, Conation

[From Rick Marken (911216)]

Bill Powers (911215.1800), in reply to Wayne Hershberger, says

> we are led to conclude that the
>object of perception and thought is a world existing inside, not outside,
>a brain. We can see how this world of experience is related to what we >conjecture to exist in a physical environment outside of us, but we can >also see that the relationship is not a simple or direct one, nor is it >wholly verifiable because of our peculiar circumstance of being inside >the very system we model and by necessity having to perceive and think >using its equipment.

This seems right to me. But I have gotten behind on this epistemology debate (or maybe I just don't understand it). Could you (Bill) or Wayne give me a short (like two sentence) description of what is being debated. Based on the above statement, I am wondering if Wayne is arguing that there is no physical environment, or that the physical environment is an unwarrented assumption, or what?

On a lighter note, I wanted to ask Wayne if he knew anything about Kathy Kolbe. Did she send a copy of her book to every author of an article in your Volitional Action book? I looked over her book and thought it looked like the typical pop psych stuff. Does this mean that your Voliitonal Action book is getting into the mainstream bookstores? Maybe it'll make the NY Times bestseller list soon.
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: Mon, 16 Dec 1991 17:15:30 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Jan-Olof Eklundh

[Martin Taylor 911216 17:00]

An open letter to Bill Powers:

Bill,
I went to a seminar this morning given by one Jan-Olof Eklundh, of the Computer Vision and Active Perception group at the Swedish Royal Institute of Technology in Stockholm. He and you need to make contact, and I told him so. Also I suggested that I would ask you to send him your Little Man Demo, since it seems to tie in with what he is doing.

One of the projects they are working on is a hardware simulation of the stereo vision system. They have built a "head" that has a pair of eyes and a processor based on transputers (one in what he showed, but they have just acquired a multi-transputer card). The head has no effectors that can move things in the world, but it can move as if it were on a fixed shoulder base, and the eyes have essentially the same movement degrees of freedom that human eyes do. The task that he demonstrated was for the eyes to remain stabilized and focussed on an arbitrary real object, such as somebody's hand, moving in 3-space in a naturally complex visual world.

One of the interesting points, in light of our discussions, was that he asserted the importance of dealing with independent degrees of freedom (scalar control systems) even though the individual degrees of freedom might compute almost the same thing and might be in conflict.
He was quite concerned to reduce "to almost zero" the kinematic analysis (outflow "control"), and had not quite got to the point of appreciating the nature of perceptual control--but nearly, and in some of the control systems I think he had it as a matter of practice if not of principle. He was quite unsure of how to deal with hierarchic control systems,
and I think it is here that the Little Man would be very helpful (as well as the Book, or tutorial papers from Gary).

Eklundh said they had the intention to mount the head on a behaving body, and the research intention is to make a "seeing" machine. It seems to me that there is an opportunity here for a substantive test of the real-world application of PCT ideas in a machine that has to deal with all of the complexity of the visual world, rather than simulated objects.

Eklundh is the person I was trying to remember last week who showed me the importance of visual phase response. I spoke to him about your work, and he was interested to know more. I suggested that he should join the CSG-list, and he may do that. His e-mail address is joe@bion.kth.se

Prof. Jan-Olof Eklundh
Dept of Numerical Analysis and Computing Science
The Royal Institute of Technology
NADA, KTH
S-100 44 Stockholm
Sweden

Please send him a Little Man, if you can.

Martin Taylor

PS. To Gary: I didn't suggest it, but it might be nice if you would send Eklundh some tutorial material on PCT. He knows all about engineering control theory, so that aspect isn't wanted (and he and the audience complained about the "closed-mindedness" of control theorists in much the way some on this group talk about psychologists!).
===

Date: Mon, 16 Dec 1991 18:29:28 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Martin Taylor stuff

[Martin Taylor 911216 17:45]
(Bill Powers 911215.1700)

Consciousness:

I'm not going to use long quotes here. I have the feeling that I must have misstated my hypothesis about consciousness. You took me to treat it as a cause of ambiguity reduction, whereas I think of it more as a byproduct that goes along with ambiguity reduction. The hypothesis is that consciousness is related to choice, which I think is not far removed from your saying that it related to reorganization. If there is a difference, it may be that reorganization would imply that the to-be-used control structure does not exist before the reorganization takes place, whereas choice would imply that it does.

If it is true that unconscious perception occurs and is consistently ambiguous where ambiguity is available in the data, then I would (again vaguely and personally) judge that the appropriate low-level hierarchies do exist and are being used, at least in imagination mode. Those that correspond to high-level references (goals, intentions...) that are close to being satisfied by the input are, in some sense, "chosen" and the others suppressed. It's more conflict resolution than reorganization, in this way of looking at it. (I am quite aware of the weasel-word connotation of "in some sense", but I do not have a precise view of what is going on.

Information and levels:

>>Bill, do you think that ALL possible sequences of events form signals >>that are appropriate and DIFFERENTLY usable by the sequence detector; >>that EVERY set of intensities leads eventually to a different category >>percept?
>

>I don't know what you mean by "the" sequence detector. In my pandemonium>style model there is a separate sequence detector for every sequence that >matters. All sequence detectors operate at the same time, independently, >even when they are logically redundant. Same for category detectors: each >detector is waiting for its inputs to indicate a member of the specific
>category it is specialized to perceive. Different category, different >physical input function.
>

I really put my foot in my mouth (on my keyboard?) there, didn't I. Of course I know that. My intention was not to single out any specific sequence detector, but the set of all sequence detectors, and likewise category detectors. Anyway, having corrected me, you answered appropriately "No", and that is all that is needed to assert that there is an
ever increasing redundancy as we go up the levels; concomitantly, there
is a reduction in control bandwidth-equivalent, and a slowing of the system responses as we go up the hierarchy. This is not to say that the information rate provides the limiting factor, but one has (read "I have") a kind of principle of evolutionary efficiency that says Nature has done a pretty good job of otimizing compromises in the face of physico-mathematical constraints. No proof, but it would not be parsimonious of Nature to provide processing capacity that could not be used, nor to provide sensory/mechanical operations for which
processing was unavailable. Sure, there are lots of other constraints.

Redundancy:

>>The amount (and perhaps type) of redundancy occurring
>>between levels of the hierarchy is, I suspect, critical for the >>stability of the whole system, and of its ability to reorganize.
>
>I don't disagree; I just don't understand. How can there be "redundancy" >between levels that handle different types of variables? I have always >understood redundancy to mean delivering the same information by two or >more means, any one of which would have been sufficient.
>
>This probably isn't a big problem, but it would be good to get our >terminology coordinated.
>

All redundancy means is that a "channel" conveys less information than it might. "Channel" is in quotes because the concept also applies to sets of objects or messages, not only to a transmission channel.

The key idea is that there are two sets of probabilities defined on the same channel. One set of probabilities applies to the capability of
the channel. For example, the channel might consist of a set of four
card suits, and then each turn of a card could convey 2 bits of information. But if it turned out that after a red card showed the next one was always black, then knowling that rule would be enough to tell you the colour
of the next card, so the showing of the card would provide only one bit.
The redundancy of the sequence would be one bit per card.

That's a trivial example, but in general the idea is that the information could, in principle, have been coded into a shorter message on the same channel, if the structure that is expressed in the redundancy were
removed.

The usual example for simple redundancy is distributional--if the letters of the alphabet were equally used, they could convey log2 (26) bits per letter. But e, t, s and the like are much more common than q, z, x,
and so a message made by scrambling all the letters of a text would convey fewer bits per letter (I've forgotten just how much).

In our discussions, we have been dealing with structural redundancy expressed in the syntax of sentences. Because "the" predicts that a noun or adjective will follow with high probability, it restricts the probability distribution for the following word(s), and thus reduces their information-carrying capacity.

Redundancy in communication is normally used to enhance the likelihood that the message received is the one the originator intended should be received. Well constructed redundancy ensures that the most probable kinds of channel errors can be detected, and with luck corrected (for example, if a red card were followed by another red card, then a third before a black card showed, one might expect that the second one was a mistake).

I conceive of the "channel" for a control system (or a set of ECS of equivalent level) as being the control systems to which is (they) supplies reference signals, and that provide it (them) with the data of perception.

Martin Taylor ==

Date: Tue, 17 Dec 1991 10:21:00 MST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Misc

[From Bill Powers (911217.0930)]

Rick Marken (911216) --

I'll let Wayne answer your question about the epistomological argument. I think there is a physical environment, although our models of it can be checked only for mutual consistency (as far as I know).

I haven't received anything from Kathy Kolbe, so she didn't sent copies of her book to EVERY author in Wayne's volume. I guess she liked you best.

Martin Taylor (911216.1700) --

Thanks for the note about Jan-Olof Eklundh. I'm preparing to leave for Boulder where I will attend my son Denny's graduation (in mechanical engineering), and won't be back until Monday, so won't start anything with Eklundh until then. From what you say, he has found his own path toward a CSG-like concept of control theory. The best indication of this is his complaint about the "closed-mindedness" of control theorists! when he does show up on CSGnet, let's throw a welcome party for him by spending some time swapping tutorials. From your description, what he is doing with vision is of great importance to PCT.

(911216.1745) --

OK on your use of consciousness. You're touching on a big area for experimentation that I've simply ignored because the basic phenomena have to be established before theorizing makes any sense. I would call it the subject of attention. I believe that you and I are talking about the same thing here in more or less the same way.

My pandemonium model is "flat" in the sense that all systems have equal importance and all of them work all of the time. The implication is that every control system we own, once turned on, remains on forever, night and day. The only way to take a control system out of action is to send it a zero reference signal. But this does not actually turn a system off, unless you make some rather detailed assumptions -- all systems are composed of balanced pairs of one-way control systems with excitatory reference signals and inhibitory perceptual signals, so that in the absence of either positive or negative acting reference signals there can be no error signal. That's not bad neurologically, except that some comparators (in the brain stem) receive *inhibitory* reference signals (e.g., from Purkinje cells in the cerebellum) and excitatory perceptual signals, so they can't be turned off that way. Maybe, in fact, they never do turn off. But that's a factual question, not a theoretical one.

Anyhow, there seems to be informal evidence that only some of the control systems we have acquired are actually in operation at any given time, and at least the most important ones are those involved in "voluntary" -which is to say conscious -- behavior. We need some experiments here, to see what happens to the parameters of control when attention is taken away from a conscious control process (for example, when another one encounters an unusual problem). How many dimensions of control can actually be in this conscious state at once? I have a hunch that the number is limited (7 plus or minus 2?). The implication is that subsidary control systems do operate outside of consciousness, at least those needed to carry out the details of a higher-order process; it's also possible that systems operating at levels higher than the level of attention at a given time continue working without consciousness, setting reference signals for the conscious level. But the general picture I get certainly does not suggest thousands of equally active systems at every level operating all of the time.

So I am in agreement with all your remarks about consciousness and reorganization even though the HCT model lacks any machinery to handle this phenomenon (other than the basic reorganizing system, which functions without consciouness but is supposedly "directed to the right place" by consciousness). If anybody wants to inject a fresh load of data into HCT, this is an excellent place to do it.

information and levels:

>I really put my foot in my mouth (on my keyboard?) there, didn't I. Of >course I know that. My intention was not to single out any specific >sequence detector, but the set of all sequence detectors, and likewise >category detectors.

When we've known each other a few more years, I will know when a statement is a slip of the head or shorthand and when it needs to be responded to literally. Don't apologize, you're among friends. At least on this net, even real mistakes are forgiven in advance. I hope.

I understand what you mean by redundancy and channel capacity now, and find no fault. I think that your basis for predicting slower operation at higher levels is probably more pertinent than mere stability requirements. Neural transit times computed just by counting synapses and measuring the lengths of pathways are far shorter than actual "reaction times," so most delays must have to do with complexity (and redundancy) of computing processes.

>The key idea is that there are two sets of probabilities defined on the >same channel. One set of probabilities applies to the capability of >the channel. For example, the channel might consist of a set of four >card suits, and then each turn of a card could convey 2 bits of >information.

It sounds as though the concept of probabilities applies mainly to the levels above categories, where we deal in discrete symbols (at least when using logical program rules). You can't compute bits directly when dealing with analogue signals at the lower levels, although I understand that conversion from continuous to information-type measures is possible. When frankly analogue processes are involved, however, it seems to me that using analogue concepts to explain them is more appropriate than trying to force them to fit a tradcition based on discrete phenomena.

Yes, redundancy in messages does help assure receipt of the intended message (as in cyclic redundancy checks). I think I was referring to
paraphrasing, which is a different sort of redundancy in that the physical messages can be vastly different, yet convey the same meaning (among other meanings). I think this is more common kind of redundancy in human communication.

>I conceive of the "channel" for a control system (or a set of ECS of >equivalent level) as being the control systems to which is (they) >supplies reference signals, and that provide it (them) with the data of >perception.

Fine -- so the channel is ultimately defined by the highest level peceptual function involved. Makes sense to me.

Best

Bill P. ==

Date: Tue, 17 Dec 1991 15:04:08 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: redundancy and level

[Martin Taylor 911217 14:30]
(Bill Powers 911217.0930)

>
>It sounds as though the concept of probabilities applies mainly to the >levels above categories, where we deal in discrete symbols (at least when >using logical program rules). You can't compute bits directly when >dealing with analogue signals at the lower levels, although I understand >that conversion from continuous to information-type measures is possible. >When frankly analogue processes are involved, however, it seems to me >that using analogue concepts to explain them is more appropriate than >trying to force them to fit a tradcition based on discrete phenomena.
>

Actually, no. The concept applies equally and without forcing whether the probabilities are discrete or take the form of a distribution on a continuous variable (or manifold). In human communcation practice, probably the greatest redundancy is in the acoustic channel, where the effective sampling rate (Nyquist crierion) for speech is over 10 kHz (nearer 40 kHz if you want to deal with everything we can hear, which makes the speech signal intrinsically 75% redundant right off the bat). But the transitions in the speech signal that are identified (rather arbitrarily) by segmantation algorithms happen at around 100 per second tops. Then, we get into sequence redundancy, and eventually to categories. I think that the redundancy gets less once we get above the level of categories, because at higher levels it will depend more on the recipient's understanding than on the transmitter's expectation as to whether the redundancy will be needed, and how it should be structured.

All structure implies redundancy. Redundancy detected implies structure. The best way to incorporate redundancy into a complex channel is to have a model or algorithm that specifies how the channel is to be used. In my mind, the difference between a scientific model and a statistical description is only that the model provides a much more redundant description of the phenomena. And that's a 20th century statement of Occam's Razor.

Really, discreteness and continuousness are irrelevant when considering redundancy. Point of view is the important thing: what do you know about how the observations COULD turn out if you knew nothing of their structure; what do you know about how they could turn out given what you know of their structure, and what do you know after you have made the observation.
The difference between the first and the second is the redundancy of the channel, in some circumstances known as syntax, but always as structure; the difference between the second and the third i the information provided by the observation, in some circumstances known as the information content of the observation.

Parenthetically, it is my belief that none of this really makes much sense in a frequentist version of probability, because observations cannot be really repeated, especially in the linguistic-social environment in which this came up first. I think you have to use a subjective probability view, based on expectations about observations (which may well derive from past frequencies of event types). But then, I think this about almost all uses of probability, so maybe my opinion doesn't count for much in that area.

Martin ===

Date: Wed, 18 Dec 1991 10:48:00 CST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: TJ0WAH1@NIU.BITNET
Subject: Kathy Kolbe

[From Wayne Hershberger]

Last week I received a complimentary copy of _The _conative_ _connection_, a Random House book written by Kathy Kolbe, a management consultant in Phoenix, Arizona. (Some of you have also received copies---Dave, don't forward anything to me; I'm sure that what you got was intended for you, not me.) In a cover letter, she mentions my book _Volitional_action:_Conation_ _and_control_ to which so many of you contributed chapters, and observed that, "You and the other contributors to your publication have taken a very interesting and important approach to a topic we all obviously agree merits considerably more attention than it has previously been given." Amen, to that.

I have given the book a quick read and culled some telltale comments quoted below that CSGers may find familiar. I can not count the number of times I've heard similar comments from CSGers, the likes of Bill Powers, Dave Goldstein, Ed Ford, Jim Soldani, Perry Good, etc. (i.e., particularly the applied contingent).

Some quotations:

Speaking about conflict she writes,

"Give others the same freedom to be themselves that you yourself need. Don't try to impose your will on another's" (p. 89).

"To remove strain, stop doing what you're doing" (p.87).

Talking about the locus of control she says,

"Computers will never replace people because they can't take initiative. That's the conative dimension that only living things can contribute" (p.75).

Referring to what we call a reference signal she says,

"By removing herself from the situation long enough to regain her centering point, Helen went back to her desk able to once again accommodate the situation" (p. 72).

Speaking of personnel management, she says,

"no matter how good the plan is or how apt the planners, if all
the players don't buy in, the show doesn't come off" (96).
"No amount of training, cajoling or force can make one person live up to the expectations of another and only letting go of unrealistic demands can remove the source of conative tension"
(p. 91).

Kolbe's says that her father, E. F. Wonderlic, is the originator of the concept of personnel testing. Perhaps this accounts for her having developed the Kolbe Conative Index, a psychometric instument (included in the book) that is supposed to measure a person's conative style. It may well be that her Index is addressing the same sort of individual differences that Tom Burbon has measured with some of his subjects, finding incredible retest reliability (what parameters were those Tom, something to do with gain and slowing?). Perhaps that also accounts for her attention to Dave Goldstein's chapter in my book in which Dave mentioned his control-theoretic psychometric instruments: the Life Perception Survey and the Life Perception Profile.

Kathy Kolbe appears to be a kindred spirit. I hope she will join us on the CSGnet. I'll get in touch with her, but our president Ed Ford, might even want to give her a personal visit. Ed, did you get a copy of her book? Why don't you send her a couple of yours?

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 Wed, 18 Dec 1991 14:03:48 EST
Reply-To "Control Systems Group Network (CSGnet)" <CSG-@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From:"CHARLES W. TUCKER" <N050024@UNIVSCVM.BITNET>
Subject:A RECENT VERSION OF THE STATISTICS SO FREQUENTLY USED IN SOCIAL SCIENCE

A TABLE SHOWING THE RELATIONSHIP BETWEEN

SEVERAL DESCRIPTIVE STATISTICS*

	r	r2	k2	k	E
1.00	1.00	.00	.00		100 %
	.9995	.999	.001	.032	97 %
	.9987	.997	.003	.054	95 %
	.995	.99	.01	.099	90 %
	.954	.91	.09	.299	70 %
	.90	.81	.19	.435	56 %
	.87	.756	.244	.493	51 %
	.865	.748	.252	.50	50 %
	.80	.64	.36	.60	40 %
	.71	.50	.50	.70	30 %
	.60	.36	.64	.80	20 %
	.50	.25	.75	.87	13 %
	.40	.16	.84	.92	8 %
	.31	.10	.90	.95	5 %
	.20	.04	.96	.98	2 %
	.10	.01	.99	.995	0 %
	.00	.00	1.00	1.00	0 %

DEFINITION AND INTERPRETATION OF THESE STATISTICS**

All of these measures describe two variables (X, Y) within a particular sample:

r is a correlation (or coefficient of correlation) which describes the linear association of one variable with another. It can also be characterized as "... a relative measure of the degree of association between two series " of values for two variables. It varies between 1 (perfect positive correlation) to -1 (perfect negative correlation). The closer this measure is to a perfect correlation the more confidence one has in "predicting" the values of one variable from another variable.

r2 is a measure of "explained" variance (or coefficient of determination) which describes "shared" variation or the amount of variance that one variable is "explained" by the other variable or the proportion of the sum of y2 that is dependent on the regression of Y on X. The larger the numerical value of this measure the more confidence one has in "predicting" the values of one variable from another.

k2 is a measure of "unexplained" variance (or coefficient of nondetermination) which describes "unshared" variation or the amount of variance that one variable is NOT "explained" by the other variable or the proportion of the sum of y2 that is independent of the regression of Y on X. The smaller the numerical value of this measure the more confidence that one has in "predicting" the values of one variable from another.

k is a measure (called coefficient of alienation) which describes the lack of linear association of one variable with another or the ratio of the standard error of estimate to the standard deviation of the variable. The smaller the numerical value of this measure the more confidence one has in "predicting" the values of one variable from another.

E this measure is computed by (1-k)100 and is called an "index of forecasting efficiency" (Downie and Heath, 1965: 226) and indicates the "improvement" for a prediction by knowing the coefficient of correlation (r) for two variables as contrasted with knowing nothing about the linear association of the two variables. For example, with a coefficient of correlation of .71
one can "predict" the values of one variable from another 30% better (on the average) than one could "predict" those values WITHOUT any knowledge of the relationship between the two variables OR one has decreased the size of the "error of prediction" by 30% (on the average) by knowing that the
correlation of the two variables is .71.

REFERENCES

Arkin, Herbert and Raymond R. Colton. 1956. Statistical Methods. College Outline series, Forth Edition, Revised.

Downie, N. M. and R. W. Heath. 1965. Basic Statistical Methods. Second Edition. New York: Harper and Row.

____

*compiled by Charles W. Tucker with the encouragment and assistance of the Control Systems Group (especially Gary Cziko) and the comments of Jimy Sanders. Other comments appreciated - N050024 AT UNIVSCVM.BITNET
**It should be noted that these descriptions and interpretations, especially those involving "predictions" are limited to a particular sample; if another sample is not a random sample from the same population then predictions about the other variable ("Y") will be unpredictably worse than the original sample.
===

Date: Thu, 19 Dec 1991 08:41:08 PST
Reply-To:Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: marken@AERO.ORG
Subject:Kolbe, language

[From Rick Marken (911219)]

Wayne Hershberger -- thanks for the info about Kolbe's book. You have an astounding ability to find the wheat in the chaff. I scanned the book (and rescanned it) and could find little of merit. You found some nice quotes. I feel like Scrooge to your Bob Cratchet. If Kolbe does get on the net, I hope she's not like the ghost of christmas future!!

To the linguistic crew -- I'm enjoying the respite from the language discussion, not because I didn't enjoy it but because I didn't always understand it. Nevertheless, while relieving myself during intermission at a highly boring meeting yesterday (requiring much coffee) I noticed a sign that seemed relevant to that discussion so I will share it with you. The sign said "Please do not throw paper or cigarettes in the urinal". The sign interested me as an example of the precedence (in this case) of meaning over grammar in language. I immediately understood the sign to mean "don't through this stuff into the toilet". It took me a while to realize that a proper grammatical reading would lead me to think that it was ok to throw things into the unrinal as long as I was not standing in the urinal at the time.

What was interesting to me was how the context of what I was doing and what I know about plumbing and stuff led to the instant creation of what I believe was the appropriate image (the image that those who made the sign intended that I have). If I were a lawyer and my client had thrown stuff into the urinal, I would take the case if he were not standing in the urinal while throwing the stuff.
What does all this mean?
That I should get back to work I suppose.
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: Thu, 19 Dec 1991 14:58:20 MST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Ed Ford <ATEDF@ASUACAD.BITNET>
Subject: Wayne, David, & Dick

from Ed Ford (911219.15:00)

Wayne Hershberger -

Wayne, I did receive a copy of Kathy Kolbe's book along with a note from her requesting we get together and talk. I was told by her office that she is out of the country and will return on the Jan. 1st. I will be meeting with her on Friday, Jan. 3rd. And, as you suggested, I sent her a copy of Freedom From Stress. I haven't had time to look through her book but plan to do so prior to our meeting. Will keep you informed.

David Goldstein -

Dave, I understand what you are saying concerning the difference between random and learned suicide. I just don't see things the same way you do. I think if people have thought about suicide in the past, it would more likely come to their mind than to those who have never thought about it. Also, it would more likely be considered or taken seriously as an option in those people in whom there would be no conflict with their higher order value or belief systems.

I think people attempt suicide when they see no better option regardless of whether they've tried or just thought about it before or not. I think when people are dealing with one or more conflicts, especially when they're both chronic and very painful conflicts, they are going to reorganize. If, when they are reorganizing, suicide is among the options that appear in their mind, then they are going to evaluate that option and compare it to their other conscious options. If they foresee, down the road, no reduction in pain from the other conscious options, then suicide is going to look like the best viable choice. The more they lack belief in their system's ability to reconcile the difference between their reference signals (i.e. their various wants and goals) and their perceptual feedback (i.e. their present perception of how things are), the more they will be inclined to seek relief just from the pain itself, having given up on the possibility of resolving the conflict(s). So now it's just a matter of reduction of pain. If, to them, suicide offers the best solution for reduction of pain, then they will attempt suicide, unless, like Hamlet, they should ponder the famous "perhaps to dream" in his famous soliloquy. And that returns us to having a conflict with values or beliefs.

Dick Robertson -

I down loaded the info on your student and now I can't find it. If you would again send me her name, mailing address plus her CSGnet address, I would be appreciative. Sorry about that. I guess it's about time to clean up my messy office. Who knows what I'll find!

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

Date: Fri, 20 Dec 1991 14:35:33 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From:
Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Kolbe, language

[Martin Taylor 911220 14:25]
(Rick Marken 911219)

(If a wierd message with the above header and some garbage characters reaches you, don't blame me, blame a bad phone line that kicked me off in the
middle of responding)

>Rick says:
> The sign said "Please do not throw paper or cigarettes in the
>urinal". The sign interested me as an example of the precedence (in this >case) of meaning over grammar in language. I immediately understood the >sign to mean "don't through this stuff into the toilet". It took me
>a while to realize that a proper grammatical reading would lead me >to think that it was ok to throw things into the unrinal as long as I >was not standing in the urinal at the time.

I would have interpreted the notice as suggesting that if you find stuff in the toilet, you shouldn't throw it.

Yes, most of what we understand is situation-dependent. There are some experiments, which I guess I could look up, that show how people can solve puzzles in logic easily if the logic conforms to the sense of the words
(i.e. it conforms to an image of what should happen) much more easily than
if it is abstract (in terms of A, B, X etc.), and very much more easily than if it contradicts what the sense of the words is.

I think that almost all language understanding is based on situation-related meaning, and that syntactic constructions are used only as a backup in case of ambiguity in the possible semantics and pragmatics. The speaker has to incorporate the syntax in what is said, but the listener (or reader) doesn't have to rely on it, and usually only uses it as a check. And, as Rick points out, the speaker/writer often misuses the syntax without misleading the hearer/reader. This is possible only because the syntax does not drive the
interpretation. If a successful interpretation can be achieved in contradiction to the syntax, so be it. The interpretation stands.

Martin Taylor ==

Date: Fri, 20 Dec 1991 14:15:29 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Syntax and meaning

[from Gary Cziko 911220]

Martin Taylor (911220 14:25) says:

>Yes, most of what we understand is situation-dependent. There are some >experiments, which I guess I could look up, that show how people can solve >puzzles in logic easily if the logic conforms to the sense of the words
>(i.e. it conforms to an image of what should happen) much more easily than >if it is abstract (in terms of A, B, X etc.), and very much more easily than >if it contradicts what the sense of the words is.

The example that comes most readily to mind is the one which asked subjects something like: "All cards with an even number one side must have a vowel on the other." Four cards are then shown: 4, E, 7, X and the subject is asked which cards must be turned over to make sure that the rule has not been violated. This is in practice surprisingly hard (try it), although logically very simple (most adult subjects get it wrong).

It becomes much simpler if you say: "All receipts over $50 must be initialed on the back by the manager." Four receipts are then shown: $75; back with no initials, $25, back with initials. This is very easy although logically identical to the number and letter task. (Was it Studdert-Kennedy who did this study?) This task makes real-world sense. The previous one does not.

>I think that almost all language understanding is based on situation-related >meaning, and that syntactic constructions are used only as a backup in case >of ambiguity in the possible semantics and pragmatics. The speaker has to >incorporate the syntax in what is said, but the listener (or reader) doesn't
>have to rely on it, and usually only uses it as a check. And, as Rick points >out, the speaker/writer often misuses the syntax without misleading the >hearer/reader. This is possible only because the syntax does not drive the >interpretation. If a successful interpretation can be achieved in
contradiction
>to the syntax, so be it. The interpretation stands.

This certainly appears to be my experience as well. There are many times I've attended lectures where the speaker said just the opposite of what he meant (using "visible" instead of "invisible" or forgetting a "not" somewhere) and nobody even seemed to notice (except me, of course, but I'm special).

But doesn't this pose a problem concerning the evolution of language? How could language develop all the very complex syntactic devices if they were not or very seldom needed for meaning, particularly before writing was invented?

There are also apparently some differences across language with respect to the attention paid to syntax. If I ask you what "The book ate the chair" means, you will probably interpret the book as the eater and the chair as the eatee. But in Italian or Spanish I don't think anyone would come up with such a syntactically driven, odd-meaning interpretation.

I suppose one could say that even if 95% of all utterances could be undertood without syntax, the remaining 5% could in itself provide enough of a selection pressure for sophisticated syntactic devices. And since syntax sometimes IS needed, it may be easier to use it just about all the time rather than switch it on and off. Hell, we've got to put the words in some order anyway, so why not use a syntactic one.--Gary

P.S. The answer to the letter and number card task is 4 and X. Most people think that you have to turn over 4 and E. (Boy, I hope I got this right. What a good way to make a public fool of myself.) For the receipt task, the $75 receipt must be turned over (to make sure it has been initialed) as well as the one without initials (to make sure it is not over $50). ==

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: Sat, 21 Dec 1991 14:06:09 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From:
Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: handedness

[Martin Taylor 911221 13:45]

What is the place of handedness within PCT? Why are there right-handed and left-handed people, and why are almost all the right-handed people left-brained for language, whereas of the left-handed people only about half are right-brained for language, if that?

I can imagine that handedness reduces the opportunity for low-level conflict among control systems that could otherwise equally well perform the same function. That notion is supported by the fact that a few things are easily done left-handed but not right, after training (e.g. in my culture one eats with a knife in the right hand and a fork in the left, and it is very awkward for me to use a fork with the right hand, even when I don't have a knife).

How is handedness expressed? Is it a question of less precise or less rapid perceptual processing? Of gain differences? or what? Is the difference among these possibilities (or others) directly testable? Remember that handedness is not a unitary function, and that there are degrees of handedness, which vary from joint to joint. A person may be strongly right-fingered, but weakly right-wristed, for example (somebody at this laboratory studied this about 30 years ago, and his conclusions may have been altered by later work that I don't know about).
Would PCT provide any theoretical background for a relationship between handedness and writing direction, or between writing direction and the
degree of phonetic representation in a script (it seems that a shift from right-left to left-right writing occured in Greek script at the same time vowels were introduced, and neither Arabic nor Hebrew have vowels represented as strongly as consonants--they write right-left). I think this paragraph represents a wild shot in the dark, but one never knows (for background,
look at "The Alphabet and the Brain" (Eds De Kerckhove and Lumsden, Springer Verlag, 1988))

But regardless of these more subtle effects of handedness, a simple testable description of it in PCT terms would be quite interesting.

Martin ==

Date Sun, 22 Dec 1991 17:22:25 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From:goldstein@SATURN.GLASSBORO.EDU
Subject:how do I get email addresses

I want to send a private message to someone on CSGnet
but the old email address is not working.

How can I check the emai address of someone?

Thanks,
David Goldstein

internet: goldstein@saturn.glassboro.edu
===
Date: Mon, 23 Dec 1991 13:26:26 GMT
From: cam@AIFH.ED.AC.UK
Subject: Re: handedness

From Chris Malcolm:

I am a right-handed person who was forced to write left-handed. As a consequence I have specialised in fine manipulation with the left hand,
but forceful arm-strength tasks with the right, i.e., I am left-handed but right armed.

AT the age of 18, faced with important exams and very bad handwriting, I taught myself a new writing script. In 3 months I could write as fast
the new way as the old. After a year I was beginning to forget the old -with my left hand. But if I tried writing with my right hand, it automatically still wrote the old way.

After about ten years I noticed that my right hand, with almost no practice at all, had largely switched over to my new style of writing, but if
I tried tracing my signature in the the dust with my feet, then I still used the old style.

After about 30 years, my feet still want to write with an amalgam of the old and new style. My right hand is now fully converted to the new style. Note that I probably write about two words per year with my right hand --
hardly enough to be (I would guess) a significant learning experience.
I have so far forgotten my old style of handwriting that when I want to remember it I experimentally trace letters in the air with my right hand, because it has retained the old skill closer to the surface of performance recall than my left (writing) hand.
No, I don't stutter, hold my pen in a funny way, or suffer in any way
from having been forced into left-handed writing.
===
Date: Mon, 23 Dec 1991 14:57:09 -0600
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: cmcphail@UX1.CSO.UIUC.EDU
Subject: Re: how do I get email addresses

\>
\>I want to send a private message to someone on CSGnet
\>but the old email address is not working.
\>
\>How can I check the emai address of someone?
\>
\>Thanks,
\>David Goldstein
\>internet: goldstein@saturn.glassboro.edu
\>
DAVID:

TRY THESE. CHEERS AND A HEALTHY/HAPPY NEW YEAR. CLARK

```
*      Control Systems Group Network (CSGnet)
*
*      Review= Public    Subscription= Open          Send= Public
*      Notify= Yes       Reply-To= List,Respect     Files= Yes
*      Validate= Store only
*      Confidential= No
*      Local= *UIUC.EDU,UIUC*,NCSA*,UIUCVMD
*      NOTEBOOK= YES,N,WEEKLY,PUBLIC
*      STATS= Normal
*      Ack = Yes
*
*      OWNER= CZIKO@UIUCVMD (Gary A. Cziko)
*
```

*
dgaw@ADS.COM GAW David
marken@AEROSPACE.AERO.ORG MARKEN Rick: Aerospace Corp, LA
cam@AIFH.EDINBURGH.AC.UK MALCOM Chris
JEVARGAS@ANDESCOL Jaime Eduardo Vargas
avery.andrews@ANU.EDU.AU ANDREWS Avery, Australian Nat. U., Can
well!dooley@APPLE.COM Dooley, Jeff San Jose State University
atedf@ASUVM.INRE.ASU.EDU FORD Ed; Arizona St U, Tempe
bclyde@AUG1.AUGSBURG.EDU CLYDE Robert
HARNDEN@AUVM HARNDEN Eric
A6212DAN@AWIUNI11 PARZER Peter
cutmore@BEN.DCIEM.DND.CA CUTMORE Tim, Environmental Med.,
cjoslyn@BINGVAXU.CC.BINGHAMTON.EDU JOSLYN Cliff
bill@BIOME.BIO.NS.CA Bill Silvert
lars@BLUEEYES.KINES.UIUC.EDU GOLDFARB Larry; U. Ill.-Urbana;Kinesi
LACOURSE@CALSTATE LACOURSE Michael; Cal. State U
USEREA4K@CC.SFU.CA WINTER A.; Simon Fraser U, BC
wet!wcl@CCA.UCSF.EDU LITTLEWOOD Bill, San Francisco
bnevin@CCB.BBN.COM NEVIN Bruce E.
rrcsg-l@CCS.CARLETON.CA Carleton Distribution
Wolfgang.Zocher@CDC2.RRZN.UNI-HANNOVER.DBP.DE ZOCHER Wolfgang cochran@CLIO.STS.UIUC.EDU COCHRAN Cynthia; U Ill.-Urbana
RYATES@CMSUVMB YATES Robert
ncon@COMPLEX.CCSR.UIUC.EDU PACKARD Norman: U. Ill-Urbana
76010.1363@COMPUSERVE.COM SATTLEY Kirk & Joanne
mar@CS.ABER.AC.UK RODRIGUES Marcos Aurelio: U. Coll
zhang@CS.UBC.CA ZHANG Ying; UBC; Computer Science
GASSER@CSGHSG5A GASSER Pius; St. Gallen CH
JUNGER@CWRU JUNGER Peter D., Case Western Res
klaw@DAVID.NEWCASTLE.EDU.AU Kim LAW
goerner@DG-RTP.DG.COM GOERNER Sally J.
mmt@DRETOR.DCIEM.DND.CA TAYLOR Martin, Environmental
eplunix!peter@EDDIE.MIT.EDU CARIANI Peter, MIT
mike@EEG.COM WARD Mike
Psy_Delprato@EMUNIX.EMICH.EDU DELPRATO Dennis, E Mich U, Psy
demoor@ESAT.KULEUVEN.AC.BE DeMOOR Bart
ray@ESPRESSO.BOEING.COM ALLIS Ray
GRANJEON@FRULM63 GRANJEON Eric
GUILLOT@FRULM63 GUILLOT Agnes
LGMUSA@GALLUA MUSA Larry
haider@GMDZI.GMD.DE HAIDER Christian
usc@GMDZI.GMD.DE SCHNEPF Uwe; Germany
MCCLEL@GRIN1 McCLELLAND Kent; Sociology; Grinnell C
n8442161@HENSON.CC.WWU.EDU TUTTLE Patrick
MFMISTAL@HMARL5 TALMON Jan
thurman%hrlot1.decnet@HQHSD.BROOKS.AF.MIL THURMAN Richard; Wms. AFB, AZ
jshsu@ICAEN.UIOWA.EDU HSU J.
YOAVSYS@ILNITE COHEN Yoav
kdeacon@INETG1.ARCO.COM DEACON Keith: ARCO Petroleum
Oded.Maler@IRISA.FR MALER Oded
ccchen@ISS.NUS.SG Chung-Chih Chen
N1.PEP@ISUMVS PATTERSON Pat
mark@ITD.DSTO.OZ.AU NELSON Mark: DSTO Salisbury,South Aus
SBROWN@KENTVM at vmd.cso.uiuc.edu Steven R. Brown
cunningb%mon1@LEAV-EMH.ARMY.MIL CUNNINGHAM Bill; Fort e, VA
eprince@LYNX.NORTHEASTERN.EDU PRINCE Eileen
bstring@MAINZ-EMH2.ARMY.MIL Robert L. Stringfield
0004742580@MCIMAIL.COM FORSSELL Dag, Valencia, Ca
4912499@MCIMAIL.COM CLARK Bob; Cincinatti OH
4972767@MCIMAIL.COM WILLIAMS Greg; Gravel Switch KY
szietz@MCS.DREXEL.EDU ZIETZ Stanley, Drexel U., Phil.PA
mcnamara@MGI.COM MACNAMARA Curt: Mgmt. Graphics, M
ian.nimmo-smith@MRC-APPLIED-PSYCHOLOGY.CAMBRIDGE.AC.UK NIMMO-SMITH Ian JEB1@MSSTATE BOGGESS Gene: Mississippi State U

21003HGC@MSU COOK H. Gary
AX53@MUSICA.MCGILL.CA AROCHA Francisco
inj3@MUSICB.MCGILL.CA WALTERS Joel, McGill U, Montreal
RAINDROP%UNCVX1.BITNET@NCSUVM.CC.NCSU.EDU J. Hsu
TJ0WAH1@NIU HERSHBERGER Wayne: Northern Illin
rlim!lenting@NLUUG.NL LENTING Jacques; NL
KHACKER@NMSUVM1 HACKER Kenneth
ASLANSKY@NTUVAX SLANSKY Nuala
DOSCSL@NUSVM CHOW Siew Loong
FCCHEN@OREGON CHEN Bob
HEIDISV@OREGON SVEISTRUP Heidi.
jmlubin@PHOENIX.PRINCETON.EDU LUBIN Joseph M.
carnahan@RIA.CCS.UWO.CA CARNAHAN Heather
goldstein@SATURN.GLASSBORO.EDU GOLDSTEIN David, New Jersey
HALEY@SDNET HALEY Bill
TBOURBON@SFAUSTIN BOURBON Tom: S.F. Austin U, Nacog
DTANNER@SJSUVM1 TANNER Don: San Jose State U
CSCHRODE@SNYESCVA SCHRODE C.
olsen@SUNA0.CS.UIUC.EDU OLSEN Ray; CS; U. Ill.-Urbana
nhi@SWANEE.EE.UWA.OZ.AU TA Nhii
powersd@TRAMP Bill Powers
===

Date:Mon, 23 Dec 1991 15:54:36 EST
From: "CHARLES W. TUCKER" <N050024@UNIVSCVM.BITNET>
Subject: Re: how do I get email addresses
In-Reply-To: Message of Sun,
22 Dec 1991 17:22:25 EST from <goldstein@SATURN.GLASSBORO.EDU>

Dear David,

I believe you can get the entire list of participants on any list bytyping this:
TELL LISTSERV AT UIUCVMD REV CSG-L [obviously this is for our list] but
you can substitute "rev csg-l" with any list name, e.g., rev quarls-l.
Try it. If you make an error you will quickly know.
Regards,
Chuck ===

Date: Mon, 23 Dec 1991 20:40:57 EST
From: goldstein@SATURN.GLASSBORO.EDU
Subject: subscribers

tell listserv at uiucvmd rev csg-l
===

Date: Tue, 24 Dec 1991 07:49:00 CST
From: TJ0WAH1@NIU.BITNET
Subject: MERRY CHRISTMAS, boss reality

[From Wayne Hershberger]
****
To the Control System Group:

MERRY CHRISTMAS TO YOU ALL!
****
(Bill Powers (911215.1800)
>We keep going around and around on the same points without
>getting anywhere.

Amen. We are back where we began. You say:

>And ultimately we are faced with a paradox, the one you and I
>have been arguing about. We find by experimentation that the >presence of certain signals in a brain is the sine qua non of >perception. Remove those signals and you destroy, as far as the >victim is concerned, a chunk of the immanent order. Yet you

>don't destroy it for anyone else. What other conclusion can we >reach but that perception is absolutely contingent on those >signals? That puts us, as perceiving entities, inside the brain. >To deny that would be to destroy the whole structure of >perceptual and conceptual organization we have so painfully >built up. That structure is at least as well worked out as any >metaphysical argument in words, and a whole lot better tested
>experimentally.

And I say, as I did back in July (Hershberger 910723) that ablation (to which you refer above: "Remove those signals"), the technique pioneered by Pierre Flouren's in the 17th century to localize mental functions identifies certain NECESSARY components of the various functions we call mental (e.g., vision). It does NOT identify the NECESSARY AND SUFFICIENT components. Without the photon there is no vision--for anyone. Ablate photons and we are all blind. This means that the PROPRIETARY aspect of our respective experience (my perceptions versus your perceptions) are contingent upon our respective brains. That is the argument you are making, right? But that does NOT put us IN our respective brains! Our feet are too big.

So, you are right, this is where we came in. Perhaps, it is time to take a different tack. Let me reciprocate by asking you what, if anything is wrong with the following remark (using your terminology): Bill Power's HCT model models an aspect of boss reality (the organism aspect), the other aspect (the environment aspect) already having been well modelled by contemporary physics.

Now that you're in Colorado, I trust your "kids" will be joining you and Mary for Christmas. Have a merry one.

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: Wed, 25 Dec 1991 14:40:05 -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: Syntax;probability & redundancy; boss Kringle

[From Bill Powers (911225.1100)]

MERRY CHRISTMAS (with suitable cultural disclaimers).

Went to Boulder to see son Denny graduate with BSME degree. Fun. Got flu on Sunday for drive home. No fun. Mary did 300 miles of it over snowy tundra and icy passes while I snored or muttered delirious advice. Better now, although the higher centers are the last to know for sure.

Martin Taylor (911220) --

>I think that almost all language understanding is based on situation>related meaning, and that syntactic constructions are used only as a >backup in case of ambiguity in the possible semantics and pragmatics. >The speaker has to incorporate the syntax in what is said, but the >listener (or reader) doesn't have to rely on it, and usually only uses >it as a check.

Maybe syntax is most useful when constructing new sentences. Very little of what passes for communication is actually new; consider TV, or those action movies where the script goes "Got him. Yeah. Freeze! You OK?"
Once you've figured out a sentence from its syntax, the next time you hear it you can just go from the word-sequence to the meaning. You don't even really need to distinguish

the individual words, as long as the whole thing sound sorts of familiar. I pledge a liegance to the US flag and all that. I used to think that forspacious was a kind of weather
phenomenon, as in oh beautiful forspacious skies. Of course this means that you reduce the range of meanings you're prepared to recognize. Someone says "Controlled variable" and you hear "control variable." Communication without, or with only a minimum of, syntax is possible, but the capacity for making and understanding fine distinctions is lost. That's my objection to ghetto talk: without the tool of syntax it's hard to say something like "If he had been willing to listen, I would have explained the difference between saying one is going to do something and actually intending to do it." Anyone ignorant of syntax can say "Pick that up. Bring it here. Dig there in the yucky part. Lunch time. Get lost." Those aren't really sentences; they're "utterances" with meanings people can learn. We can discern syntax in them if we look for it, but it's not really essential.

One historical reason for developing syntax as a formal tool of communication may be legal: "I don't know what the salesman told you, Mrs. Hammurabi, but this is the contract that you imprinted." The rules of games and traffic laws similarly require some formal way, some agreeable way, to narrow down the possible meanings in case of disputes. If someone wants to make a big deal out of whether I said he could have "a" goat or "the" goat (that we were looking at), he'd better arrange for "a" and "the" to have formally different significance, in advance.

Gary Cziko (911220) --

Same subject. It's very hard to spell out rules verbally so they can't be reinterpreted. For example, in your first game, "All cards with an even number [on] one side must have a vowel on the other," the solution assumed that all cards have a letter on one side and a number on the other. I'd turn over the 7 to make sure there's not an even number on the other side. In the supposedly easier one, "All receipts over $50 must be initialed on the back by the manager," I'd check the $25 receipt, too, just to make sure the manager wasn't simply initialing all the receipts without looking at the amounts, perhaps before the receipts are filled in, and I'd turn the initialed ones over to see if he's initialing the front, too. The rule doesn't say he should initial ONLY the backs on ONLY the receipts over $50, AFTER they are filled in. All rules seem to involve an unspoken background of assumptions, which are sometimes pretty hard to track down to eliminate the last ambiguity. Sometimes the only way to do it is to port the statement over to mathematics or symbolic logic. Should the manager initial a receipt for $750,000?

Martin Taylor (911221) --

Handedness wouldn't seem to come out of control theory as an a priori principle: it's just a phenomenon. If you practice doing everything with your right hand, you'll probably be right-handed. If there's a bias in the nervous system (or in the environment) that makes learning things left-handed easier, you'll be left-handed. I don't think that handedness per se exists; it's just that something may exist that leads to using one hand more than the other. When we find it (assume we bother to look) we will then be able to explain handedness and probably a few other things. If it weren't for scissors and words like "sinister" I doubt that handedness would count as a very important phenomenon.

If you want to study handedness CT-wise, set up a task that you can model, determine the parameters of control for each hand or body part, and see how the parameters change with training.

David Goldstein (911223) --

The "tell listserv at" command only works with some mainframes; it means nothing to the Fort Lewis VAX system, or I suspect to most others. To get the csg list you have to send a message to the address "listserv@uiucvmd". The "subject" doesn't matter. The message should be review csg-l (country

The (country part is optional -- it tells you some details about how many subscribers there are and where (bitnet only).

Martin Taylor (911217) --

>>When frankly analogue processes are involved, however, it seems to me >>that using analogue concepts to explain them is more appropriate than >>trying to force them to fit a tradcition based on discrete phenomena.
>
>Actually, no. The concept applies equally and without forcing whether >the probabilities are discrete or take the form of a distribution on a >continuous variable (or manifold).

Without forcing? I think we're diverging into two different universes of discourse here. At the lower levels, where continuous perceptual variables are the phenomena to be accounted for by a model, the most likely model (in my opinion) is an analog model: input neural frequencies convert to chemical concentrations which govern the output rate of firing of the neuron. Knowledge of the physical processes involved should be enough to provide a description of how an output frequency depends on a set of input frequencies. It's true that because of the underlying granularity of the signals, there will be inherent noise, and thus a given output signal at a given time has only a probability of bearing one exact relationship to the inputs. But that probability, in the normal range of most perceptions, is very high. Consider any visual image, any taste, any sense of skin pressure or effort: these impressions have essentially no discernible noise in them. By concentrating on the probabilistic aspect of a low-noise signal one runs the risk, it seems to me, of overlooking the significance of the signal itself: i.e., the 97 per cent that is not noise.

> In human communcation practice, probably the greatest redundancy is in >the acoustic channel, where the effective sampling rate (Nyquist >crierion) for speech is over 10 kHz (nearer 40 kHz if you want to deal >with everything we can hear, which makes the speech signal intrinsically >75% redundant right off the bat).

I don't buy the analogy. There is, as far as I know, no "sampling" at all in the acoustic channels, at 10 KHz, 40 KHz, or any other rate. Why not 10,000,000 KHz? This is not a synchronous detector or an A-to-D converter. You're extending the properties of a digital sampling system to a nervous system, I think illegally. You could say that a transistor amplifier "effectively samples" the music input signal at a rate of 200,000 KHz, but if it's not a sampling system to begin with, the model is simply wrong. There would be, for example, no aliasing problem in the transistor amplifier. This is somewhat like the TOTE unit, which Miller, Galanter, and Pribram wanted to apply to everything, even spinal reflexes, as if all systems in the body were little programmed wait loops. But why not use the right model instead of one that's a little like the right one in a few respects? The properties of the right model (i.e., one that's appropriate to the known mechanisms) are likely to apply correctly; those of the wrong model, that assume nonexistent mechanisms, are likely to apply incorrectly.

>All structure implies redundancy. Redundancy detected implies structure.

This is beginning to sound like one of those truths that are so true that they mean nothing. If these statements are true, where does this get us in understanding how the perceptual system works? It seems to me that using only the notion of redundancy, we couldn't tell the difference between perceiving a taste and perceiving a procedure. Both kinds of perceptions represent convergence from multiple detailed inputs to a single dimension of output. I guess I'm asking "so what?"

>In my mind, the difference between a scientific model and a statistical >description is only that the model provides a much more redundant >description of the phenomena. And that's a 20th century statement of >Occam's Razor.

When you say redundant here, do you mean in the sense of offering more description than is necessary to pin down a unique phenomenon? I don't think I would interpret Occam's

Razor that way. If I remember right, the saying is that we mustn't multiply entities NEEDLESSLY. In my opinion, we need enough entities to make all the distinctions that matter -- all the distinctions neccessary to allow building a model that reproduces all the phenomena of interest under all circumstances of interest. I don't see how a statistical description can do anything but discard details. Enlighten me, please.

>Really, discreteness and continuousness are irrelevant when considering >redundancy.

They aren't irrelevant when one is trying to build a model that deals in continuous variables at one level, and discrete ones at other levels. So are you saying that redundancy isn't relevant to such a model?

>Point of view is the important thing: what do you know about how the >observations COULD turn out if you knew nothing of their structure; what >do you know about how they could turn out given what you know of their >structure, and what do you know after you have made the observation.

You're speaking of the analyzer of the system: "you" is not the system itself. What YOU know has no effect on how the SYSTEM works. A squarerecognizer doesn't have to know it could have been looking at a triangle or a chocolate ice-cream-cone.

Somehow I feel that your arguments are getting so general that they don't really deal with the problem I'm interested in, which is how the system is put together to work as it does. Everything you say may be true, but it doesn't help with the design. Or if it does, I don't see how.

>I think you have to use a subjective probability view, based on >expectations about observations (which may well derive from past >frequencies of event types).

This would make a shambles of my hierarchy, because it puts into the lowest levels of perceptual organization functions I reserve only for the highest ones. There's something going badly amiss here between our points of view. Maybe it's the virus. ----------------

----

Wayne Hershberger (911224) --
Yeah, Boss Reality says Merry Christmas back. Really.
>And I say, as I did back in July (Hershberger 910723) that
>ablation (to which you refer above: "Remove those signals"), the >technique pioneered by Pierre Flouren's in the 17th century to >localize mental functions identifies certain NECESSARY components >of the various functions we call mental (e.g., vision). It does >NOT identify the NECESSARY AND SUFFICIENT components. Without the >photon there is no vision--for anyone. Ablate photons and we are >all blind.

What are these mythical "photons" of which you speak? I don't know anyone who has the ability to "ablate photons." We can perform various acts, like shutting our eyes or pulling the chain on a light, that result in loss of vision, but to attribute that loss of vision to a loss of
"photons" goes far beyond anything that is observable.

By accepting "photons" as a necessary precursor to vision, you are leaping ahead to the conclusion you want to reach; namely, that photons actually exist just as we imagine them to exist. But I can't accept that mode of argument: I want to know the operational basis for every critical entity you use in your proofs. I will accept that turning off the lights results in loss of vision. Those are both observables, perceptions. I do not accept that you have shown photons either to exist or to have anything to do with this phenomenon -- not until you tell me your basis for knowing that.

>This means that the PROPRIETARY aspect of our
>respective experience (my perceptions versus your perceptions) are >contingent upon our respective brains. That is the argument you
>are making, right? But that does NOT put us IN our respective
>brains! Our feet are too big.

Cute comment but irrelevant. Our perceptions that we call "feet" are certainly not too big to fit into a brain: they are precisely small enough to pass through a neural fiber. All aspects of our perceptions are proprietary, including our convictions that some are not. If that were not true you would have convinced me by now. But you have nothing objective to show me to help make your case.

You ask,

>... what, if anything is wrong with the following remark (using your >terminology): Bill Power's HCT model models an aspect of boss reality >(the organism aspect), the other aspect (the environment aspect) already >having been well modelled by contemporary physics.

Sensing a bear-trap, I answer cautiously. Both the HCT model and the physics model purport to represent aspects of a boss reality. Both are tested (by a person, using a brain and body) by assuming the model to be correct, and predicting the effects of actions on this boss reality that have consequences we can perceive. It is the boss reality that determines whether our predictions work out as we expect, or whether different consequences occur. If the consequences are different, we modify our imagined pictures of the boss reality in a direction that promises to lessen the difference. This process converges to some minimum-error condition where we declare ourselves satisfied with the models. According to both the physics model and the HCT model, this process of acting and testing takes place inside a brain. It is not necessary to assume that the model in the brain has any particular correspondence to the boss reality. It is necessary only to assume that whatever that correspondence may be, it is stable over time.

I ask a similar question of you: is it fair to say that WH believes
(a) that there are NON-proprietary aspects of our respective experiences and (b) that we can say unequivocally what they are?

>Now that you're in Colorado, I trust your "kids" will be joining
>you and Mary for Christmas. Have a merry one.

We saw two in Boulder for Denny's graduation, and the third arrives with grandchild genius elf in an hour or so for official Christmas such as it is in this deplorably secular household. Merriness, however, seems to be nonsectarian. HO ho to you and yours.

Best to all,
And most heartfelt thanks to Gorby, who ended it today, for changing history and I think saving the world from disaster.
P.S. Doesn't anyone remember that it was Gorbachev who, in his U.N.
speech a year or so ago, called for a "new world order?"
Bill P. ===

Date: Thu, 26 Dec 1991 13:00:40 CDT
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: Re: Kolbe

From Tom Bourbon [911226 - 12:46]
Back from travels. Wayne Hershberger [911226] described Kolbe's
book on _The Conative Connection_. Wayne identified some nice quotations, very much in line with ideas in PCT. Then Rick Marken [911219] came back with a remark that he had seen the chaff in Kolbe's book, where Wayne had seen the wheat. I guess I had sen both -- some interesting and promising similarities, set in a context that reflects the kinds of loose talk people must use when trying to market self-improvement indices and methods. Kolbe's emphasis on the importance of establishing what PCT calls reference signals, then letting the behavior take care of itself seems sound.
But the reference signals of greatest iimportance inn her presentation fall into four major categories -- shades of William Glasser, and scads of other self-help, self-improvement marketeers. And people also conveniently sort out into a few "types," in terms of "behaviors."
    Close, but oh so far away.

Wayne, your remark that David need not forward something to you seems to imply that we all received a copy of the letter Kolbe sent
to you -- I will not forward mine.
Best wishes,

Tom Bourbon <TBourbon@SFAustin.BitNet>
Dept. of Psychology
Stephen F. Austin State Univ.
Nacogdoches, TX 75962 Ph. (409)568-4402
==
Date: Thu, 26 Dec 1991 13:57:01 CDT
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: M. Taylor and Uncertainty

From Tom Bourbon [911226 - 13:43]
For some time now, a discussion has centered on the topic
of information theory and its relevance, or lack thereof, for work in PCT. Martin Taylor is the major advocate for the putative importance of information concepts in PCT. (Many posts, but see as examples [911213] and [911216].) At a very abstract level, I see some relevance for information-theoretic (IT) concepts in
general discussions of PCT, but, for any problems modeled to date, I see no role for IT concepts.
Martin, could you help some of us better appreciate your position by picking any published example of quantitative
modeling with PCT, then point out specifically how you would
use IT concepts to either duplicate, or improve on, the results published for the IT-free versions of the PCT model? I am open to any quantitative demonstration that IT concepts (information, uncertainty [actual or maximum], redundancy, channel capacity, or any others) duplicate, or improve on, the results of garden variety PCT modeling. Absent such quantitative demonstrations, I
am lost, trying to see the relevance of information theory at the level of simulating working models of control systems.
Best wishes,

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

Date: Thu, 26 Dec 1991 12:53:51 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: left handedness

[Martin Taylor 911226 12:45]
(Bill Powers 911225)

>Martin Taylor (911221) --
>
>Handedness wouldn't seem to come out of control theory as an a priori >principle: it's just a phenomenon. If you practice doing everything with >your right hand, you'll probably be right-handed. If there's a bias in
>the nervous system (or in the environment) that makes learning things >left-handed easier, you'll be left-handed. I don't think that handedness >per se exists; it's just that something may exist that leads to using one >hand more than the other. When we find it (assume we bother to look) we >will then be able to explain handedness and probably a few other things. >If it weren't for scissors and words like "sinister" I doubt that >handedness would count as a very important phenomenon.

But handedness seems to be more built-in than that. If it were just a matter of training, there should be a better balance between right-handed and left-handed people,

and I doubt that there should be so few ambidextrous people. Yes, for sure training helps, but it seems to be much the same
as the training that allows infants with severe cortical damage to compensate by using undamaged parts of the brain. They learn, but not as well as they otherwise might, and the deficits relate to the normal function of the damaged part of the brain, not to the general overloading of work onto a small brain part.

If evolution has made handedness such a strong factor in humans (and other primates and a few other species), it must be more important theoretically than you suggest. I guess the difference between our views is that you don't think handedness exists except as a training phenomenon. But even if that were true, I would have thought naive PCT would have said that there should be only random imbalances between the hands, rather than a usually global
one that has a strong correlation with the side of the brain that deals
with language.

Martin Taylor ==

Date: Thu, 26 Dec 1991 12:44:27 EST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>
Subject: Re: Syntax;probability & redundancy; boss Kringle

[Martin Taylor 911226 12:00]
(Bill Powers 911225 Xmas Day!)

>There's something going badly amiss here between our points
>of view. Maybe it's the virus.

I think it must be. You were happy enough with my description of redundancy before your happy yule trip. In this posting you seem to have lost it again.

On sampling: The sampling rate statement has nothing to do with whether sampling is actually done. It says that if one samples at a certain rate
one has enough information to reconstitute the original continuous signal
exactly. Any further samples give no more information. They can be considered as a redundant point of view on a wider bandwidth channel (the actual signal
is known to be limited within bandwidth B, but oversampling would allow you
to exactly specify a signal of wider bandwidth W. B is redundant from the viewpoint of W).

On generality: I am indeed trying to be general, and pointing out what MUST
be true about the information available to the elementary control systems
in the hierarchy. I do not assert that the constraints imposed by
considerations
of information actually are the limiting constraints, but if there are limiting constraints derived from other considerations (kinematics, for example), they must lie within the constraints imposed by information rates.

On modelling and Occam's razor: Occam may have said that entities should not be needlessly multiplied, but a more modern formulation can be based on Kolmogorov complexity, which is what I was doing. The issue comes down to
how much information beyond the <model|description|theory> you have to supply to describe the specific instance, and how wide a range of specific instances the
<model|theory|description> purports to describe with how large a description. Naturally, the description is redundant, and the best
encryption of it, removing the redundancy, describes its Kolmogorov complexity. The shorter the better. So, Occam's razor can be restated as asserting that the simplest, widest ranging, and most accurate description is the best
(and those concepts being measured in the same informational terms, they
are commensurate and can be traded). In this view, there is no qualitative difference between a statistical description, a model, or a theory. The distinction is quantitative, and the quantitative difference is usually large. Models tend to be simpler and more accurate than statistical descriptions, and theories tend to be more wide-ranging and simpler than models, with no particular bias as to accuracy. I am

applying information theory to your model. It carries no implications as to the necessary form of your model, but there shoudl be compatibility.

On continuity and discreteness: There is no difference in principle between a probability for a discrete event and a probability distribution over a continuous set of possible events. So there is no forcing in using the concepts in the continuous domain.

On subjective probability: I do not intend to refer to an EXTERNAL observer's subjective probability. There can be no such thing, except as part of a model of that external observer. Only the one exposed to the situation can have
a probability estimate. Also, I do not require that there be any consciousness associated with the subjective probability. There is a very difficult issue
of words here, and to pursue it would lead into much the same argument you are having with Wayne on boss reality. We may have to do it, but I think
it is a side-track to the main object of enquiry, the properties of PCT. For that, it matters little whether the probabilities concerned are subjective
or objective, provided that they exist. The fact that one does not *observe* a probability distribution in a percept does not mean that it is not there. It means that there are differences in input that do not cause observable differences in behaviour. That relates to, but is not the same as
categorization,
and may be the reason you and Bruce were having an argument about the place of the category level.

>This would make a shambles of my hierarchy, because it puts into the
>lowest levels of perceptual organization functions I reserve only for the >highest ones.

I can see how you would think that, given the word "subjective". All it means is "from the viewpoint of the place where the probability is used". It does not imply a function, so much as a capability or perhaps a predisposition.
To the retinal cone, there is a certain probability that a photon will cause a particular chemical event, and that probability varies both over time and over the wavelength of the photon. The retinal cone does not have to be
manipulating anything for this to be so.
Happy Boxing Day. Why do you work on Christmas? Especially with the flu!
Oh, well. I am under the influence of a mild flu, kept in check by a suitable pill, so maybe all the above makes no sense, either.
Here's hoping it does make some sense.
Best wishes for a convergent New Year.
Martin Taylor
Please forgive intrusive or altered characters. The phone line from home is inserting them, and I don't know which are in the text that goes out.
===

Date: Fri, 27 Dec 1991 15:47:13 EST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: goldstein@SATURN.GLASSBORO.EDU
Subject: Kolbe

To: CSGnet people
From: David Goldtein
Subject: Kolbe
Date: 12/27/91

Wayne, Rick and Tom have given their impressions of Kolbe. I have also received her book and will, in the near future, comment on it. For now, I can say that I like what she has done and think that it has a lot of possibilities.

Right now, I wanted to pick up on some comments by Tom. He noticed that Kolbe has four individual difference variables in her classification system. He spoke of the varibles as types. One point is that each person is described in terms of how much of each variable characterizes the person. In other words, a person does not fall into one of four categories as Tom's remarks seem to imply.

A second point relates to the issue of modeling. I have been thinking about how to model the self-image control system(s) in a person. [Bill has advised me to wait until I have applied the approach I am taking to more people before trying to create a model.] All I want to say is that from what I have learned from a self-image study, I can see how it is possible to take the variables of Kolbe, which I think are talking about systems level perceptions, and incorporate them in a modeling approach.

In summary, I think we should stay open-minded about Kolbe's work and should not let any negative preconceptions about psychometric approaches in general, for example Rick's reactions, or negative preconceptions about nonmodeling approaches, for example Tom's reactions, stop us from mutually exploring our common interests which, as Wayne has pointed out, seem to be promising.
===

Date: Sat, 28 Dec 1991 15:12:29 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: Re: Syntax;probability & redundancy; boss Kringle

Bill,
Good to hear from you - we just got back from a genuine aussie Xmas at the beach ...

It's going to take me a while to sift through the thru the net stuff.
Something I'm sort of interested in doing is picking thru current `philosophy of mind' (Fodor, Dennet, Sterelny, etc.), & seeing at what points their ignorance of control systems messes them up. I'm on the whole sympathetic to most of what they do, but there's definitely stuff they don't understand.

Cheers, more later, love to all

Avery ===

Date: Sat, 28 Dec 1991 13:44:53 -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: Martin's ideas

[From Bill Powers (911228.1000)]

Martin Taylor (911227) --

> If evolution has made handedness such a strong factor in humans (and
>other primates and a few other species), it must be more important >theoretically than you suggest. I guess the difference between our
>views is that you don't think handedness exists except as a training >phenomenon.

I meant that handedness itself is just a trait, not something to be used in an explanation. To the extent that there's an assymetry in learning skills with systems on the left and the right, the explanation isn't to be found in the trait, but in an assymetry of organization.

I guess I'm not very interested in handedness because it isn't a nice clear-cut phenomenon. There's very little one does with the dominant hand that can't be done, with practice, with the other hand. The decussations are not absolute so that ALL signals to and from the right side connect to the brain's left side. Lots of people do some things left-handed and others right-handed -- there doesn't seem to be any rigid restriction. When I was working on modeling tracking, I did most of the runs lefthanded because putting the joystick on the left kept the wires from draping over the keyboard. Later I switched to a stick on the right, and was soon showing essentially the same behavioral parameters. There weren't any interesting differences. Of course I haven't surveyed any populations.

As to "the side of the brain that deals with language," I'm not so sure there is such a thing -- depends on what you decide to call "language," and also on what you accept as a "fact." Practically all the left-brainright-brain stuff I've seen has been statistically derived, meaning that slight preponderances of effects on one side are customarily elevated to absolute dichotomies. Tom Bourbon pointed out that in brain-activity maps, presentations of results use a lot of suppression of readings below an arbitrary threshold -- it appears that there's a clear-cut boundary, but that's just showing where the threshold is set. The whole brain is active all the time, with some parts being a LITTLE more active than others. I don't like data of this kind. It's like looking at the noise and ignoring the signal. The effects I'm interested in are big and obvious.

Redundancy, sampling, etc.

>>There's something going badly amiss here between our points
>>of view. Maybe it's the virus.

>I think it must be. You were happy enough with my description of >redundancy before your happy yule trip. In this posting you seem to >have lost it again.

I'm still reasonable happy with your description of redundancy, to the extent that I have a clear picture of it. What's puzzling me is how to use this concept in the HCT model. I have a lot of questions about the interface between our ideas.

>On sampling: The sampling rate statement has nothing to do with whether >sampling is actually done. It says that if one samples at a certain >rate one has enough information to reconstitute the original continuous >signal exactly.

How would this apply to the operation of a neuron? Suppose that the output frequency is created by a relaxation oscillator whose firing threshold is set by a smoothly-varying chemical concentration at the axon hillock. We thus have a relationship between concentration and impulse rate. At the receiving end of the signal generated by this oscillator, we have the complementary operation: a series of chemical jolts results in a concentration that is the superposition and smoothing of a train of exponentially-decaying spikes of concentration. As a result, an output train of spikes becomes a function of an input train of spikes, one frequency being converted into another without any necessary synchronization of output spikes with input spikes.

Alternatively, we could consider the chemical concentration inside neurons as the "continuous signal", so that one chemical concentration is made to be a function of a remote chemical concentration with the oscillator output and synapse to the next neuron being the mediator. Now the neural signal is somewhat analogous to "sampling," although it is the "sampling" rate itself that carries the information.

In either case, in what sense is there an equivalent sampling system that, sampling at a given rate, yields "enough information to reconstitute the original continuous signal exactly?"

>On generality: I am indeed trying to be general, and pointing out what >MUST be true about the information available to the elementary control >systems in the hierarchy. I do not assert that the constraints imposed >by considerations of information actually are the limiting constraints, >but if there are limiting constraints derived from other considerations >(kinematics, for example), they must lie within the constraints imposed >by information rates.

Once we understand the computations by which specific perceptions at one level are derived from specific perceptions at lower levels, calculations of the relative information content at the two levels may well prove to satisfy the informational constraints at least as upper bounds. So I'm not casting doubt on your claims concerning information content. I'm just wondering what the constraints we can calculate from information theory (in principle) can do toward pointing us to better forms of the model. It seems to me that information considerations must be inherently ad hoc; one can't tell what measures of signals to use until the perceptual (and other) computations are known, and so can't know what equivalent virtual sampling rates would be appropriate.

>On modelling and Occam's razor: Occam may have said that entities >should not be needlessly multiplied, but a more modern formulation can >be based on Kolmogorov complexity, which is what I was doing. The issue >comes down to how much information beyond the <model|description|theory> >you have to supply to describe the specific instance, and how wide a >range of specific instances the <model|theory|description> purports to >describe with how large a description.

Sounds very imposing. It would help me if you could say what constitutes a "specific instance." For example, suppose we are trying to come up with a <model|description|theory> of balancing behavior (i.e., standing up). Modeling simple balancing under small perturbations might involve nothing more than sensing pressure distributions on the bottom of the feet and
adjusting torques at the ankles to maintain the center of pressure within the footprint. A successful control-system model would be able to maintain balance indefinitely under any arbitrary force perturbations, changes of muscle response, and small tilts of the floor, within the outer limits of control. How would we determine the Kolmogorov complexity of a control-system model of this behavior in comparison to that of a treatment in terms of redundancy or statistical perception?

>On continuity and discreteness: There is no difference in principle >between a probability for a discrete event and a probability >distribution over a continuous set of possible events. So there is no >forcing in using the concepts in the continuous domain.

"A continuous set of possible events" sounds like an oxymoron to me. Suppose we are talking about control of body temperature. The temperature variable, as I would model it, is not an event but a variable. A steady temperature is just as possible as a varying one. Variations do not take place as little jumps, but simply as smooth changes in a measure on the real number scale. Similarly for a neural signal representing temperature. The frequency of impulses is continuously variable, not quantized; frequency can change by any amount including zero. When you compute probabilities in a case like this, do you have to go to the level of single impulses? And how do you treat the state of the temperature variable, or other variables such as distance, angle, or velocity, which are continuous without entailing any events? Are you, as an external observer, arbitrarily marking off a continuous process into packages in order to create events?

>On subjective probability: I do not intend to refer to an EXTERNAL >observer's subjective probability. There can be no such thing, except >as part of a model of that external observer. Only the one exposed to >the situation can have a probability estimate. Also, I do not require >that there be any consciousness associated with the subjective >probability.

"The one exposed to the situation" is ambiguous in terms of HCT. Do you mean the whole person, using all levels of intellectual capacity, or each subsystem using only the computations that define it?

If you mean each subsystem, how does subjective probability enter into the operation of an input function that (for example) weights a set of intensity signals and sums them to form a chocolate-signal? Does this input function do the probability estimates itself? Are these estimates expressed in some form other than the chocolate-signal itself? Is the probability calculation something other than the weighted summation? And just what is it that is "probable?"

Enough questions!

>Why do you work on Christmas? Especially with the flu!

Was that "working?" Oh, dear. Actually I just wandered downstairs and turned on the computer and played around a little. Please don't call that working. I'm not supposed to be working any more.

Best,

Bill P. ==
Date: Sun, 29 Dec 1991 12:52:47 EST

Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: syntax

>>Rick says:
>> The sign said "Please do not throw paper or cigarettes in the >>urinal". The sign interested me as an example of the precedence (in this >>case) of meaning over grammar in language. I immediately understood the >>sign to mean "don't through this stuff into the toilet". It took me
>>a while to realize that a proper grammatical reading would lead me >to think that it was ok to throw things into the unrinal as long as I >>was not standing in the urinal at the time.

>I would have interpreted the notice as suggesting that if you find stuff in the >toilet, you shouldn't throw it.

Rick's observation is wrong in detail, although fine in what I take to be its ultimate point. His first parsing of the sentence treated the prepositional phrase (PP) as a directional complement, designating the path travelled by some entity involved in situation described by the sentence. Then, for some reason, he thought that it ought to be taken
as a locative adjunct, indicating where there event or situation described occurs. The reason is perhaps the use of `in' rather than `into' as the preposition, which is possible from some verbs but not others (c.f.
Dan sent the spies in/into Russia) - I don't think anyone knows what determines the possibilities. The third interpretation has the PP as a modifier of the object noun-phrase. All three are grammatically possible, although only the first one is pragmatically plausible.

The problem of `PP attachment ambiguity' is a standard issue in computational linguistics, and shows clearly that a workable parser can't produce full conventional grammatical structures and then present them for pragmatic evaluation--a system that tried to work that way would get buried in trash. The viable possibilities are

a) grammatical and pragmatic processing are interleaved in some manner.

b) grammatical processing comes first, but produces an `underspecified' structure that is vague or ambiguous about certain matters, such as PP attachment. `Common sense' then chooses the most appropriate from the grammatically available possibilities.

I'd also suggest that it's at least as likely that grammatical processing is tried first, with common sense as the backstop. The reason is that grammatical processing is basically pretty easy--even crummy little
microcomputers can be programmed to do it, but common sense is quite another matter.

Avery Andrews (Avery.Andrews@anu.edu.au)
==
Date: Sat, 28 Dec 1991 23:07:02 EST
Reply-To: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender: "Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET> From: goldstein@SATURN.GLASSBORO.EDU
Subject: more on Kolbe

To: CSGnet people
From: David Goldstein
Subject: Kolbe
Date: 12/28/91

I have had a chance to read "The Conative Connection" by Kathy Kolbe. Recall that she sent several of us this book and expressed interest in the CSG approach. Ed Ford has a meeting set up with her in the early part of January.

1. She has developed a psychological pencil-and-paper test(Kolbe Conative Index or KCI) which claims to: tell you what a person, left to his own devices, will do. The test is 36 questions long.
Here is one item:
36. If free to be myself, I would get things done by:
researching (fact finder)
planning ahead (follow through) hard work (implementor) challenges (quick start)

The test taker is supposed to indicate which choice is most like him, and which choice is least like him.

Each choice indicates one of four possible modes of action: fact finder, follow through, quick start, and implementor.

Kolbe does not provide a scoring key, but I have put in parethesis, what I think each choice indicates.

The meaning of each "action mode," and the link to Perceptual Control Theory, is suggested by the following quote:

"Achieving a goal requires making it a priority (fact finder), focusing energy on it (follow through), getting to its bottom line (quick start), and making it happen in tangible terms (implementor)."

The test taker receives a score (1 to 10) on each variable. Scores of 1,2,or 3 mean a person resists activities involving that action mode. Scores of 8,9 or 10 mean the person insists on doing activities in that action mode. And scores 4,5,6, or 7 means that a person accommodates when doing activities in that action mode.

Example: A person's "MO" may be--2498. This means resistence in the fact finder mode, accomodation in the follow through mode, insistence in the quick start mode and insistence in the implementor mode.

2. Kolbe states that the KCI measure is unrelated to measures of intelligence and personality. She does not report the research in her book but mentions in her letter that she has lots of cases, 30,000 people, and is in the process of writing a second book supporting this assertion as well as others she makes in the book.

4. Mainly, Kolbe uses the KCI to match a person with a job. She has a second test, which measures the functional requirements of a job in terms of the four variables mentioned above.

Mismatches between a person and a job result in stress for the person and business organization.

5. Some thoughts I have had about relating Kolbe to PCT are:

(a) The Kolbe variables which a person insists on or resists are controlled perceptions(experiences), probably at the principle level of perceiving.

(b) From the self-image study which I have carried out recently, which resulted in three self-images, I disagree with Kolbe that the variables which she has identified are unrelated to selfimage (and therefore, personality). I can see how one self is quick start/follow through, one self is fact finder, one self is implementor. This makes me question her ideas about these variables measuring constant qualities about a person.

(c) In short, I think that the Kolbe variables are aspects of the self-image. As I outline in the paper about the self-image study, I think that a modeling approach of the self-image control system is possible.
===

Date: Sun, 29 Dec 1991 07:30:47 -0600
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>

Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Robert K. Clark by way of Gary A. Cziko g-cziko@uiuc.edu"
<0004912499@MCIMAIL.COM>
Subject: LOOPS 3 & 4

[from Robert K. Clark 911228]

REACTIONS TO & COMMENTS ON CLOSED LOOP #3 & #4

Receiving #3 led me to realize that I really should get back in touch with the current state of the work Powers and I developed together some 30 years ago. That I am at all aware of what's going on is due to Greg Williams,
whom I thank very much for his thoughtfulness and consideration. In 1987 I had a brief correspondence with Bill, but it went nowhere.

As you might expect, I have been using, applying, and developing these early concepts on my own during this period. My contacts with Greg have led me increase my activities along these lines.

In beginning to contact the CSGNET I offer some comments on Loops #3 & #4.

My main comment on Loop #3 is the lack of discussion of the hierarchical levels that could/would be involved! Instead, much of the discussion recapped the standard arguments for/against governmental control -especially of the economy. It seems to me that a more precise and accurate definition of the original Hierarchical Orders would be very helpful.
Since I find re-examination and refinement of these concepts continue to be interesting and useful in my own life, I plan to offer some of these thoughts to the Net from time to time.

#4 seems to me to go around and around because of insufficient recognition of the role of the engineer in designing and operating his system (whether it be "open loop," or "closed loop,"). Indeed many discussions omit the critical role of the observer, the experimenter, engineer, etc. In fact I once wrote a paper on a related subject: A systems view of psychophysiological experimentation (with McFarland and Bassan, "Integrated Data Collecting and Processing Systems in Psychophysiology" presented at the New York Academy of Sciences in 1964). Not a very good paper, but points out some of the levels of interaction normally omitted from discussion.

The key in #4 seems to me to lie in the phrase -- mentioned several times in the discussion -- "point of view." The professional engineer takes his own role for granted: it is not part of the system he is designing. In his design work he has, as was noted in the discussion, full access to all aspects of his work and hence can use the terminology as he sees fit. But the moment the engineer himself is included, most of his hierarchical structure is in action!

More of my present views will be forth-coming as time permits. I am now (retired for some years) and very busy. I am the Secretary, Treasurer, Editor for a Museum (a bit over 2 years old) interested in Rotorcraft. Since I knew how to work with the IRS, we became a tax exempt public foundation.

The National Popular Rotorcraft Ass'n held its Convention nearby last summer, and plans to return in '92. I was Financial Chairman in '91 and am again in '92.

But I retain my original interest in Human Control System Theory and we shall see how this all works out.

My regards to you all -- Bob Clark
\
===
Date: Sun, 29 Dec 1991 07:50:48 -0600
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Philosophy of Mind

[from Gary Cziko 911229]

Avery Andrews (911228) says:

>Something I'm sort of interested in doing is picking thru current `philosophy >of mind' (Fodor, Dennet, Sterelny, etc.), & seeing at what points their >ignorance of control systems messes them up. I'm on the whole sympathetic >to most of what they do, but there's definitely stuff they don't understand.

I think you will much less sympathetic to Dennet and Fodor when you see them through Perceptual Control Theory eyes.

I have been recently reading Dennett in his _Intentional Stance_ as well as some Fodor and even though I'm sure that they are a hell of a lot smarter than I am, I can see some very real problems with their stances. Dennett seems to have absolutely no understanding of how purpose or intention can play a role in a living organism. He says essentially that we can pretend for now that humans have intentions to make sense out of what they do but that ultimately as our understanding advances we will able to eliminate the notion of intention from human behavior!

A very nice critique of Fodor's positions is provided by Mark Bickhard, particularly in his book with Richie on _The Nature of Representations_ where he critiques Fodor's critique of Gibson theory of perception (he critiques Gibson as well). Bickhard's interactivist view of representational is nicely compatible with that of PCT and so the critiques he makes of Fodor are also the kinds of critiques that PCT would make. Essentially, it is an argument against what Bickhard calls an "encodingist" view of mental representation and is very much the same argument that Bill Powers has made a few times on the net. The trouble is that while Bickhard makes what I feel are important arguments against the mainstream "encodingist" view of mental representation, he doesn't seem to understand PCT either although he is familiar with it.

Maybe the best way to get some discussion going here would be for you to state in what ways you are sympathetic to the views of people like Dennett and Fodor (indeed, I see Dennet and Fodor as having quite different views where mind is concerned).--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: Sun, 29 Dec 1991 23:17:40 -0600
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Sentence Meaning

[from Gary Cziko 911229.2315]

Just a few quick examples to follow up on recent comments by Powers and Taylor on the role of syntax in the meaning of sentences.

1a. My husband likes football, but I couldn't care less about it.
1b. My husband likes football, but I could care less about it.

Funny how "I couldn't care less" and "I could care less" mean the same thing, even though one has a "not" in it. Compare "I like football" and "I do not like football." Quite different meanings there.

2a. No head injury is too trivial to ignore.
2b. No head injury it too trivial to treat.

Again, "ignore" and "treat" are close to having opposite meanings, yet in these sentences they mean quite the same thing--take care of all head injuries. (Although I bet that one of the linguists on the net can give syntactic reasons for the convergence in meaning here.)

So it seems that we often make sentences mean what we want them to mean, regardless of their formal structure.--Gary

----

Gary A. Cziko
Educational Psychology   Telephone: (217) 333-4382
University of Illinois   Internet: g-cziko@uiuc.edu 1310 S. Sixth Street
Bitnet: cziko@uiucvmd
210 Education Building   N9MJZ
Champaign, Illinois 61820-6990
USA

==

Date: Sun, 29 Dec 1991 21:37:50 PST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: marken@AERO.ORG
Subject: Language, Kolbe

[From Rick Marken (911229)]

The ol' computer has been down for some time. Good thing the net is slowing down. I just wanted to thank Avery Andrews (whomever you are) for that wonderfully hilarious reply to my urinal sign post. Also a big thanks to Gary Cziko for bailing out my intuitions about the primacy of imagery over (ug) grammar. The "could/couldn't care less" example was wonderful.

Re Kolbe: David, I looked at the book again and I think you should go with my intuition -- forget it (unless, of course, you want to make a bundle; but I think that s--t will be more difficult to sell in the 90's. But who knows.) I think she needs to learn a bit more about variable means (to counter disturbances) -- she seems to have intended results down pat.

Hasta manana

Rick

==

Date: Mon, 30 Dec 1991 22:21:17 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: Re: Sentence Meaning

I think the ultimate in polarity reversals are the Eastern Massachussetts dialects where people say things like:

these gloves cost $10 and so don't the others.

Avery Andrews ==

Date: Mon, 30 Dec 1991 13:38:00 GMT
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Hortideas Publishing <0004972767@MCIMAIL.COM>
Subject: WELCOME, BOB

From Greg Williams

I, for one, am happy that Bob Clark has joined the Net. His experience of several years of ruminations regarding living control systems -- in parallel and independent of Bill Powers' ruminations -- should enrich the dialog (two eyes are better than one, and all that). Go to it, Bob!

And since MCI somewhat mysteriously credits me with $100 from out of the blue, maybe I will have a bit more to say than in the past. BTM, I'll be lis[listening -- darn these line editors!], so don't make any mistakes, folks.

Happy New Year to all....
==

Date: Mon, 30 Dec 1991 11:39:37 -0600
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: semantics

[from Joel Judd]

Bill, Bruce, Martin, Rick, Tom, other theorists and philosophers:

I have been resisting the temptation to get too caught up in the recent linguistic/syntax discussion in order to try and work out a general position on language on which I can center what is apparently going to be my future in academe. I think I have finally found (after already writing 100+ pages) a fairly succint argument, which captures the glowing generalities I love to use, as well as the line of reasoning sure to conflict with the greatest number of people in SLA. But there seems to be much here beyond SLA and into psychology in general, which is also why I like it, and would like to share it with the net (in E-flat, Tony--not too fast...). Really, the ideas are nothing novel, but it's the combination that finally seemed to make sense, as part of a theoretical "package."

This line of thought started a couple of months ago with my reading of several respected researchers in SLA who have been trying to describe what theory should look like. By looking over the last thirty years, they saw two "kinds" of theories (using the term loosely)--those which purported to describe/explain processes (eg. how one learns negation in Japanese); and those which described states, or properties (eg. government and binding rules). A few people, most recently and notably Kevin Gregg at St. Andrews University in Japan, have tried to make sense of such a distinction, arguing for the need for both kinds (though sort of saying that a field like Second Language _Acquisition_ would find _process_ theories more interesting, as they are the theories that explain change), but also arguing for their fundamental incommensurability (incompatibility?). This latter point has stuck in my mind, and I have wondered why, if one is trying to develop an overall picture of the dynamic nature of human behavior, one would care to make, as well as steadfastly maintain, such a distinction. This is why I asked some time back about the way we sometimes talk about "knowledge" as a sort of static thing that, given some learning, changes to another state.

Enter at this point Gary's mention of Mark Bickhard and an in-press article which led to looking at a couple of older books also by him, _Cognition, Convention and Communication_ (1980); and _Knowing Levels and Developmental Stages_ (Campbell and Bickhard, 1986). Gary has mentioned some of his arguments against traditional encoding theories of knowledge. I have enjoyed his desire to clearly define terms often used in science, especially in social science, and that's why I'd like to get reactions from specialists in other branches of psychology and sociology etc. to the following summary and argument.
Bickhard begins by pointing out how the history of language study can be characterized as "prescriptivist;" that is, linguistics was as often as not the description of language as it "ought to be," not how real speakers actually used it from day to day. Over time, such practice served to "[idealize] away from the actual variations and processes of that which is' (1980, p.3). He argues that even since Chomsky much of this prescriptive attitude has remained, even though linguists have claimed to have dropped the idealizations. However, they still talk of a "speaker-hearer," which seems to be another form of idealization based on their descriptions of language. So there is danger of falling into a circular definition of language: language is what linguistics decribes it to be (1980, p.3). I think language learning fields have borrowed the same concept in the notion of a "native speaker" of an L2.

As the actual language which served as a basis for description became obscured, the possibility that description could be confused with explanation became real. Again, I think this danger has become especially relevant in Lang. Learning fields, where descriptions of what Bickhard labels "task capabilities" have often been used as explanations of the learning process. This is one of Campbell and Bickhard's strongest arguments against both Piaget and Chomsky in the 1986 book--the description of a system's capabilities at a particular point in development, even if the ability can be demonstrated species-wide, does not tell us how that ability got there in the first place, or why the system needs it or uses it.

By using examples from two of this half-century's most notable human scientists, Bickhard makes a good case for the distinction between description and explanation. And his reason for doing so is where I think CT and interactivism should find an understanding: because internal (unobserved) processes are responsible for development and behavior. Many of his arguments for understanding behavior from an internal perspective parallel those given in defense of CT; additionally, he offers some ontogentic explanations for developmental issues regularly raised on the net. His point of view on understanding human behavior is stated as follows:

"...(1) we can attempt to describe [a] system's class of interactions [with the environment] from an external perspective; and (2) we can attempt to explain those interactions in terms of processes of the system." (1980, p.103)

What is required for the second perspective of course is a model of the underlying processes, something which CT offers. The model needs to be testable and falsifiable, and that is where most, if not all, of SLA's models go down the tubes. They are virtually all so general as to be empirically useless for _explanation_. What they capture quite well in some cases are descriptions of important aspects of language learning: Yes, we can often catch ourselves before/after an utterance and correct the verb form used or the vocabulary item or the subject verb agreement, etc.(The Monitor Model). Yes, our attitudes towards Japanese will play a role in how we learn Japanese (The Acculturation Model). And on and on. But where is the underlying psychological process? For example, the Monitor Model offers an "affective filter." When it's "up," we learn laboriously and non-naturally. When it's "down," we learn more naturally. What does that mean?

In order to provide an adequate _explanation_ of linguistic processes we must, to paraphrase Bickhard, know what such processes ARE, and to know that, we must explore the ontology of those processes. Never mind the fact that we cannot observe them directly, we can still propose testable models of the processes (which is what CT claims as well). What we then observe in testing our models serves to modify or change them, so there is a dialectic between description (of a test) and explanation. To quote again from Bickhard:

"_Explanatory_ accounts of underlying processes and their properties are constrained by _descriptive_ accounts of the manifestations and potentials of these processes. That is, accounts of underlying processes that produce effects under appropriate conditions (uh...) are constrained by accounts of the range of possible effects. In the case of psychological processes, which are unobservable, descriptive accounts typically take the form of characterizing task performances. These descriptions have theoretical force because they cover a potentially infinite set of possible task performances, and they are falsifiable. Such descriptions of task capacity restrict explanatory theory because the explanatory theory must account for the described capacity; explanations that cannot yield the capacity have to be rejected as inadequate. However, descriptive capacity accounts are not explanations. They do not specify a generative mechanism, or an underlying process, that could produce the task accomplishments described, or account for their lawfulness." (1980, p.15)

Getting back to the origins of my interest in this whole mess, I understand from the above quote that what is usually of interest in process accounts--CAUSALITY--cannot even be approached without some explanation of the process. Recent SLA discussion alluded to above deals with this by saying a property theory is no good without a property one. But again I ask, can we usefully maintain the distinction in the first place? I think that doing so betrays (1) a wish to pursue a causal model of behavior and (2) confusion as to what is _description_ and what is _explanation_. A property account, it seems to me, can only be an account of an idealized property (or properties) of a system, NOT an explanation of the processes which led to the emergence of the property, or an explanation of why it is used by the system--there is no ontological account of the processes, hence no "property thoery"--AND (and this is where I think everything came full circle, and I saw "Sylvania" over my head) because statistical measurement of the observed property(ies) obscures any individual understanding of the processes. There seems therefore to be no way to successfully maintain a traditional scientific perspective and hold out the hope that it will explain human behavior--ideas of causality, observation, logic, and quantification are all tied up in an incoherent view of psychology.

This whole account I summarize in two points: (1) An understanding of language, (and other psychological processes), requires a dialectic between _descriptions_ of observed behavior which inform and constrain theoretical _explanations_ of unobserved internal processes. Hence, the term "theory" should be reserved for explanations, not descriptions. Historically, descriptions have formed the bulk of work in SLA to the point where they have been too often confused with explanation, with the result that now it is difficult to make the distinction. (2) An explanation of learning and development, by its very nature, requires a "developmental," or in other words an "active" perspective. This is so because the description of what people _do_, or their "task capabilities" to use Bickhard's terminology (like Piaget's stages), is _not_ an explanation of the psychological processes which account for those capabilities. To put it bluntly, the observation that at point X in development someone learning English correctly uses the 3rd person -s ending on verbs does not explain WHY or HOW he has learned to do so--the observation is not an account of learning. Such an account now seems to me to be a whole 'nother ballgame.
So is all of this ______ ?
A) painfully obvious
B) bordering on CT hereticism
C) evidence that I'm thinking too much
D) evidence that I'm finally catching on E) A & D
F) B & C ==

Date: Mon, 30 Dec 1991 11:14:45 PST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: marken@AERO.ORG
Subject: Control of behavior

[From Rick Marken (911230)]

After years of procrastinating, I am going to try to write a book about the control of human behavior. Of course, the book will really be about PCT. But I want to focus on what PCT has to say about human interaction (which often involve attempts by one person to control the behavior of others).

At one time the debate about controlling human behavior was quite intense; it was led, of course, by the great champion of "scientific" methods of behavior control -- B F Skinner. But the debate seemed to die out long before Skinner did. I think it had something to do with the success of teh cognitive "revolution" in psychology. I was never sure why cognitive psychology obviated concerns about behavior control. I have only one book that explicitly reals with this question -- U. Neisser's "Cognition and reality" (1976). He has about 5 pages on the implications of cognitive psych for the possibility of "social control". His conclusion seems to be that cognitive psychology implies that people are too smart to be controlled. At one point he says "the psychologist cannot predict and control anyone who knows more about the situation than he does" -which is true. But Neisser never explains why this is true. In fact, he never explains what control is -- and how this differs from influening or causing behavior.

I wonder if anyone on the net is aware of any relatively recent critiques of the notion of behavior control. Has there been anything published within the last 5 years or so that either advocates or trashes the control of behavior by scientific means. I would be especially interested in hearing about cognitivist advocates of behavior control. I bet there are some juicy tidbits in the sociology literature; any references to that literature would be most appreciated. Also, there must be some advocates of behavior control in the clinical field. Any recent books or articles out there?

Any help with references -- or net discussion of this topic -would be most welcome. Thanks a bunch. If I don't post before that magical moment, have a happy and safe transition into the New Year.

Regards

Rick
===

Date: Mon, 30 Dec 1991 16:13:35 CDT
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Comments: Please Acknowledge Reception,Delivered Rcpt Requested
From: RLPSYU08 <TBOURBON@SFAUSTIN.BITNET>
Subject: Re: Kolbe;semantics;control;M.Taylor

From Tom Bourbon -- [911230 -- 15:42]

David Goldstein [911227 & 911228] rose to champion Kolbe's book on the "conative connection." David, I wonder what I said in my earlier remarks about the book that led you to say that I have "a negative preconception about nonmodeling approaches." On that topic, I have no negative preconceptions -- my remarks about Kolbe came _after_ I scanned the book, not before. As I said in my previous post, I saw evidence that Kolbe accurately describes the importance of prior conceptions of results -- what we call reference signals -but Kolbe also gets many things wrong. Some of the things she says about behavior are not true of control systems. Rather they are loose ways of talking about behavior. (I am at a disadvantage My copy of the book is at home, where I am presently unable to log on to the network, so I cannot give specific quotations right now. They will follow later.)

I have scanned the book again and I stand by my first reaction -- close, but no cigar.

Also, I am far less wiling than you to attribute good understanding of the principles of control to people who use some of the jargon -perhaps I have read too many undergraduate papers to continue making that assumption. For example, several times in your posts, you say you assume or think Kolbe is talking about systems-level concepts, but you do not know that. She never says anything specific about how or why behavior occurs and she does not explain behavior in terms of a hierarchy like that in PCT. Ask her what she means, but do not assume she has it down pat. There is too much evidence that she does not.

Joel Judd [911230]. I vote "D."

Rick Marken [911230]. Go to it, on the book. But, before you do, check out the last umpteen years, right down to the present, in the radical behaviorist literature. That is a thriving group and its members insist that control of behavior is the supreme test of any scientist's understanding of behavioral phenomena -a test even more telling than prediction of behavior. Look in journals like _Behavior Analyst_, and _Behaviorism_.

Along with a few students, I will make a presentation on controlling the behavior of others, at Southwestern Psychological Association, in April. Then I plan to publish -- the first submission will be to _Journal of the Experimental Analysis of Behavior_, the behaviorist holy book. Any bets on the results? (The paper demonstrates how one person can control the actions of another, by disturbing a variable controlled by the person who is the object of control -- the controller can indeed control the actions of the controllee, but only so long as the controllee can control his or her controlled variable, and only so long as she or he does not know about, or care about, the controller. A lot of "ifs." Also, one model can control the actions of another, under the same conditions, and a model and a person can swap off playing the roles of controller or controllee.)

Martin Taylor. My plea [Tom Bourbon, 911226] still stands, especially in light of subsequent posts by you concerning your conviction that information-theoretic concepts can improve the performance of PCT models. Please show us how information-theoretic concepts can improve the results of any published account of quantitative modeling with the PCT model. I am admitting my own inability to see how or why information measures could produce such improvements. If they can, I want to understand how and why.

Prospero ano nuevo.

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

Date: Tue, 31 Dec 1991 13:07:37 EST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: philosophy of mind

Well, what's the PCT position on the `language of thought' issue? It seems plausible to me that among other kinds of goals, people have what one might call `propositional' goals, where what is intended is that some proposition be true, regardless of what kind of perceptions serve as evidence for this. For instance, if my goal is `fridge-door be open' (so that I can get a beer), I don't really care whether my evidence for the openness of the door be visual, haptic, or even olfactory (if I'm sufficiently single-minded in my desire for beer). I think I can envision a control theory for propositional goals, where the goals go in the `want box', and what the system does is try to get the contents of the want box to be included in the contents of another box, the belief box, these contents being sentences in a language of thought. But the LOT is a tricky issue, and I'm happy to read anyone who can tell cogent-seeming stories for or against it, e.g., the usual philosophers of mind.

Avery Andrews (avery.andrews@anu.edu.au)
==

Date: Mon, 30 Dec 1991 18:20:59 PST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: marken@AERO.ORG
Subject: Control of behavior

[From Rick Marken (911230)]

Tom Bourbon (911230) -- thanks for the info about behavior control. I would sure like to see a copy of your paper on the topic, the one you plan to submit to JEAB (my prediction -- it will be enthusiastically accepted and the entire staff of JEAB will join CSGnet on the same day). I do know about the journals that you mentioned. I know that there are still many behaviorists who believe in this stuff -- but is there any book that is like "Beyond freedom and dignity"(BFD) but is more recent. Perhaps a recent cognitively oriented rebuttel to Skinner's "classic". Maybe BFD (great initials for a book) is the only popular work of this sort. If so, that's ok. I'll work from there.

I just thought of a way to demonstrate the kind of operant behavior that the behaviorists love -- frequency of a discreet response -with a tracking type task. Maybe this is what you are doing? Instead of using a continuous control device (like a handle or mouse) use space bar presses that are integrated over time. The integrated output of the space bar presses is then added to the disturbance to produce the position of the cursor. Even with no disturbance the subject would have to keep pressing the bar at some minimum rate in order to produce an integrated output that keeps the cursor on target. Schedule effects can be simulated by varying the integration factor (for example) so that a higher or lower rate of pressing is required to keep the cursor on target. The data could be presented as "cumulative records". If this is something like what you've done, let me know. Whether it is or isn't, I think I'll set up the program to do this so that I can start getting some operant conditioning data without having to clean up any cages (except my own, of course). This would be a fun way to demonstrate control of behavior (or the appearance thereof) because it would be easy for an observer to watch changes in bar press rate as you change the schedule

(integration factor) or disturbance.
One last note -- I agree completely with your comments about Kolbe's book. I especially agree with your point that understanding PCT has to involve a bit more than making utterances that sound vaguely symathetic to our point of view.
Please send me a reprint of your behavior control paper -- or, better yet, post me a copy.

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: Mon, 30 Dec 1991 22:00:30 EST
Reply-To"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sende"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: goldstein@SATURN.GLASSBORO.EDU
Subject: last comment about Kolbe

To: Rick and Tom

The problem is that you are follow through types. Get a little quick start in you, will yuh. Or use a little fact finder to read the book more thoroughly. When all else fails, punch a bag to let out the implementor energy.

Best,
David

==

Date: Mon, 30 Dec 1991 23:46:02 -0600
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Behavioral Control; Description & Explanation

[From Gary Cziko 911230.23.22]

Rick Marken (911230):

>Also, there must be some advocates of behavior
>control in the clinical field. Any recent books or articles out
>there?

In education, it seems that the behavior mod types are almost always working in special education where they use conditioning techniques to get mentally subnormal individuals to do things like greet people and tie their shoelaces. This is interesting in itself, since it is consistent with Neisser observation that (normal) people are too smart for behavioral control techniques. I don't know the names of the journals, but I can get some of these if you wish.

==

Greg Williams (911230)

> I, for one, am happy that Bob Clark has joined the Net.

Me, too. I have had several interesting phone conversations with Bob over the past few months. I am very pleased that two of the three original developers
of PCT are now on CSGnet (Powers and Clark; we learned several months ago on CSGnet that the third, R. L. McFarland, has passed away).

> And since MCI somewhat mysteriously credits me with $100 from out of the >blue, maybe I will have a bit more to say than in the past.

So maybe there IS a Santa Claus after all. But instead of reindeers and sleigh he now uses electronic fund transfers! If so, good to know he is a friend of CSGnet.

==

Joel Judd (911230) (with hint to Greg Williams)

I'd vote D with Tom Bourbon, although I'm not sure that I've yet caught on enough to PCT to know if you have finally catching on.

Concerning the distinction between description and explanation, Greg Williams read a charming story he wrote about Fred (as in B. F. Skinner) and Bill (as in Powers) at our Durango meeting last August. I think his story captures some of what you are concerned with and perhaps his MCI Mail gift from Santa will motivate him to post it on the net. What say, Greg? It will be your Christmas present to the net. If you're still too cheap, send me the diskette (3 1/2") and I'll post it for you and even send you back the diskette.

Best to all in the New Year, Gary.

Gary A. Cziko
Educational Psychology   Telephone: (217) 333-4382
University of Illinois   Internet: g-cziko@uiuc.edu 1310 S. Sixth Street
Bitnet: cziko@uiucvmd
210 Education Building   N9MJZ
Champaign, Illinois 61820-6990
USA

===

Date: Tue, 31 Dec 1991 08:55:55 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: goldstein@SATURN.GLASSBORO.EDU
Subject: new year

I just wanted to wish all a Happy New Year.
Live long and may all your error signals be
small.
David Goldstein ===

Date: Tue, 31 Dec 1991 09:25:30 -0600
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: "Gary A. Cziko" <g-cziko@UIUC.EDU>
Subject: Re: philosophy of mind

[From Gary Cziko 911231.0930]

Avery Andrews (911231; Hm, how did you send this today when I read it yesterday?)

>Well, what's the PCT position on the `language of thought' issue?

While I don't know what any "official" PCT position would be, it seems that
Fodor's basic claims are quite incompatible with PCT and its view of reorganization. As I understand it, Fodor claims that the mind cannot create or generate ideas, concepts,

etc. any more complex than what it starts out with. I think that this is seriously mistaken. If it is logically impossible as he claims, then biological evolution couldn't work either and we woudn't be here wondering about this type of thing.

>It seems plausible to me that among other kinds of goals, people have >what one might call `propositional' goals, where what is intended is >that some proposition be true, regardless of what kind of perceptions >serve as evidence for this.

How could we possibly satisfy our intention that a proposition be true if we couldn't somehow perceive the proposition as true?

>For instance, if my goal is `fridge-door
>be open' (so that I can get a beer), I don't really care whether
>my evidence for the openness of the door be visual, haptic,
>or even olfactory (if I'm sufficiently single-minded in my desire
>for beer). I think I can envision a control theory for propositional >goals, where the goals go in the `want box', and what the system does >is try to get the contents of the want box to be included in the
>contents of another box, the belief box, these contents being sentences in >a language of thought.

Perhaps you should explain what you mean by propositional goals. Do you mean by this something that wouldn't fit into one the 11 levels of the type already included in "standard" PCT? Are you not just describing how the hierarchy would work in PCT with intentions at higher levels setting reference levels at lower ones?--Gary ---------------
----

Gary A. Cziko
Educational Psychology   Telephone: (217) 333-4382
University of Illinois   Internet: g-cziko@uiuc.edu 1310 S. Sixth Street
Bitnet: cziko@uiucvmd
210 Education Building   N9MJZ
Champaign, Illinois 61820-6990
USA

===

Date: Tue, 31 Dec 1991 10:02:00 -0600
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: jbjg7967@UXA.CSO.UIUC.EDU
Subject: season's greetings

HAPPY NEW YEAR!! ==

Date: Tue, 31 Dec 1991 11:43:16 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: happy new year!

[From: Bruce Nevin 911231 1124]

I have been following ongoing discussion with interest albeit without time to respond. I intend to take some time at my PC tonight and tomorrow and hope to get some responses off Friday. I will have more liberty than usual at home since Sarah and children stayed in NY after our Christmas gathering with family there--I'll drive down Friday night to bring them home. Sarah's mother is helping to see that she finally does get clear of pneumonia, certain conflicts of goals so far having interfered. One other touch of personal news, I wrote what I thought were satisfactory answers to the PhD preliminary exam in historical and comparative linguistics the 24 hours of the 18th and 19th, though I won't have official results until toward the end of January. (It would be nice to get good news on my birthday!)
Given respite, I have been reading a copy of Ben Franklin's autobiography that I picked up, and came across the source of a long-time favorite quote, which I offer you in

context of New Year's resolutions, and which I suppose bears on the relationship reference signals might have to self-instruction, "propositional goals," etc.:

> So convenient a thing is it to be a _reasonable creature_, since it enables one to find or make a reason for every thing one has a mind to do. (p. 43 of the Everyman edition)

May all your reorganizations seem after the fact as though purposeful!

Bruce Nevin
bn@bbn.com

==

Date: Tue, 31 Dec 1991 18:12:19 +0100
From: Oded Maler <Oded.Maler@IRISA.FR>
Subject: New year too

Dear camrades. I'm very busy lately trying to control my future position (together with my principles) so I'll just join the wishes (although, you know, we have another calendar, etc...) for a happy new year, "may you never understand that you are just a small component in a much higher level (social? cosmic?) system.."

And it's a nice tradition to recommend books so let me recommend Quine's Ontological Relativity, which I'm sure will be found very interesting by those interested in meaning, language etc.

All the best

--Oded ==

Date: Tue, 31 Dec 1991 10:56:00 MST
From: PETERS_R%FLC@VAXF.COLORADO.EDU
Subject: Language etc.

[From Bill Powers (911231.0900)]

Joel Judd (911230) --

Choice "D" from me, too, although I'd say you have caught on to a lot long before this.

>As the actual language which served as a basis for description became >obscured, the possibility that description could be confused with >explanation became real.

There are two ways that this possibility becomes an actuality. One is through statistical research; the other is through generalization and abstraction.

In the social sciences, the word "theory" is used to describe a proposed statement of relationship: people who have characteristic X exhibit a tendency toward behavior Y. I would call this a proposed fact: either X's show Y or they don't. If they do, we now have an observed relationship (never mind how reliable it is) that demands theoretical treatment. The corresponding theoretical statement would tack on "because ..." to the observation, and propose a mechanism that accounts for the observed dependency.

The other way in which description is confused with explanation is through the manipulation of categories. A specific instance of behavior by a specific person (Joe opens a door and walks out of the room) is converted to an instance of a class of behaviors of a class of persons (a male college sophomore exits from an enclosed space). The specific antecedent conditions are also converted to a category: "the room contains 400 people" converts to "the population density in the enclosed space is more than 2 persons per square yard." Now the happening becomes "A white male sophomore exits from an enclosed space when the population density exceeds 2 persons per square yard." This now looks like a more general statement that will apply to more people than just Joe and more instances of crowding in larger and smaller rooms. In many branches of the social sciences this is considered to be an explanation.

Of course the statistical approach and the generalization approach are used together.

Bickhard hits the nail on the head:

> "...(1) we can attempt to describe [a] system's class of >interactions [with the environment] from an external perspective; and >(2) we can attempt to explain those interactions in terms of processes >of the system." (1980, p.103)

This is exactly what I mean when I say that the theoretician has to take the point of view of the behaving system. When you imagine being a particular control system, you realize that the actual environment is almost irrelevant: all you can know about it is contained in the perceptual signal, and the relationship of the perceptual signal to external processes and entities depends entirely on how the input function is organized. So the control system can control only its perception; the effects it has on the external world while doing so are unknown to it.

> In order to provide an adequate _explanation_ of linguistic processes >we must, to paraphrase Bickhard, know what such processes ARE, and to >know that, we must explore the ontology of those processes. Never mind >the fact that we cannot observe them directly, we can still propose >testable models of the processes (which is what CT claims as well).

Precisely. The key is not so much being able to prove that the model is right, but simply understanding how to propose processes in such a way that they COULD be right. This amounts to appreciating what sort of thing has to be accomplished by the system in order for its externally-observed behavior to be as it is. We may not know how to build a general configuration-perceiver, but at least we know that the input has to be a set of sensations, and the output has to be a signal that covaries with our own sense of configuration. If we can think of one mechanism capable of doing this in one instance, that is better than not knowing of any mechanism. And when we have one mechanism that works, we can try to find another one that works better, seeing how the first one fails. And so it goes until we have a good model.

>To put it bluntly, the observation that at point X in development >someone learning English correctly uses the 3rd person -s ending on >verbs does not explain WHY or HOW he has learned to do so--the >observation is not an account of learning. Such an account now seems to >me to be a whole 'nother ballgame.

It is. To make models, we have to develop a skill, a sensitivity to what is unspoken or unexplained. If someone learning English puts an -s on third person words, we must first explain how that person recognizes a "third person" word. It's not enough to say that a social convention requires the use of such endings with or without apostrophe in such circumstances. We have to say how a person knows when to apply one convention or a different one. So this leads us to realize that we have to deal with meanings, not just statistical aspects of language usage.

The externalized viewpoint conceals such unspoken problems. To the external observer of language, it just seems that third-person words exist and are used in certain contexts, according to certain conventions. These are just facts of nature; it isn't obvious that the phenomenon of a "third-person word" reveals something about processes inside the speaker. The only way you can appreciate the problem is to get inside the speaker and catch yourself using a third-person word, and ask "How did I know that I should use that word?"

Thanks for a very interesting exposition.

Rick Marken (911230) --

It's not new, but you might look through _Beyond the punitive society_ (Wheeler, Harvey; San Francisco: W. H. Freeman, 1973). It's full of talk about "nice" control, and ends with B. F. Skinner's "Answers for my critics."

I'm sure you will say this, but just to make sure:

People are perfectly justified in disliking the idea of "control" in connection with human behavior. The problem is to convince them that control is a real and fundamental phenomenon of behavior, and that they dislike it because it has been misunderstood and misused. This is a delicate point. Lots of people see living a "controlled" life as being the opposite of freedom, creativity, and spontaneity. This is because they've been taught self-control in a way that guarantees internal conflict, and have been taught to try to submit to external control, which is impossible. In rejecting these teachings, they've rejected too much. They've failed to see that it was their own ability to control that was being violated in both cases, so they think that control itself is the Bad Thing. This subject is worth a long and careful exposition.

I really want to read this book. Hurry up.

Avery Andrews (911230) --

An aside. Mary's maiden name was Andrews, as in Mary A. Powers. One of her brothers is Avery Andrews, a dean at George Washington University. One of Avery's sons was Avery (911230) Andrews. So this Avery, a linguist, is my nephew, whom I admire. A hearty welcome to CSGnet, Avery. Please meet 120 of my friends.

>It seems plausible to me that among other kinds of goals, people have >what one might call `propositional' goals, where what is intended is >that some proposition be true, regardless of what kind of perceptions >serve as evidence for this.

Yes. To fit this into the HCT (Hierarchical Control Theory) framework, it's necessary to do a slight amount of translating. A proposition can be a reference condition, or it can be a perception. If it's a perception, it is being used as a description (or perhaps a "belief"): a conversion
from lower-level perceptions into the type of perception (symbolic expressions) handled at the category, sequence, and program levels of language. The perceived state of the world (inner or outer, present-time or imagined) is converted into a statement describing that state. The descriptive statement (proposition) is then compared with a reference-statement (proposition), and the difference is noted. If there is a difference, two paths may be followed depending on what you're doing. You could re-describe the lower-level situation to bring the description closer to the reference-description, or you could act on the world to make a different description appropriate, again bringing the description closer to the reference statement.

Making a CT model of linguistic processes means, among other things, trying to work out how the difference between two statements can be used to institute a corrective move at lower levels. Errors at one level should lead to changes in choices of statement elements at a lower level, the tricky part being that the changes should alter the perceived statement in a way that lessens the discrepancy. I don't know how this would be accomplished, but I suspect that the linguists on this list are working on it. If they aren't, why not, you guys?

> For instance, if my goal is `fridge-door be open' (so that I can get a >beer), I don't really care whether my evidence for the openness of the >door be visual, haptic, or even olfactory (if I'm sufficiently single>minded in my desire for beer).

This gets into the link between the purely linguistic hierarchy (dealing in words and word structures) and the general perceptual hierarchy (in which words are just another kind of perception). You can visualize a beer in your hand without any linguistic aids, and act to bring about a real matching perception. Or you can say "I wish I had a beer in my hand," which entails quite a different problem. Somehow the STATEMENT has to be translated into a specific visual/haptic/olfactory goal with which the actual experience of a beer in the hand can be compared: like must be compared with like. You can't compare the taste of the beer with the word "fizzy." You can only compare "flat" with "fizzy." The perception and the reference signal have to be in the same perceptual space, at the same level. So you can taste the beer, convert one aspect of the taste into a description, "flat," and THEN compare the description with the verbal reference signal (good, it's "flat", just the way I wanted it).

>I think I can envision a control theory for propositional goals, where >the goals go in the `want box', and what the system does is try to get >the contents of the want box to be included in the contents of another >box, the belief box, these contents being sentences in a language of >thought.

I'd revise that a little, viewing the "belief" as a perception of a state of affairs (or of a description of it). So the goals go in the "want" box, and the current belief (i.e., what you believe is actually going on, a perception) is compared with the goal. Action is taken until what you believe (perceive) to be present or occurring matches what you want to be present or occurring.

So the proposition "I'm drinking a beer," as a goal, is compared with the proposition "I'm drinking a ginger-ale," and the action is to throw out the ginger-ale and get a beer. Getting a beer converts the perceptual proposition, the belief or description, to "I'm drinking a beer," which matches the goal proposition, if all the attributes of the beer meet their respective standards at the lower levels, thus meriting the descriptive term "beer."

When Bruce Nevin gets back on the net, I'm going to propose some steps toward building a real HCT model of language. Bruce follows the Harris approach, and I know you follow the Chomsky approach while Martin Taylor uses probabilistic concepts, but I think we can sketch in a way to test
all these concepts within a control-system model. Just as a preview, parsing would be considered a mode of perception.

R. K. Clark (911229) --
Welcome, Bob. You're right that the discussion of control-system engineering would have benefited from speaking about the levels, but the main problem as I saw it was trying to get a control-system engineer to see that his diagram and the CSG-type diagram are really the same. Obviously, my attempt didn't work. Our friend eventually signed off the net, wishing us luck but saying that as an engineer he simply HAD to think of controlled variables as outputs. Sound familiar to you?
I hope I'm talking about the same thing you mentioned.
[Bob Clark and I developed the basic control-system model in the years from 1953 to 1960. I have never properly acknowledged his part in this development, which was major.] ------
----
Gary Cziko (911230) --
I'm sending you the stuff on floppies. There's no guarantee that the transmission from here to Ft. Lewis is error-free. The Little Man version 2 (in C, with dynamics) isn't quite ready -- still looking for that perfect way to stabilize the arm. I promise to let it go before it's perfect, though.

Best to all,
Bill P. ==

Date: Tue, 31 Dec 1991 13:43:07 EST
From: "Bruce E. Nevin" <bnevin@CCB.BBN.COM>
Subject: untrustworthy sort

[From: Bruce Nevin (911231 1324)]

The tradition of recommending books appeals to me. A number of us might enjoy Jim McCawley's "How far can you trust a linguist?," in which he discusses "didespread disparities between what linguists believe and what they say they believe." His remarks on language acquisition are germane to your present concerns, Joel. The paper is at pp. 75-87 in

Simon, Thomas W. and Robert J. Scholes. 1982. _Language, mind, and brain._ Hillsdale, NJ: Lawrence Erlbaum Associates.

This collection stemmed from an interdisciplinary symposium on the topic of the title sponsored by the Sloan Foundation in 1978. The intro is by Steven Harnad.

I should say that McCawley is a linguist in good standing at the University of Chicago. With a fine sense of humor.

Bruce Nevin
bn@bbn.com ==

Date: Tue, 31 Dec 1991 14:04:16 EST
From: Martin Taylor <mmt@DRETOR.DCIEM.DND.CA>

Subject: Re: Kolbe;semantics;control;M.Taylor

[Martin Taylor 911231]

(Tom Bourbon 911226 and 911230)
> Martin Taylor. My plea [Tom Bourbon, 911226] still stands, >especially in light of subsequent posts by you concerning >your conviction that information-theoretic concepts can >improve the performance of PCT models. Please show us
>how information-theoretic concepts can improve the results >of any published account of quantitative modeling with >the PCT model.

I haven't been avoiding the question. I've been both trying to think about it and trying to shake a mild flu that has been interfering with all but the simplest thinking.

Your question is very interesting to me, because I haven't ever thought about using information theory in the way you suggest. I have always used it in a qualitative way, as, for example, in showing why higher-level control systems should be expected to work more slowly than low-level ones, and why there should be levels of abstraction in communication. The published studies of control, as far as I know, have been dealing with low-level control systems, whether coordinated or not. My understanding of those is that the accuracy with which the models replicate the performance of the human subjects is based on the estimation of parameters such as gain and transport delay, and their optimization. Information theory has nothing to say about the shapes of models, and optimization of the models would be the same (I think) whether the behaviour within the models were described in terms of uncertainty or in terms of variance. The accuracy of the model predictions is based on point estimators. It is conceivable that the predictive errors might be used to estimate uncertainties in different parts of the control loops concerned, and that is where my thoughts have so far been (unsuccessfully) trending in trying to answer your question.

I don't think my posts have been aimed at improving the performance of PCT models, so much as trying to place constraints on what models should usefully be considered. For the most part, the result is simply too confirm that Bill Powers' ideas are at least consistent. Nothing that I can see has so far suggested that any of what he has said could not be so.

As I remember, this whole information and uncertainty thread arose because of what I considered to be a misuse of the term "information" in the discussions on language and PCT. There turned out to be less understanding of information theoretic ideas within the (responding) community than I had supposed, so I have tried to improve the situation. With luck, the concepts will become sufficiently intuitive that modellers will unconsciously embed them into their models.

A New Year's Resolution: I will try to see whether anyone has used information-theoretic ideas in the analysis of control systems (at least at a mathematical level I can understand). In statistical analysis, I-T ideas are often more powerful and general than linear analysis such as ANOVA, because they make no assumptions about linearity. The

analyses do, however, take into account what is known about one thing when estimating another, and that should be relevant in the control-system analysis. Although I have used I-T concepts consistently throughout my career, it is decades since I have worried about their exact mathematical form, and this question of yours intrigues me enough to go back and try to relearn it properly. If time permits... As always when I post from home, please forgive extraneous characters in the above. I don't know which are outgoing and which come only on the echo from the other end.
Happy New Year to all.

Martin ==

Date: Wed, 1 Jan 1992 10:44:02 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: philosophers of mind

>While I don't know what any "official" PCT position would be, it seems that >Fodor's basic claims are quite incompatible with PCT and its view of >reorganization.

My guess would be that Fodor is right about short time-scales (minutes & seconds), but wrong about longer ones (days ...). Notice that a lot of vocabulary acquisition happens on the short time scale. More generally,
I certainly wouldn't want to claim that any particular philosopher is right about everything, or even very much, but just that there are useful ideas there.

>How could we possibly satisfy our intention that a proposition be true if
>we couldn't somehow perceive the proposition as true?

Well, it seems to me that in the general case we don't *perceive* propositions as being true, but *judge* them to be true (I don't perceive the gold ring
I just bought to be made of gold, since I wouldn't have a clue as to how to tell real gold from a serious attempt at fakery, but I believe it to be gold because I just bought it from a jeweler said it was, believe that jewelers operating from shops in Ozzie shopping malls don't lie about what their selling, etc. Judgements obviously can be regarded as `higher order perceptions', but it seems to me to be a reasonable bet that they are a `natural kind' of H.O.P. that philosophers actually know quite a lot about (although distinguising the actual knowledge from the trash might not
be easy).

>[Ahierarchy would work in PCT with intentions at higher levels setting >reference levels at lower ones?

Does anyone have a clear idea of how to build a `fridge door open'
(or `leopard nearby') detector along the lines of the hierarchy? My understanding of the higher level portions of the hierarchy was that they are rather tentative suggestions, and I confess to having never
been very happy with them (what's the latest version, anyway? the latest I've seen has 10 levels, not 11). Consider the the LOT-based propositional goal satisfier to be an alternative tentative suggestion about how
some of the higher level stuff might work.
Avery Andrews (avery.andrews@anu.edu.au)

dd
==

Date: Wed, 1 Jan 1992 16:23:14 EST
Reply-To:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
Sender:"Control Systems Group Network (CSGnet)" <CSG-L@UIUCVMD.BITNET>
From: Avery Andrews <andaling@FAC.ANU.EDU.AU>
Subject: Some of Bill's Responses

Powers (911231) --

>The only way you can appreciate the problem is to get inside the speaker >and catch yourself using a third-person word, and ask "How did I know
>that I should use that word?"

No. The first wisdom of linguistics is that speakers are always wrong when they try to explain why they say what when (the stories are pathetic, and tend to fail within 30 seconds). Joel is quite right
in saying that linguists tend to confuse description and explanation,
but getting behind the descriptions to the explanations will be
reverse engineering all the way (I regard current linguistics as being essentially descriptive, in spite of the presence of a lot of talk about explanation). One important difference between
talking and ordinary behavior is that such control systems as there
are to govern the structural aspects of language use are limited, and don't work very well. On the one hand there are no
significant disturbances to prevent you from saying something that means one thing rather than another, and on the other hand it's
incredibly difficult to tell what the people who you are talking to actually make of what you are saying, so it, I would say, primarily a matter
of flying blind via feed-forward (which is one reason why most people do it so badly).

>Making a CT model of linguistic processes means, among other things, >trying to work out how the difference between two statements can be used >to institute a corrective move at lower levels. Errors at one level >should lead to changes in choices of statement elements at a lower level, >the tricky part being that the changes should alter the perceived >statement in a way that lessens the discrepancy. I don't know how this >would be accomplished, but I suspect that the linguists on this list are >working on it. If they aren't, why not, you guys?

Well, it's a hard problem, and it seems to me that there's a tremendous amount of work in AI directed towards solving it. There are at least
two issues to be dealt with: a) what is the structure of the representations in the `reference signal' (desire box) and `input signal' (belief box)
b) how are the actions calculated to get them converged. (a) is something which linguistics can contribute to, but not (b), I would say.

Removing my professional linguist's hat, the best I can suggest re (b)
is something along the lines of how PROLOG programs work. Suppose that the goal is to `have coffee'. If you are in possession of a Vesuviana coffeemaker, this goal can be obtained by satisfying the following sequence of intermediate goals:

1) there be ground coffee in the coffee-receptacle
2) there be water in the bottom reservoir
3) the device be assembled and on the stove
4) the burner the device is on be on.

[and then waiting--I don't know how to put this bit in]
This is a lot like the pattern-action rules of some AI programs,
or a sequence of statements in a conventional programming language, except that in the latter cases the feel seems to be wrong,
due to emphasis on the events to be initiated rather than the resultant states to be achieved. Two of the many troubles with my little story]
is (a) it begs the question of where the `programs' come from. Lizards get theirs from there genes (if you want to be cooler, head for the shade), but we are able to make ours up. (b) in interesting cases, there are a
horrifyingly large number of different paths what would have to be explored to figure out how to satisfy the desired goals, and it's not at all clear how we are able to find our way through the maze.

>Somehow the STATEMENT has
>to be translated into a specific visual/haptic/olfactory goal with which >the actual experience of a beer in the hand can be compared: like must be >compared with like.
I'd suggest that David Marr's work on visual perception & Ray Jackendoff's on natural language semantics (Jackendoff 1987, Consciousness and the Computational Mind; 1991

Semantic Structures) is quite relevant to the conversion problem. A basic idea of this work is that we construct a
3-D model, mostly from visual and haptic information. Categories can be seen as labels pinned to objects in the 3D model, relations as lists of pairs, triples, or whatever, of objects from the model, & so
forth. I have a paper draft touching upon some of this which I could send to anyone who's interested.

I have a theory about the comutational rationale for the propositional (categories, relations, programs -- I'm not convinced that these belong
on different levels) level, which goes like this. The path from the actual situation to an acceptable one is often very indirect, & organisms can't afford to explore all the possible ways to go with their bodies. So
we have a simulator (`imagination'), to calculate the likely outcomes of various courses of action. But full analog simulation is also extremely expensive, whence the utility of propositions and logic
as a cheap (though often nasty) substitute. To predict what will happen when I pick up a bucket containing tennis-ball and carry it around the house I
don't have to try to imagine in detail how a tennis-ball will behave in a moving bucket, but can using the principle:

```
bucket(b)
thing(e)
in(e,b)[t1]      (e is in b during time period t1)
goto(b,p)[t2]
within(t2,t1)
after(t3,t2)
within(t3,t1)

      ->
      at(e,p)
```

In effect, propositional representations save computuational effort by omitting details that are (usually) irrelevant to the important features
of outcomes.

>So the proposition "I'm drinking a beer," as a goal, is compared with the >proposition "I'm drinking a ginger-ale," and the action is to throw
>out ..
I'd doubt whole propositions have to be compared here. E.g. `I'm drinking a beer' implies `I taste beer', which is perceptually false, while the
original propositional goal is known to not be satisfied. It won't
be easy to figure out how the interfacing and the logic works.

Avery Andrews

Gary A. Cziko Internet: g-cziko@uiuc.edu
Associate Professor Bitnet: cziko@uiucvmd
University of Illinois at Urbana-Champaign
Bureau of Educational Research
1310 S. Sixth St.--Room 230 Champaign, IL 61801-6990 USA

Telephone: (217) 333-4382 FAX: (217) 333-5847

==

Date: Sun Jan 05, 1992 8:14 am EST
From: Control Systems Group Network
EMS: INTERNET / MCI ID: 376-5414 MBX: CSG-L@vmd.cso.uiuc.edu

TO: Dag Forssell / MCI ID: 474-2580
TO: Robert K. Clark / MCI ID: 491-2499
TO: * Hortideas Publishing / MCI ID: 497-2767
TO: Multiple recipients of list CSG-L
EMS: INTERNET / MCI ID: 376-5414
MBX: CSG-L@UIUCVMD
Subject: Levels, Description and Explanation

[from Gary Cziko 920104.2045]

Greg Williams (920101)

Greg, I'm still mulling over your correspondence with Dennis Delprato which you shared with us on the net on the first of the year. I find I need some more help to understand all the implications of your view on description and scientific explanation. To set the context:

Press RETURN for more; type NO to stop:

> In sum, then, my notion is that, contrary to the view
> generally held by scientists, genuine (extrapolative, rather
> than summarizing) prediction and control of phenomena at
> level n can be achieved only by theories couched in terms of
> level n-1.

Now the puzzle:

> This implies that empiricist theories in
> psychology can be used to (genuinely) predict and control
> (and explain, as I use the term) sociological phenomena, NOT
> psychological phenomena. Empiricist theories in physiology
> are required to (genuinely) predict and control psychological
> phenomena.

My problem is understanding how description in psychology can be used to explain sociology. For argument's sake, let's say that people are as Skinner conceived them to be and that we have data showing that they will produce certain behaviors to get certain rewards (money, food, sex, etc.). Now you say this can't be used to explain individual behavior but CAN be

Press RETURN for more; type NO to stop:
used via sociology to explain some aspects of group behavior.

But if the empiricist psychology of Skinner is just a summary of observations and gives you no basis for generalization (and you must always have generalizations to predict since conditions are never exactly the same), then how can this be any better for sociological prediction? I can't see how moving up a level to n+1 solves the problems that are there at level n. You will now want to be able to predict that given a bunch of people under certain conditions they will interact in a certain way as a group. But how can this be done if your psychology in inadequate to begin with? Why don't the inadequacies at any level n entail inadequacies at all levels greater than n? You can see the pit I am falling into here. Can you stop my fall, or at least provide a soft landing spot?--Gary

P.S. Maybe some examples would help me to understand these ideas better.

Gary A. Cziko
Educational Psychology   Telephone: (217) 333-4382 University of Illinois
Internet: g-cziko@uiuc.edu
1310 S. Sixth Street   Bitnet: cziko@uiucvmd
210 Education Building   N9MJZ

Press RETURN for more; type NO to stop:
Champaign, Illinois 61820-6990
USA
