Thread on levels of perception

Unedited posts from archives of CSG-L (see INTROCSG.NET):

For details of modeling a functioning hierarchy, and discussions of how they work, see several posts in documentation of PCT demos, file ARM2\ARM2_CSG.NET, and the file BehaviorOfPerception.pdf.

Date: Tue Mar 24, 1992 2:51 am PST
Subject: Levels of perception
From Bill Powers (920324.0300)] Reference: Mark Olsen (920323)

You ask about the functions relating one level of perception to another. This is indeed the question that HPCT poses -- but doesn't answer. What lies behind HPCT is not any proposal as to how each level of perception is derived from the one below it, but a proposal as to what the levels of perception are and how they are related. This is the phenomenon that any model must in the end explain.

The "H" part of HPCT can be taken in two ways: first, as a general sketch of a hierarchy of control in the abstract, with the communication between levels consisting of a series of perceptual re-representations of reality and a corresponding set of reference signals used to control lower levels; second, as a series of proposed levels of perception (and control) based directly on an analysis of experience with the hierarchical control concept as a guide. This is a beginning model; there may well be other modes of communication between levels, but the basic one is probably valid.

The definitions of levels define the modeling problem. We can see that the sensation level is probably derived by weighted summations of intensity signals, the weights defining a vector in a perceptual space having fewer dimensions than there are different sources of intensity signals. But that answer to the modeling problem comes after noticing that sensations seem to depend on intensities in a particular way, a way that could be modeled as weighted summation. The phenomenon to be modeled comes before the model.

And that's as far as I can go. I don't know how configurations are derived from sensations -- how it is that we can get the sense of, say, a particular person's face over a range of distances and orientations and expressions. If signals standing for the dimensions of a face existed, then it's possible to make a rough guess that transitions of the face from one state to another would be sensed using time functions and partial derivatives; that's a feeble start toward a functional model that you could run on a computer. As to the rest of the levels, the kinds of computations involved are mostly a mystery to me. The few guesses we have come up with are strictly stabs in the dark. You can use words like "integration" to describe how some kinds of perceptions are put together to create others, but the word is just a noise. It doesn't tell us anything about the processes involved.

Behind this exploration of perception lies a fundamental postulate; if you don't internalize it, I don't think you can even get started on the problem of modeling the brain's perceptual systems, or for that matter, in understanding HPCT. The postulate, simply put, is this: it's all perception.

By that I mean that no matter what you attend to in the world of experience, whether you refer to inner or outer experiences, concrete or abstract, verbal or nonverbal, the object of your attention is a perception. You are looking at or otherwise experiencing the brain's perceptual activities, not the objective world itself.

Vision is the most important sense to understand this way if you're sighted; understand vision and the rest (touch, taste, sound, etc.) will follow. The world you see begins as pixels (individual picture elements). The pixels are so close together that you see no spaces between them, although the sensory nerves do not overlap and in fact do not completely fill the retina. There's a world between the pixels, but we don't see it unless the view shifts slightly -- and then what we had been seeing disappears into the cracks between the

pixels. This is invisible to direct experience; the world seems continuous over the whole visual field. We get a sense of seeing the world at infinite resolution, and can't imagine what the whole field would look like if we had, say, ten times as many retinal receptors and the optical acuity and brain power to take advantage of them. This would be like seeing the world through a magnifying lens, except that the whole world would look that way, not just one little part of it (which we still see at human resolution). The only way to imagine this is to go the other way: view the world at a lower resolution, as in a halftone photograph or a television screen seen close up, and imagine that the result is the only world you can ever see. That's how our picture of the world would look to a different organism with higher visual resolution. But we experience it as having continuous detail right down to the level where it appears smooth. I suppose the fly sees the world in the same way. But its world is smoother than ours.

Building up definitions of the rest of the levels in the hierarchy is then a matter of noticing persistent types of structure in this world of picture elements. The first level above the pixels themselves is sensation, a type of perception that can't be analyzed in any way except into variations of intensity. Color is a sensation, as is shading.

Perhaps things like edges are sensations, derived in one step from the pixel distributions. When analyzing perceptions, however, don't use any data but your own experience. Theory and neural data will tell you that in the visual field, in the retina itself, all edges are enhanced, so that there is a strong outlining effect. But look at the edge of a sheet of paper on a dark tabletop. There is no outline. The closer you look at the edge, the more nearly it seems to be an infinitely sharp line separating uniform white from uniform dark. The edge itself is there -- but you can't see it as an object. It's just a sense of edgeness. Only under special conditions, as in looking at a smooth gradient of illumination going over a relatively short distance from white to black do you see edge effects like the "Mach band", the only clear subjective evidence of edge enhancement. However those neural signals enhanced at edges are processed, the result is that step changes look like step changes, not outlines as in cartoons. Whatever model we come up with for how the nervous system processes pixel information, it must result in edges that look this way, without borders. If it doesn't, the model is wrong.

The next step is to notice that the edges and corners and broad white areas of the piece of paper add up to -- a piece of paper. If you've made this transition properly, it will come as a surprise. Where did that piece of paper, or piece-of-paperness, come from? It wasn't there in the edge, or the corner, or the whiteness, or the darkness. It comes into being only when all those elements are seen grouped into a thing, a configuration with a familiar shape, orientation, distance, size, and so on. The Gestalt psychologists of old spent a lot of time looking at things like these. They should have kept going. Or perhaps they shouldn't have been cowed by the behaviorists.

You have to go slowly and by the smallest steps you can devise. If you go too fast you'll miss the smallest steps; if you miss the smallest steps you'll lose the sense of examining perceptions, and start projecting the visual field into an external world again. You'll jump to the more abstract levels and lose the connection from one level to the next. This is, if you like, a form of meditation on experience in which you distance yourself from experience and look at it merely as a display. You're not trying to see anything about the world, but only something about the display. You're trying to see what features the person who constructed it thought of putting into it, just as when you read a program you think to yourself "Now he's setting up an array to hold the results" instead of just reading the code, or when you read a novel as a literary critic you think "Now he's introducing tension" instead of just getting tense. Who the "he" is is immaterial -- the point is to see what is before you as a construction that has inner organization, and try to see how it is put together.

The general principle is that when you have found a level, like sensation, the next level is going to depend on it; also, the current level depends on the one below it. If you analyze a perception to see what it is made of, at first you see just more perceptions of the same level -- big configurations are made

of little configurations. But when you analyze in just the right way, you suddenly realize that all configurations, of whatever size or kind, are made of sensations, which are not configurations of any kind. And you realize that if it weren't for the presence of those sensations, there couldn't be any configuration to see: a field consisting of a single sensation, such as white, can't lead to any sense of configuration. There's a relationship between these levels of perception. That gives us a hint about building models of perception, a hint about how the brain's perceptual system is constructed.

Sometimes you will identify what seems to be a higher level of perception, some characteristic common to all perceptions, unconnected to lower levels you have previously seen. Then you can use this kind of analysis to try to fill in the gap. What is this new perception made of, besides smaller perceptions of the same kind? When the gap is large, the missing steps are obvious. You can, for example, look at spatial relationships such as "on" -- something being "on" something else. You can see the on-ness clearly, it's right in front of you. But what is it made of? If you said "sensations" you would clearly be making too large a jump, because on-ness involves objects, things, configurations. Some kind of object is "on" some other kind of object. If it weren't for the impressions of distinct objects, there couldn't be any sense of the relationship between them. But is that step small enough? I've had to put two levels between relationships and configurations: transitions (which can be zero), and events (which can be as simple as mere duration). Seeing something "on" something else involves more than a brief contact; there must be duration.

Perhaps someone else could find smaller steps still, or would characterize the intervening steps differently. There's still a lot of room for improving the definitions of the phenomena we're hoping ultimately to model.

I'm not talking here about the models themselves. I'm talking about the attitude you take toward your own experiences when you're trying to notice phenomena that need modeling. If you were a physicist you wouldn't be taking this attitude. You'd treat the world of perception in the normal unanalytical way as if it lay outside yourself where everyone could see it, and you'd search for laws relating changes of one kind of perception to other kinds of perceptions. You would call these "natural laws" or "behavioral laws" and assume you were discovering truths about an objective universe.

As a CT psychologist, however, you have a different objective: to grasp the natural world as a manifestation of human perception (your own), and to ferret out of it some regularities that tell us about perception rather than about the world perceived. If you stumbled onto this attitude accidentally, without understanding what you were doing, you might well find yourself in a state with a clinical name: dissociation. I don't recommend this attitude as one suitable for ordinary living. It's difficult and uncomfortable, and it tends to strip the meaning from experience (until you get past a certain point, after which you realize that meaning, too, is perception, and let it back in). If you're afraid that understanding your girl friend as a set of intensities, sensations, configurations, transitions, events, relationships, categories, sequences, programs, principles, and system concepts in your brain might strain your feeling toward her (and hers toward you), don't do this with your girl friend. Do it with somebody else's, or a laboratory rat. It doesn't matter who or what you do it to, because you're really talking about your own perceptions. This is a private experience valid only in one person's world. It can become public only to the extent that different people independently arrive at the same analysis. I've always hoped for that, but only a very few people, to my knowledge, have tried this for themselves. Most people just memorize my definitions, which unfortunately are in words. It's easier to push words around than to shut up and examine direct experience.

You'll hear objections to this process alluding to introspectionism, which failed to get anywhere a long time ago. But introspectionism didn't fail because it looked at the kinds of things I'm talking about here. It failed because it confused the subjective with the objective (and so did its critics). The world that I'm speaking of examining here would be called, by most conventional scientists, the objective world, not the subjective one. I'm not recommending shifting attention off the objective world and plunging into

the dim and uncertain world of inner phenomena -- or what we imagine to be inner phenomena. I'm recommending a change of attitude toward the world we normally consider to be the objective one, which includes the world outside us and our bodies as we experience them. I'm saying that you will learn something if you look on this world as directly experienced evidence about the nature of your own perceptual system, and only in a conjectural way about the world that is actually outside you.

Instead of treating relationships like on, beside, after, with, and into as properties of the external world, look on them as perceptions constructed on a base of lower-level perceptions. Instead of seeing categories as made of things that are inherently alike, think of categories as ways of perceiving that MAKE things appear to be alike -- things that are actually, at lower levels of perception, different. Instead of seeing sequential ordering as a fact of nature, see it as a way of putting ordering into an otherwise continuous flow. In short, take nothing about experience for granted, as if some aspects of experience were really outside and others were inner interpretations. Put the whole thing inside, and see what you come up with when you understand that it's all perception. All of it.

Final notes:

In HPCT diagrams, we show signals coming out of perceptual functions and going into higher-level ones (as well as the local comparator, if the signal is under control). I think of these lines as representing single neural signals that vary in only one dimension: how much. This can be confusing, because we don't experience single signals under normal circumstances (when we do they cease to be meaningful). Instead we experience all the signals within the scope of awareness, at every level in the state we call conscious. To understand what the single-signal concept means, you have to break this world of simultaneous perceptions into its components, the individual and independent dimensions in which the totality of perception can vary. You have truly identified one isolated perception when it can vary only in the degree to which it's present, which we experience as its state. If the perception varies without in the slightest changing its identity, you have probably noticed a single signal.

This can be important when you talk about control. We talk loosely about controlling "a dog," for example. But that way of talking is really lumping many independently variable aspects of the dog together. You don't control its species, or its eye color, or the length of its tail. You don't even control its behavior. If it's behavior you're controlling, you always control SOME PARTICULAR VARIABLE ASPECT OF THE DOG'S BEHAVIOR. You may control the radius within which it can move, by putting it on a chain. You may control its speed of walking by saying "stay" or "follow," and its path by saying "heel." Whatever you control, it must come down to a single variable or small sets of variables independently controlled. If you're controlling in more than one dimension, you must sense more than one variable, and have a control system operating independently for each one. That's because independent dimensions can be independently disturbed; you need independent control systems so that a disturbance in one dimension can be corrected without necessarily causing an error in another dimension.

None of this answers your question as to how perceptual signals in a diagram depend on perceptual signals lower in the diagram. The only general answer I can give is that some computation lies between them. The input data consists of lower-level perceptions; the output data, the higher-level perceptual signal, represents the value of the function being computed over and over or continuously. At each level, I presume (judging from the way the context changes every time you consider a higher level), a new type of computation is involved, not simply a repetition of the kind of computation at the lower level. The process of deriving categories from sets of relationships can't be carried out by the same kind of computation that derives relationships from sets of events or lower perceptions. There is no one kind of computation that could serve at all levels.

But as I say, I am, we all are, a very long way from grasping what these kinds of computations are. Every time people come up with a new computer program for

recognizing objects, they try to establish this new computation as the blueprint for the whole perceptual system. This is a waste of time. The blueprint changes with every level. Weighted algebraic summation is simply not going to suffice to model our capacity to recognize and execute a program described in words: a rule. Even though such networks are purported to recognize categories, I think that the categoryness is read into the results by a human observer. I don't think that any category- recognizing back-propagation model will actually create what human beings experience as categories -- for example, the category "wife." Of the eleven levels of perception in my model, I think we know how model two of them, the first two. All the rest of our modeling presents to us what a human being might recognize as a higher-level perception, but which the circuit or program itself does not recognize -- or control.

In that I could be wrong, of course, because I speak the truth when I say I don't know how the higher levels of perception work. That means I don't know how they don't work, too. I'm just expressing a hunch.

It's late and I've posted this so I could get to sleep (some ideas just have to leak out through the fingers before they'll let you alone). I'll get to comments on other interesting mail tomorrow.

Best, Bill P.

Date: Tue Mar 24, 1992 5:12 am PST
Subject: RE: Bill on Levels

From Pat Williams (920324)

I really liked your explanation of "Levels of Perception," Bill. Hearing how you arrived at the levels you have found, rather than just names and descriptions of the levels makes it much clearer to me. It is really fascinating to me to think about how perceptions can be combined to form higher level perceptions. I have a hunch that computer programmers (at least good ones like you) may be better at this kind of thinking than most people, since they have to break things down to minute details to make anything work. Just determining the simplest perception like edge recognition is amazingly complicated. I'm currently working on an automatic curve tracer for PictureThis. Determining the edges, corners, and intersections of curves when you only have local pixels to work with is incredibly difficult. It seems fairly trivial until you try it and find all the exceptions. And of course that is no where near as complicated as what you are trying to figure out.

Best wishes, Pat

Date: Sun May 31, 1992 6:54 am PST
Subject: Abstract, concrete, HPCT

[From Bill Powers (920531.0700)] Copy to CSGnet.

Greetings from CSGnet. My name is Bill Powers. I have just received a copy of your delightful paper (with Agre), "Abstract reasoning as emergent from concrete activity," from my nephew Avery Andrews, who is a linguist residing in Australia. How is "Agre" pronounced? Is it "ah-gruh" or "aeger?" Or something I haven't guessed? I think I can handle "Chapman."

There are some points of contact between your ideas and the basic model that's behind CSGnet (a Bitnet-internet list). The CSG stands for "control systems group", which is a collection of people (including many off the net) from many disciplines who have taken up some ideas I developed (with Clark and McFarland) in the 1950s, and have been working on since then.

There are three aspects of this theoretical framework.

One, called CT, or control theory, is just the basic body of theory developed by control-system engineers in the 1930s and 40s to describe and predict the behavior of closed-loop negative feedback systems -- servomechanisms, regulators, and such.

The second is PCT, or perceptual control theory, which is the adaptation of CT to the universe of organismic behavior (starting with Wiener, Rosenbleuth, and Bigelow but branching off quite early from cybernetics). The basic idea behind PCT is that living control systems act to bring perceptual representations of external variables into a match with internally-specified reference signals, maintaining them in a near match despite changes in the reference signals and occurrence of external disturbances tending to alter the perceptions. "Perception" is used in a generic sense to mean all experiences from raw sensory input to abstract representations. We talk about PCT when we mean to indicate only that some perception is under control by behavior, the kind of perception being secondary.

The third aspect is HPCT, meaning hierarchical perceptual control theory. This is not really control theory per se, but an attempt to introduce facts of experience and some neurological facts into the general model, to make it specific to human experience and human architecture. I'm going to bore you with a rather detailed description of this hierarchy, because unless you understand it you won't see how it relates to your work.

The concept behind HPCT is a hierarchy that runs in two directions: a perceptual hierarchy building upward, and a control hierarchy building downward. A given level (containing many control systems) receives inputs that are copies of perceptual signals of lower order, some under direct control and some uncontrolled. A perceptual function in a specific control system generates a new signal that represents a variable of a new type, derived from lower-level perceptions (or sensors, of course, at the lowest level). A comparator compares the state of this signal with a reference signal received from systems of a higher level. The error signal resulting from the comparison goes to an output function that ends up distributing reference signals to control systems of the next lower level -- the same level where the perceptual signals originated. Only the lowest level of outputs generates muscle action.

So each level of system acts to match its own perceptual signals to reference signals received from higher levels, and acts by means of varying reference signals for systems at the next lower level. The result is a hierarchy of goal-seeking and goal-maintaining control systems with many systems at each level and many levels. The highest level of reference signals has to be handled in a special way, of course, which I won't get into here.

The first level of perception is called the "intensity" level. The perceptual signals at this level are generated by sensory nerve endings (a perceptual signal is measured in impulses per second -- individual spikes have no significance in this theory). Each first-order perceptual signal therefore represents the intensity of stimulation in one sensory ending. As neural

signals vary only in magnitude (carried as a frequency), they are one-dimensional: they represent how much stimulation is present, but not what kind. So the first level of perception is a collection of millions of signals representing pure magnitudes: essentially, positive numbers. This first level of perception contains all possible information about an external world, as far as the brain is concerned (meaning, of course, as far as we are concerned). Some first-level perceptions are under direct control: primarily, those representing muscle stretch and tension. We experience these as "efforts."

Second-level perceptions are functions of sets of first-level perceptions. The functions are probably weighted sums. The signals that result are called "sensations," which are vectors in little subspaces made of a few independently variable intensity signals. Taste, for example, seems to be a function of only four kinds of intensity signals. Color seems to be a function of three kinds. Second-level sensations are controllable by varying reference signals for those first-level perceptions that are under control: muscle tensions. Most second-level sensations are not under control. There are probably uncontrolled perceptions at every level, although fewer at the higher levels.

Sensation-signals, just like intensity signals, can vary only in magnitude: one signal can represent only how much of the particular sensation is present, not what kind it is. The "kind" is determined by the weightings applied to the intensity inputs in the perceptual function. So this is a pandemonium model: one control system controls only one kind of perception, and controls it strictly with respect to its magnitude. This holds true at all levels.

Third-level perceptions, functions of sets of sensation-signals, are called "configuration" signals. I don't know the nature of these functions, or of any perceptual functions from here on up. At this level, the world of objects and static patterns comes into being. But there are also sound configurations (phonemes, chords), tactile configurations (a squeeze), and somatic configurations (internal feelings like nausea) -- all sensory modalities are involved. A given configuration signal has a magnitude that indicates the degree to which a given kind of configuration is present. One signal can represent only one kind of configuration.

This is the perceptual world that we think of as consisting of "concrete objects." You see where I'm going -- this is one of the lowest levels of the same world you refer to as "concrete."

The next level is concerned with something like "transitions," which could mean rates of change (like rate of spin) or partial derivatives and integrals -- paths from one configuration to another. The shapes of paths can be altered smoothly, as can the speed and direction with which paths are traversed. You can traverse a path partway, stop, and reverse to the starting configuration. So the control of transition-perceptions involves at least the dimensions of shape, direction, and speed. The "shape" dimension may simply be an underlying configuration perception.

Next comes "events." An event is a unitary set of transitions, configurations, sensations, and intensities perceived as a space-time package. An example is "jumping." Below the level of events, the underlying perceptions flow smoothly from one state to the next. At the event level we make arbitrary divisions of this flow into sections that we perceive and control as a single thing happening.

Above events are "relationships," which are derived from perceptions at the event level on downward. Relationships are things like on, in, beside, before, after, inside, outside, between, and so forth -- not as named, but as perceived. Control of relationships is involved in most behaviors. The means of control is to vary reference signals for events, transitions, configurations, etc.

Above relationships are "categories." This is the first "digital" level: all the levels below are basically analog. At the category level we perceive different things as examples of the same thing: we perceive dogs instead of

individual instances of dogs. And at this level, I believe, we begin to symbolize: substitute one representative perception for a class of perceptions. The "representative" perception can be arbitrarily chosen: a representative perception standing for many different configuration signals that look different but are classified as the same might be the configuration of marks that looks like this: "dog." A word is simply a perception used as a symbol for -- used to indicate a category of -- other perceptions, the symbol in this case being a visual configuration perception. Any perception can be used as a symbol for any other perception or class of perceptions. I am not, by the way, very satisfied with the definition of this level, particular the process of naming. There could be a missing level.

The category level, once defined, leads to a re-evaluation of the lower levels: we realize that lower level perceptions, in themselves, are neither names nor categories. One of the difficulties in parsing experience into levels of organization is that we often apply an inappropriately high level of perception in trying to grasp the nature of a lower level. A configuration perception, for example would not be "a dog." It would be that configuration, directly experienced, that we put into a category with other configurations and refer to with a name, "dog." I think you allude to this problem in your paper.

Above categories we have "sequence," or "ordering." I think this is also what Common Lisp users mean by a "list." It is not the elements of the list; it is the sense of "listness" or ordering itself. It is a perception standing for a set of lower-order perceptions with regard to their sequence of occurrence. It is NOT a "program," because it contains no choice-points. A sequence is like a recipe: break two eggs into a bowl, stir well, add milk, pour in frying pan, add bread, etc., with the final element being called French toast. The elements of this sequence are categories of relationships among events consisting of transitions from one configuration to another, all built out of sensations having variable -- and controlled -- intensities. There is control at each level, but the highest level of control involves assuring that the perceived sequence is of a particular recognizable kind.

Category-names in sequences become the elements of "programs." A program is a network of choice-points. To perceive a program is to perceive a particular recognizable network: not any one path through it, but the entire module with all its branches at once. Each element in the network can be anything from a sequence, a list, on down. This is the main level, I think, where "abstract reasoning" takes place (although of course the elements with which reasoning deals are sequences of symbols for categories of ...).

Above this level (!), I believe, is a level at which we perceive "principles." Other words might be "generalizations" or "heuristics." These are things that human beings have no trouble recognizing and controlling for, but which we have as yet not succeeded in getting hardware to do. We can generate programs that are EXAMPLES of principles (successive approximation, for example, which you mention), but those programs are not the principles. Similarly, our names for principles are really names for lower-level situations that constitute instances of principles, as a particular set of sensations is an instance of a configuration, with other sets of sensations being instances of the SAME configuration.

And finally, at the top (as far as I know now) we find "system concepts." These are things like "physics" and "government" and "AI" and "self." They are entities perceived as functions of sets of principles etc. The system concepts for which we control determine what happens at all lower levels -- in general, although not, of course, in detail.

These levels were defined on the basis of subjective experience, but also meet some communicable criteria for a hierarchical control relationship. A perception at any level, if analyzed into elements other than smaller perceptions of the same type, proves to consist of sets of perceptions of the next lower level and of a different type. This is a subjective call, of course, and my analysis might not exactly match someone else's. But so far there seems to be pretty good agreement with others who have looked

critically at the same aspects of experience. I expect all the definitions to change, eventually, as we explore them experimentally.

The other criterion is that in order to control a perception of any given level (act to bring it to a specific state), it is necessary to VARY the target-states of lower-level perceptions. To alter the visual configuration we call (at the category level) "squareness" to make it a little less square, we must alter the sensations that constitute its sides and corners. CONTROLLING any given level of perception requires VARYING lower level perceptions.

I think that my definitions of levels meet these criteria. The only way to check this out, of course, is to look for yourself.

You have probably noticed that in this hierarchy of perceptions, the entire world of experience, everything from the most concrete stimulus intensity to the most abstract system concept, appears as a perception in the brain. The "outside world" doesn't come into it at all. When you lean your bicycle against the wall, you're controlling one configuration perception to bring it into a specific perceived -- but not necessarily named -- relationship with another configuration perception. When you worry about how to lock the door without letting the bike fall and spill groceries everywhere, you're sorting through sequence perceptions, trying to find one that will work (in imagination, a subject we'll skip but that is in the model). And the sorting is done in terms of the NAMES of CATEGORIES of lower-level perceptions, these names becoming symbols that are handled by some sort of logic, under control of principles such as "don't blow it."

What's going on in the outside world while you're controlling all these levels of perceptions is a good question. I think it can be answered only in terms of models of possible realities. What we experience consists of neural signals.

Well, in a very small nutshell, that's HPCT. I haven't talked about the logic of control, or the kinds of experiments one does to establish what in fact is being controlled with respect to what reference state, but perhaps this is enough to tell you that we may have some common interests. I've probably given the impression that the theory is much better developed than it really is, particularly at the higher levels. But in Big Picture terms, perhaps you get the point. In a phrase that I'm trying to discourage the use of, because it's turning into a slogan, it's all perception (and control of perception).

Control theory says that control systems VARY their actions in order to CONTROL their inputs. Not their outputs. What others see as controlled output -- as behavior -- is really just an indirect effect of controlling perceptions. Another way to say this is that control systems control OUTCOMES rather than MEANS. This is why some of your buddies at MIT are on the wrong track: they're trying to build models of motor behavior that specify outputs, where the real system works by specifying inputs. They're forgetting that between muscle tensions and their final effects are many other unpredictable influences that also contribute to the outcome. Regular outcomes can be produced only if they are sensed, and if control is centered on matching what is sensed to some reference state. To produce the same outcome twice, in the real environment, you must NOT produce the same outPUT twice.

As you can guess, HPCT has a lot to say about AI. And a lot to learn from it.

If you want to look in on our list, the listserver is at listserv@vmd.cso.uiuc.edu (U. of Illinois) Send the message (to the above address, not to me, as you no doubt know) SUBSCRIBE CSG-L lastname, firstname, location It's an open forum, and pretty active (a megabyte per month, sometimes). You might find any subject at all being discussed, but all in terms of control theory. Don't hesitate to start a new thread -- or to just listen if that's your preference.

I think you may find HPCT a great tool for saying all those good things you have to say.

Best, Bill Powers

Date: Mon Aug 24, 1992 8:38 am PST
Subject: From Intensities to Sensations

[From Rick Marken (920824.0930)]

Gary Cziko (920824.0010) asks

> With nine intensity signals, why does "nine independent sensations signals at a given time" mean? Do you mean nine sensation signals at the same time?

Yes; Nine linear combinations of nine variables can be INDEPENDENTLY and simultaneously brought to nine different values. It's just basic linear algebra -- you can solve, at most, N simultaneous linear equations when you N variables (and, of course, N unknowns).

> With nine intensity signals ranging from zero to some maximum neural frequency it seems to me that there would be a lot more than just nine sensations possible just using weighted sums.

Yes - there are an infinite number of ways of combining the nine inputs linearly to produce sensation signals. But you can only control nine of these linear combinations simultaneously (that's why there are only six sensations in my excel spreadsheet -- because there only six intensities; but you can still find sets of 6 linear combinations that are not orthogonal (they are linearly dependent) so you don't really have 6 independently controllable sensations. This is how you can have conflict at level N+1 even though there are plenty of available inputs at level N).

penni sibun says (re:Preston paper):

> you saw the abstract. she started off w/ saying that behaviorism and cognitivism are looking at different aspects of the same thing.

Well, memory is the first to go (luckily). When did I see the Preston abstract? Is it in an earlier post. It does sound like she is saying what I am saying (behaviorism and cognitivism are looking at different aspects of the same thing -- but what is that thing, according to Preston?)

Best regards Rick

Date: Mon Aug 24, 1992 11:12 am PST
Subject: Degrees of freedom; Unhappy pasts

[From Bill Powers (920824.1100)] Avery Andrews (920823) --

> First, I am very skeptical that people normally perceive the wheel angle.

In HPCT, all that is required in order to say that a perception exists is that a neural signal exists in a perceptual pathway. This has nothing to do with consciousness. In a spinal control system, the signals standing for muscle stretch and tendon tension are present at all times, at some magnitude. Copies of those signals rise to the brainstem, the cerebellum, the midbrain, and (by some direct pathways) to the motor cortex. The control systems involving those perceptions are always active. But we are seldom conscious of those signals unless we deliberately attend to them, or something goes wrong that draws our attention to those levels of organization.

Second, while taking a long drive yesterday I noticed something that doesn't fit into my story as told, which is that on long distance drives at least, most of my (and my wife's, as far as I could see) steering movements were quick & small adjustments, whereby the wheel was turned various distances, but at what seemed subjectively to be a pretty uniform rate. These motions seemed to be completed before there

> was any noticeable change in the heading of the car. I don't yet have a real story about what's going on here, but I think it involves perceiving the car-heading to be wrong, ordering up more-or- less enough path-curvature to change it quickly enough, and then repeating this to straighten the car out again.

There are other controlled variables at lower levels, important ones being the senses of sideward and rotational acceleration that indicate the start of a movement of the car. If something accelerates the car to the left (a bump), you feel an acceleration to the left as the car presses sideways against you, accelerating your body toward the left. You immediately turn the wheel to the right, reducing your body's sideward acceleration (which you feel as a force). This doesn't completely prevent the car's direction from changing, but the higher- level systems based on vision can correct the residual error. So one of the lower-level reference signals that is set by the higher-level driving system is "zero lateral acceleration."

An interesting demonstration of this sort of effect shows up when you accelerate the car with the foot-pedal. At the same time that the speed-control system increases the reference level for visually- detected speed, it raises the reference signal for forward acceleration, which you feel as your stomach muscles tighten to force your trunk forward and thus keep it stationary relative to the seat and steering wheel. If the transmission happens accidentally to be in neutral, the reference signal for greater forward force causes your body to pitch forward, as if someone had slammed on the brakes.

Normally we are conscious as though from the viewpoint of a higher- level system -- which level depends a lot on the level you habitually adopt as a point of view. People like us spend a lot of time doing logic and verbal manipulation, so very often we are unaware of the workings of the lower level control systems. They operate, however, just as well without awareness, and possibly better.

So you're right in being skeptical about people perceiving wheel angle. If, however, you asked them what the wheel angle was, they could tell you by paying attention to the positions of their hands or by attending to the part of the visual field where the steering wheel is. This doesn't mean that wheel angle isn't being controlled even when they're not attending. If you're the passenger, just reach out and tug at the wheel -- you'll feel resistance from the driver even before the car has begun to deviate, and before the driver yells at you. If you keep your disturbance small, in the range of normal disturbances that arise from little irregularities in the road, the driver

might not even notice -- but the disturbances will still be resisted. That tells you that wheel angle or at least angular velocity is under active control by a system that's currently not in awareness, but is still a necessary level in the steering hierarchy.

I'm delighted that you're thinking about control theory while you drive. You'll learn a lot more about it by sorting out real control experiences than you will writing model programs. Just don't forget to attend to the higher levels every second or so!

You'll notice an implication that we can attend to perceptual signals that are not at the highest levels in the brain. I'm not sure that's true, but it seems to be true. ----------

Gary Cziko (920824.0010) --

This one seems to have kept you up pretty late.

> I get especially nervous when I don't understand something which you preface with "clearly."

Roger. I get the same feeling in reading philosophy when I run across italicized words. Invariably, these are the words most in need of definition and most lacking in it. Example: intention is a sense of _aboutness_.

> With nine intensity signals, why does "nine independent sensations signals at a given time" mean? Do you mean nine sensation signals at the same time? With nine intensity signals ranging from zero to some maximum neural frequency it seems to me that there would be a lot more than just nine sensations possible just using weighted sums.

Yes, there could be jillions of different sensation signals. But only nine AT A TIME can be controlled with respect to independently-set reference levels without creating conflict. This is because independent control amounts to solving simultaneous equations. You have nine (perceptual) functions of the nine intensity signals that must have specified values at the same moment. With nine equations in nine unknowns, a solution is possible if the perceptual functions are linearly independent. If you try to control 10 or more functions of 9 variables, you won't find a solution.

More than nine sensation-signals can be present at the same time, even with only nine different intensity signals at the level below. But only nine of them at a time can be controlled relative to specific reference levels.

> Consider at all the colors we perceive with only three (I think) intensity signals related to color. What am I missing here?

This has worried me, too. If we just look around at the environment, we see an incredible number of hues of color. Is there a control system for every hue? Part of the answer is in the fact that when we ADJUST color (for example, with a TV set's color controls), we attend to only one place in the visual field, around the center of vision. It's as if we can exert active control only for what is in the center of attention. All the other perceptual variables at the same level just sit there in whatever state we left them (if they'll stay that way -- otherwise they drift). When a lot of variables of the same type need controlling simultaneously, you get the one-armed paper hanger effect (OAPHE). Our attention jumps around among the variables, and we have to switch our outputs from one reference signal to another, trying to keep them all appropriately set as multiple disturbances upset all the variables.

This is clearly an important part of an HPCT model that doesn't exist in the present form. In a way it makes the modeling job easier -- it says that we don't need a separate control system for each parallel instance of a given type of perception. We seem to be able to keep some small number -- 7 plus or mine 2? of control processes running in parallel at the higher levels, at least within consciousness, but not more. At the lower levels, we can

apparently run far larger numbers at the same time, but mostly because we don't need to be aware of them.

At or above some level, it seems that we can set up control processes to run at the same time, but without attention they tend to decay and drift. It seems to me that this is an area where we need a lot of experimental data. We can measure control parameters rather easily and quickly. Wouldn't some research group like to investigate how control varies with attention? There are gobs and gobs of vital information to be obtained here. This is a high-priority subject. Before we try to construct models that can do this sort of "scanning" process, we have to know what phenomena need modeling.

Best, Bill P.

Date: Wed Dec 02, 1992 6:46 pm PST
Subject: Where are reference signals and error signals?

[From Bill Powers (921202.1730)] Audra Wenzlow (921202) John Gabriel (921202)

The puzzle you are working on is the same one I went through in developing the definitions of perceptions in the hierarchy. I expect that Rick Marken will be along with a message similar to mine shortly.

Audra says:

> For instance, in the rubber band experiment, I don't really care where the knot is, only how far I perceive it to be from the reference point. In other words, I am not controlling my perception, but the difference between my perception and my reference level.

When you describe it this way, it seems that the knot's position is the controlled variable, while the position of the spot where you are keeping the knot is the reference signal. This creates a puzzle, because in the HPCT model reference signals pass downward from higher systems into the comparators of lower systems, and don't come in through the senses. Only perceptions originate in the senses. So is a reference signal also a perceptual signal? This leads immediately to

> I understand that how I see this error is also a perception, but the distinction between "the perception of how far away my perceptions are from my reference signal," and my perceptions themselves should be made.

So now we also have error signals originating in the senses and becoming perceptual signals. If you try to draw the control-system diagram so that not only perceptions, but reference signals and error signals are inputs from the environment, you will soon end up in a conceptual mess.

This is not a trivial problem; you are astute to have pursued it to this point.

Behind my answer to it lies a fundamental postulate of HPCT, which is that the world we experience exists ONLY in perceptual signals. We do not perceive error signals. We do not perceive reference signals. The original reason for proposing the imagination connection was precisely to provide a way to get the information in a downgoing reference signal, which is not perceived, into the perceptual channels where all perception takes place. Only in this way can we maintain consistency with the postulate that all experience is of the perceptual signals, and still explain how we sometimes -- I stress sometimes -- can know the reference condition directly. And I believe that this postulate is essential in maintaining the overall consistency of the model with experiment and experience.

If the postulate is true, then the target spot where the knot is supposed to be is not a reference signal. It is a perception. So is the position of the knot; that's another perception. The answer to the puzzle is now staring you

in the face. To show you what it is, all I have to do is change the instructions (or you can do it yourself): keep the knot 4 inches to the right of the spot.

This makes it clear that you are not controlling the knot in isolation. You are perceiving both the knot and the spot, and you are controlling a RELATIONSHIP between them. The reference condition for this relationship, one assumes, often unconsciously, is "knot over spot." But that is just one possible state of the relationship; if you always pick that relationship, you will not realize that other reference- relationships are possible, and the role of the spot will seem ambiguous. In fact, the reference condition can be any state of the knot-to-spot spatial relationship.

When you pick a target position like "knot 4 inches right of the spot", you can now realize that you do not perceive the reference relationship. You perceive only the actual relationship. If the knot is 8 inches right of the spot, that is what you perceive, and nothing else. You do not see a knot that is 4 inches to the right of the spot; the only knot you see is 8 inches to the right of the spot (alternatively, the spot is 8 inches left of the knot). You "know" somehow that the knot and spot are too far apart (or too close together), but you have no picture of the correct relationship in your perceptions. This sense of "knowing" that what you see isn't "right" is as close as you will get to perceiving the reference signal or the error signal with your eyes open. The only way to get closer is to close your eyes and visualize the knot in the relationship you mean by "four inches to the right of the spot." Now the reference signal is routed into the perceptual channels, and you are perceiving the reference condition. Not everyone can do this easily; some people seem unable to do it at all with visual images.

As soon as you open your eyes, the imagined relationship is replaced by the real one; the knot is now too far from the spot. You are no longer imagining the reference condition, but perceiving the actual "wrong" condition. As you act, the sense of wrongness diminishes and finally vanishes -- but you are never perceiving anything but the actual relationship.

This is why we have a model. The reference signal and error signal in the model are not part of experience. They are an explanation for how action and experience come to be related as they are. We can't verify their existence by looking at experience, because all that experience contains is a perceptual report on the actual current state of affairs. We can only test the conceptual structure of the control system indirectly, by showing that it accounts for what we observe. Once in a great while we can trace out some neural circuits, like those of the spinal reflexes, and show that the physical architecture is consistent with the model.

Relationships are the most difficult types of perceptions to understand in the PCT model, because when we think of how the model works, we are using our own relationship perceptions very heavily. The process of comparison involves a relationship between the perceptual and reference signals: perception smaller than, equal to, or greater than the reference setting. But that relationship is detected automatically, outside the purview of direct experience. It occurs at EVERY level, not just the relationship level; spinal motor neurons carry out this process of comparison while controlling mere intensity signals. We have to distinguish carefully between the behavior of the model, all parts of which we view in the mind's eye using all our natural levels of perception, and our experiences of the world and our actions on it, which we view directly. Just remember that in this model all experience is perception, and all perception is the output of a sense-organ or a higher input function. Reference signals move in the opposite direction, outward or downward, and save for imagination do not appear in the perceived world.

It's ALL perception.

Best, Bill P.

Date: Mon Feb 01, 1993 10:38 pm PST
Subject: Levels; feedback too slow; open-loop models

[From Bill Powers (930201.1900)] Bob Clark (930131) --

> I am leaning toward designating Seventh Order as the Order including perceivable variables off Personality. This would imply considering Eighth Order as pertaining to Character.

I'm glad you put it that way, because it brought into focus a difference between the way you're characterizing higher orders of organization and the way I'm thinking of them.

To speak of "personality" and "character" is to take an external view of someone else's organization. That is, you seem to be looking for levels that will apply to "psychological" aspects of a person, to explain the how and why of that person's behavior.

I'm taking a different viewpoint: my definitions of levels are meant to describe how the world appears from the standpoint of the person regardless of the context. When I speak of "system concepts," I'm referring not just to things like a self or a personality or a character, but to ALL system concepts. To a physicist, for example there exists something called physics, a discipline. This is, of course, a perception. The entity called physics, I have proposed, is a concept build from a set of principles and generalizations, which both provide the material within which the entity physics is perceived, and which, as goals, are specified by the goals we have for physics -- that is, for what kind of entity we want it to be.

The principles and generalizations, in turn, are built out of a set of rational, logical, reasoned mental processes that I call, generically, "programs." In a set of programs we can discern general principles; at the same time, the principles we wish to maintain in force determine what programs we will select to use.

My intention in proposing these levels of perception was to provide a framework within which we might understand all human experiences, no matter what they are about. If the subject matter is one person's experience of other individuals, then what I call "system concepts" would correspond to what you term "personality," and perhaps what I call "principles" would correspond to your "character," and my "programs" to something like "habits or "abstract skills" or "intelligence." These are ways of perceiving other people.

But these general classes of perception and control include more than our experiences of other people. As I said, they include all experiences of all kinds. To a manager, the system concept called "my company" is as much an entity as "my children." To a patriot, "my country" is a real living entity. To a sociologist, "society" is a system concept with as much reality as "self." And to a chemist, chemistry is an entity with characteristics that depend on principles that are implemented as programs, without any organisms in the picture.

So what I am most interested in are the general classes of experience, not specific contexts in which we might give them more specialized names. The concepts of "character" and "personality" are inventions, but they are examples of fundamental classes of perception shared by the educated and the uneducated alike, and constant across cultures (I sincerely hope).

Date: Tue Feb 16, 1993 3:52 pm PST
Subject: My levels and Bob Clark's

[From Bill Powers (930216)] Bob Clark (9230205 and later) --

I wasn't accusing you of beginning with psychological constructs and then filling in lower-level systems. My point is different.

Some time between 1960, when we parted company, and 1973, when I published BCP, a change in my thinking about the levels seems to have occurred. Or maybe, being on my own, my direction of thought became clearer. This all seems to be clearer now that you're describing your hierarchical concepts once again.

At any rate, the "pre" idea was much like yours, that we were attempting to characterize human beings by identifying levels of control with various aspects of human functioning. Somewhere in that 13 years, I realized that this was not the right problem.

As I now think about it, the problem in understanding human nature is not so much to understand human beings as to understand the world that human beings experience. In this world I include not only the three-dimensional world around us, complete with living color, stereo sound, smellivision, and so forth, but also the "inner" world of imagination, memory, thought, reasoning, understanding -- the whole world of inner commentary on sensory experience. In short, the world of experience includes everything experiencable, whether we think of it as being "inside" or "outside."

This world, to the best of my knowledge, originates in signals emitted into the nervous system by sensory receptors. That observation seems fundamental to me; to deny it would be to wreck the entire structure of physical theory, which I do not propose to do just yet. There is no way for the state of the world outside the nervous system to be registered in the brain without first appearing as a set of raw unanalyzed sensory signals. Nothing by way of information about the outside universe can get into the brain in any other way.

This means that the world we experience must consist of sensory signals and other signals derived from them. The "other signals derived from them" include the totality of what we can experience, from the taste of chocolate to Fermat's Last Theorem, as well as our experienced "interest" in that Theorem, if any, and any "thoughts" we may have about it. Nothing is exempt.

When I say "it's all perception" this is what I mean. We live inside a nervous system and all we know is what goes on inside that nervous system. Even our idea of the existence of the nervous system exists as a set of neural signals, perceptions. The physical world outside us is a network of hypotheses existing in neural networks in the brain. Part of this neural hypothesis is a conjecture to the effect that there is an objective physical world outside our sensors. Sciences like physics and chemistry are very well worked out neural hypotheses. At bottom, they rest on sensory experience and all that the brain can make of such experiences. Our very attribution of physical theory to objective phenomena is itself a phenomenon in the brain.

This changes the problem. Now the problem is to classify all of experience, not just experiences of other people. We may perceive another person driving a screw into a piece of wood as showing a "skill" type of control, but this leaves unexplained the screwdriver, the screw, the piece of wood, and the relations among them. Those are also perceptions, and they are being controlled. The term "skill" refers mainly to something about the person's organization, but to explain how a skill like that is carried out we have to explain the screw, screwdriver, wood, and relationship as well. The perceptual organization needed to represent these four things explains their existence for the actor; the actor's behavior is explained, in PCT, as control of these perceptions. Whether we characterize that control as constituting a "skill" is more or less beside the point. If we can explain the behavior in terms of controlling perceptions of wood, screw, and screwdriver individually, and in terms of adjusting those controlled perceptions to maintain control of a

particular space-time relationship among them, we have explained "skill," too. But we have also explained how any person interacts with the world, whether the immediate world contains other living systems or not.

What I attempted to do with my definitions of levels was to represent the way the world seems to appear to us -- meaning, to myself as a representative human organism. This was very much an idiosyncratic first try, and it has undergone revisions as I have attempted to refine the descriptions. The process involved was quite unscientific, in that I didn't take any polls or do any objective experiments. I simply looked and listened and felt, and tried to understand what was going on from the standpoint that I was an observer watching the outputs of neural data-processing functions. "What am I taking for granted?", I asked over and over. What is it that I'm doing or experiencing that is so familiar and so self-evident that I don't even recognize it as a perception? What part of my experiences am I setting aside as having some special status, or treating as the background of more important things, or brushing out of the way so I can look at something more interesting?

The "relationship" level was a latecomer to the hierarchy. I had spent a lot of time looking for relationships between one perception and others, and between action and perception, but it took years for me to realize that relationship ITSELF is a perception. The same was true for all the levels added or modified since 1960. I had spoken for years about the "principles of control," without realizing that principles can't exist unless we perceive them, and to perceive them we necessarily have to have principle-perceiving functions. Similarly for "physics." What is physics, that I can know it exists? It's a perception, of course. If I couldn't perceive such a thing, it wouldn't exist for me. So what sort of thing is it? I have proposed calling such things "system concepts," for lack of any better term. And what other sorts of experiences are of that same sort? There are many, once you realize that this IS a sort of perception.

I think that the key to understanding how I think of the levels is to get into a mode of observation in which, as they say in Washington nowadays, "everything is on the table." No thought, no concept, no background perception, can be let go because it "doesn't count." Everything noticeable counts. Everything noticeable is evidence about what at least one brain is doing. If you accept the basic premise, that the experienced world begins as a set of unanalyzed sensory signals, the only conclusion is that everything noticeable is activity in a brain, and hence has to have a place made for it in a model of a brain.

I don't think that I've characterized the higher levels of perception very well. The most I hope to get across by the terms I use is the approach, the idea of calling into question everything we normally take for granted, all the operations and perceptions that we use in thinking about and acting on something ELSE. I don't think we'll arrive at a consensus on the levels until more people go through this very personal sort of exploration and report their findings.

Best, Bill P.

Date: Tue Oct 12, 1993 10:59 pm PST
Subject: Clark's levels

[From Bill Powers (931012.2000 MDT)] Bob Clark (931012.1730 EDT)

> I've made several attempts to address your "difficulties," but I find I'm trying to guess the sources of the problem(s).

My basic problem, Bob, is that you seem to be classifying the names of perceived behaviors rather than the basic neural/mental processes required to produce those behaviors. I don't see "Mechanical Skills" or "Person Skills" as defining basic mental (or control) processes, but more as classifications of control behaviors according to the context in which they occur.

If there are basic types of perceptions and control systems to go with them -- an unproven but attractive hypothesis -- then we would expect to see the very same types of perception and control in all contexts, whether a person is dealing with things or concepts, inanimate nature or other living systems. I think it likely that ALL the levels of perception and control are involved in ALL behaviors of an adult person, regardless of context or subject matter.

Consider the level I call "relationships." Relationships are typically referred to with prepositions: on, beside, under, after, because of, inside, and so forth, as well as by quantitative terms like greater than, equal to, symmetrical with, etc.. You can say that the potential energy of a rock at the top of a mountain is greater than that of the same rock at the foot of the mountain, a statement containing one major relationship perception -- greater than -- and several subsidiary ones. Or you can say that you feel more at ease when your friends are supportive of each other, a much more complex relationship but still a relationship. In my scheme, both require the ability to perceive (and often control) perceptions of the same level: relationships. This type of experience appears in all contexts.

You classify perceptions having to do with physics as belonging to Mechanical Skills, and those having to do with how people get along as People Skills, so in your scheme there is nothing in common between the implied control processes -- they are at different levels simply because the subject matter is different. Yet I see the capacity to perceive relationships as essential in either context.

To me, this means that there isn't anything unique about your classifications. People skills vs. mechanical skills could be covered by slicing the pie differently: for example, one-actor skills vs. multiple-actor skills, or skills in dealing with predictable versus unpredictable behavior of other entities. The basic problem with verbal taxonomies of behavior is that we have so many different ways of verbally classifying the same things that the same territory can be covered with many alternative schemes, all of which seem to apply perfectly as soon as you think of them.

I don't know if you remember our earliest days in the second subbasement of the Argonne Cancer Research Hospital, before we moved to the V.A. Hospital. We had a portable blackboard on which, at our weekly meetings, we started listing terms for perceivable and conceivable things, shuffling lists of words around, looking for some natural kind of hierarchy. But every time we would come up with what seemed a reasonable scheme, one or the other of us would come back the next week and say "Wait a minute -- this would make just as much sense if we took this top level and stuck it in the middle, and moved the bottom up to the top." It all depended on what you were thinking of as examples. Is nourishment a subset of organic materials? Or is it that organic materials are a subset of nourishment, when you consider not only human beings but bacteria? Are objects an example of sensations, or are sensations an example of objects (of awareness)? Or is awareness an object, since we can perceive it? Or can we perceive it?

We scribbled and erased and interchanged and substituted for months, and eventually the whole project just collapsed. Only when it collapsed did I realized the nature of the problem: we were trying to think of a hierarchy of control based on words, whereas the real system had to be based on the perceptions that the words were trying to indicate: on neural signals that stood for experiences, not on the names we gave to those neural signals.

That's when I started to learn about analog modeling, analog computers, analog simulations (although then the word was analogue). I learned how signals could be made to depend on other signals via computing functions, without ever being converted into symbols. I began to see how all these control processes could work directly with the signals, how the signals themselves must be what we experience and control. Even the words we were shuffling around were just more signals, attached arbitrarily to other signals. I saw that to understand the control processes we had to look beneath the words, directly and in silence, at the analog processes that took place without any need for symbol manipulation as an intermediary.

I bring this up now because I have never been sure that you shared these realizations; neither am I sure that you share them now. The way you're approaching the definitions of higher levels takes me back to our days scribbling on that blackboard looking for the right words. It's likely that I was never able to articulate what I had seen about words versus perceptions -- it wasn't until shortly before completing BCP that I first tentatively thought of the category level as a place where words could enter the picture as names of categories, 11 or 12 years after we parted company (and when I realized that categories and names, too, are perceptions). And I never have been able to communicate well the way in which one looks past the words at their meanings, which are the perceptions themselves. This concept, which is at the heart of PCT, doesn't seem to get across to many people when I try to explain it. So I can't blame you for, seemingly, not getting it either.

In my attempt to define levels, I have tried to find types of perceptions first, then terms that seem descriptive of them. I have been looking for simple obvious things -- obvious once you manage to notice them as aspects of perception instead of projecting them into a taken-for-granted objective outside world. I have been looking for types of perception that are so fundamental that they appear in all of experience, no matter what you are doing, no matter what kind of environment you're in, and no matter what kinds of systems you encounter in that environment. And no matter what you say about them, or call them. I'm not totally sure that the hierarchical idea is right, or if it is, that I have identified all the levels or got them in the right order. But the one thing I am sure of is that PCT is about controlling the world of direct perception, not words about perceptions (except of course when we use some perceptions as indicators of others).

> For the most "general" interpretation of the verb [generalize], My Dictionary gives: "generalize, -ized,-izing. v.t. 2. to infer (a general principle, trend, etc) from facts, statistics, or the like." And also: "infer, v., to derive by reasoning; conclude or judge from premises or evidence."

If you start with "general" (in my dictionary), the first definition is "of or pertaining to all persons or things belonging to a group or category." Just to follow one trail, "category", we find that the first two meanings refer to a classificatory system, or a basic classification of terms. "Classification" leads to various forms of "class", and under "class" we find "1. A number of persons or things regarded as forming a group by reason of common attributes [etc.]." A "group" is "1. any collection or assemblage of persons or things; cluster; aggregation." And "collect" is "1. To gather together,assemble."

You can go on from there with other branches, starting with other key terms like "attribute" or "all". The circles soon become very tight; the dictionary runs out of words. As you keep following the trail through the key words, you feel that you're getting close to something basic, but just as it seems you're about to get there the dictionary loops back into itself and you're back where you started.

The dictionary is of no use at all in helping us to understand the experiences that all those words try to communicate. When you follow up on definitions of definitions in the dictionary, you always end up in circles, and the circles always occur at the interface between words and perceptions. After following any chain as far as you can, it all comes down to either knowing what perception a word indicates, or not knowing. No dictionary can help you over that barrier. Either you understand what a category or group is -- that is, you can recognize one when you perceive it, before you know what name to apply -- or you don't. If you can't recognize a group or assemblage, you can't know what "group" or "assemblage" refers to. What such terms refer to is not a word, but a perception.

> From this it seems that "Generalizations" do not necessarily all belong to any one level.

This can also be interpreted the other way around: what you are defining as "levels" are each actually entities formed from words referring to many basic levels of perception, one of which is the ability to perceive generalizations, or as I call them, principles. The fact that generalizations or principles appear as a necessary component in many of your levels suggests that the basic perception is that of the principle level, for the opposite does not hold: Neither "People Skills" nor "Mechanical Skills" appear across sets of principles, because the term principles applies in other contexts like "Things equal to the same thing are equal to each other" and "What goes around comes around," which may or may not be applied to people or nonliving systems. If you can't perceive in terms of principles, you can't generalize either about interactions among people or interactions among physical variables. The capacity to perceive in terms of principles is fundamental in all contexts. "People skills" are not.

Best to all, Bill P.

Date: Tue Mar 07, 1995 11:59 am PST
Subject: category perception

[From: Bruce Nevin (Thu 950306 11:08:16 EST)]

To each level of the perceptual hierarchy, the level below it is the environment.

On the level of intensity perceptions we live in a world of sensory inputs from the environment beyond our sensory organs.

On the level of sensation perceptions, we live in a world of intensities.

On the level of configuration (or transition) perceptions, we live in a world of sensations.

On the level of transition (or configuration) perceptions, we live in a world of configurations (or transitions).

On the level of event perceptions, we live in a world of transitions (or configurations).

On the level of relationship perceptions, we live in a world of events.

On the level of event perceptions, we live in a world of relationships.

On the level of category perceptions, we live in a world of events.

On the level of ...

Now wait a minute. The experience of yellow light is a sensation perception. Is the category "yellow" about sensation perceptions like this, or is it about some event? If an event, how?

One might make a series of statements. If he repeatedly and variously undoes what I do, I am disturbing a perception that he is controlling. He repeatedly and variously undoes what I do. Therefore I am disturbing a perception that he is controlling. This is an experience on the level of logic or "programming". Is the category "modus ponens", of which this argument is an example, about logic perceptions like this, or is it about some event? If an event, how?

Am I overreaching? Is this picture of the relationship of levels too strictly ordered? Or does it support what Martin is saying about category perception?

Bruce

Date: Tue Mar 07, 1995 1:17 pm PST
Subject: Re: category perception

[From Oded Maler (950307)]

(Bruce Nevin (Thu 950306 11:08:16 EST)

> Am I overreaching? Is this picture of the relationship of levels too strictly ordered? Or does it support what Martin is saying about category perception?

One man's ceiling is another man's floor. For me it is clear that the location of various categories in the perceptual hierarchy is a completely private matter. The process of category formation depends on the personal history of experience of every person (which includes all the historical and technological context). The category of, say, mouse and window (in the computer context), will be realized lower in the hierarchy of my children than in mine. What I perceive of the mountain where I live is not located at the same level as it is located in the hierarchy of a French farmer living by. I look at the mountain and the clouds with my perception standing on the shoulders of what I know of Math, vision processing, general geographical knowledge, chaos, PCT and what not.

Words, which are our greatest invention, are also the most confusing, because they denote such a diversity of perceptual variables which might have very little in common except for the letters and sound. (As noted before in this forum, this is a reason for many of the great debates in this forum).

There are probably also some genetic factors in certain properties of neurons that tend to influence the type of perceptual variables that the individual is likely to form.

Haug, --Oded

Date: Tue Mar 07, 1995 4:34 pm PST
Subject: Levels of perception; musings

[From Bill Powers (950307.1625 MST)]

Bruce Nevin (Thu 950306 11:08:16 EST)--

> Now wait a minute. The experience of yellow light is a sensation perception. Is the category "yellow" about sensation perceptions like this, or is it about some event? If an event, how?

A problem I confronted some time ago.

You have to understand how I got to these levels. It was all purely subjective. I started with configurations, I think (it was a long time ago and my only notes are in my head). Or really with the relationship between configurations and sensations. I realized that if you ask what a configuration is made of that isn't just another smaller configuration (as a chair is made of legs, seat, back etc., more configurations), what you end up with is sensations. Shadings, colors, edges, and so forth.

So this gave me the idea that some perceptions are functions of sets of other perceptions. You can perceive colors, shadings, and edges without seeing configurations, but you can't perceive configurations unless there are at least some different colors, shadings, etc.. After fifteen or twenty years, I realized that intensities are the substrate from which sensations are created; you can perceive intensities without perceiving specific sensations, but not vice versa.

In the middle of that fifteen or twenty years, I saw more or less at random other levels of perception that were related the same way. Relationship was fairly early, I think: relationships were composed of configurations, the objects which stood in relation to each other. No objects, no relationships,

but it was possible to perceive objects without perceiving relationships among them.

The problem you have noticed showed up quite early, because after a while I realized that there were transitions, which are functions of sets of configurations (as in stroboscopic motion), between the relationship level and the configuration level. You can have relationships among transitions (faster, slower) as well as relationships among configurations (bigger, smaller). This implies that the relationship level can receive signals not only from the transition level just below it, but from the configuration level, two levels below. And once you see that, you also see that there can be relationships among sensations (saltier than) and intensities (brighter than), too.

And then sequence got stuck in between transitions and relationships, and then the sequence level got changed to the event level, with the sequence level being redefined and ending up two levels above (because of some questions Gary Cziko raised, if I remember right).

What we have now is an 11-level hypothesis in which each level is supposedly dependent on the existence of the level below and is derived from it. There is no way to prove that this dependency exists except by looking at your own experiences and seeing if the idea still holds up. There's no single rule that will allow you to deduce the next level up from the existing levels; you just have to look and see what's there, in your own experience.

What all this suggests is that any level of perception can be a function of perceptions of _any lower level_. Why not higher levels, too? Well, I've tried that on, and all I can say is that I can't make sense of it. Maybe someone else can.

I think it is fairly easy to find examples at each level in which a familiar perception unpacks into perceptions of one or more levels below. So you can have categories of intensity (blinding, bright, medium, dim, black) and so on all the way up to relationships -- but not categories of categories (whatever Bertrand Russell said) or categories of sequences (orderings) or higher things.

But I make no claims to have covered all real cases or possibilities. The hierarchy is simply what I have noticed to hold true for all the experiences I've applied it to. It's a naturalistic, phenomenological system, not an orderly scientific or mathematical system. People seem to think it makes sense at least in small subsets. It does cover a lot of stages that seem necessary between thought and action. It seems relatively context-free, and to the limited extent that I have questioned people from other cultures, culture-free. But is it right, is it self-consistent, is it complete? I have no idea. I have presented it as a guess, and until we find some systematic way to check it out, a guess it will remain.

Oded Maler (950307)--

> Words, which are our greatest invention, are also the most confusing, because they denote such a diversity of perceptual variables which might have very little in common except for the letters and sound. (As noted before in this forum, this is a reason for many of the great debates in this forum).

> There are probably also some genetic factors in certain properties of neurons that tend to influence the type of perceptual variables that the individual is likely to form.

Yes. Unfortunately, language reflects the beliefs of the current and immediately past generations. We inherited a language in which behaviorism is firmly entrenched, as well as other points of view. So we can say "You're making me angry," and "I felt his grief," and see nothing wrong.

As to the genetic factors, I quite agree -- I think we are genetically inclined to construct perceptions in the classes intensity, sensations,

configuration .... system concepts, although not to construct any particular examples of these classes. On the other hand, I don't rule out anything but inherited perceptions that are peculiar to the world a person happens to be born into.

Best to all, Bill P.

Date: Thu Mar 09, 1995 9:16 am PST
Subject: Re: categories

[From Bill Powers (950309.0815 MST)]

Bruce Nevin (various posts) --
RE: categories

There is one possible solution to the category problem, which was called "order reduction" in the 1960 paper by me, Clark, and MacFarland. I think you may have mentioned this solution in passing quite some time ago, but if so it has been lost.

When you name something, such as "jump", you are linking a word- perception to another perception, so either perception can evoke the perception of the same category. After you have named a few other things like "John" and "in" you can construct sequences like John jump in.

But now you have three objects to play with, the three words. They still exist individually as low-level perceptions. You can now manipulate them _as objects_ which are independent of their meanings. For example, you can say that in the sentence "In jump John", the word "in" comes _before_ the words "jump" and "John", a temporal relationship. You can see that uttering each word constitutes an event. By writing each word you can reduce it to a static configuration (which may account for why we can read faster than we can listen). You can see transitions between the utterances. In other words, you can treat the words as you treat any other objects, perceiving them at higher levels in a way that is independent of their "meanings" -- the categories of experience that they are normally used to evoke. You can even re-categorize the low- level experiences of the words: "In jump John" is a _sentence_. You can also regroup the words, so that "In jump John" becomes a single configuration, which can be linked to an experience of a specific person doing a specific act on a specific occasion, or a class of persons doing a specific act on any occasion.

And you can do this again, and again and again (although perhaps not again). The process, I think, is called "verbal abstraction." A word is made to stand for a category of perceptions. Then the word is treated as an object in itself and linked to other words treated the same way, to create new categories which can be named ("adjective"), and so on. Each time a kind of order reduction takes place, in which a word standing for a category is manipulated along with other words at all the usual lower levels of perception, AS IF the categories were being manipulated. The relationships among the words then are treated as if they are relationships among the categories. However, they are not really relationships among categories, but relationships among _names of categories_.

It is quite possible to manipulate words in a way that suggests manipulations of categories that are not actually valid. If this order reduction has gone on several times, decoding a verbal statement may involve translating each word into examples of other words that have been classified with it, then decoding those, and finally decoding the last layer into nonverbal perceptions. Only then can the original statement be compared with experience to see if it holds up.

In the world of verbal communication this verbal abstraction goes on all the time. But I think we can lose track of how many levels of abstraction lie between our symbol-manipulations and the world of direct experience they are supposed to mean. Verbal reasoning can become an end in itself, in which words are shuffled around according to adopted rules or social conventions and are never expanded into their detailed experiential meanings. "How are you?"

When this is done systematically, you have mathematics. Solving equations written in matrix algebra notation is done by treating the matrix symbols simply as symbols which are manipulated according to special rules to produce more symbols at the same degree of abstraction. The actual manipulations take place at a low level in the hierarchy; writing down or imagining symbols in certain patterns and relationships.

In mathematics, as opposed to verbal communication, there are specific and explicit connections between the symbols at one degree of abstraction and the symbols at a lower degree. At any time, the matrix equations can be expanded into their detailed algebraic equivalents, and the expanded version can be shown also to express a true equality. It is a requirement of mathematics that manipulations at the matrix level expand into VALID manipulations at the algebraic level, and that those manipulations in turn obey the basic mathematical rules relating quantities.

Verbal abstraction very seldom obeys any explicit requirements that what you say at one degree of abstraction remain valid when expanded into the next lower degree. This is my general objection to verbal abstractions. It's not that abstraction can't be done in a valid way; only that it is seldom done in a way that can be shown to be valid when reduced to the basic connections between words and nonverbal experiences. Statements that seem true at a high degree of abstraction often prove to contain many falsehoods -- counterexamples -- when expanded into terms of a lower degree. This, I suspect, is one of the motivations for seeking statistical truths. Statistical truths are not bothered by counterexamples.

This is why science really can't be done verbally. If we were all constantly conscious of abstracting (as Korzybky recommended), perhaps words would be as useful as mathematical symbols. But we aren't; we have never developed any formal system for making sure that what we say at one degree of abstraction is literally true at the next lower degree.

I've strayed a bit from the original point. To get back to it, if we analyze what we say about higher levels of perception, particularly those levels above the category level, we will find that we are really talking about words, objects at a much lower level of perception. We are talking about the NAMES of levels of perception, and the NAMES ARE NOT THE PERCEPTIONS TO WHICH THEY REFER. The only way to appreciate the higher levels is to experience them directly and wordlessly. "Look before you leap" sounds like a principle, but it only refers to a principle. The principle being indicated can be experienced only by attending to that sense of caution and alertness that governs our actions when the principle is in effect. The sentence does not mean that we should be sure to examine the surroundings before we engage in saltation. It means that we should adopt a general attitude that can be experienced, that can be recognized, that can be alluded to in various ways by means of words, but that is not itself a word or a collection of words. It is a principle.

Best, Bill P.

Date: Fri, 26 May 1995 08:09:26 -0600
Subject: Points of view

[From Bill Powers (950526.0700 MDT)]

Bill Leach (950526.2244) -- Oded Maler (950525)

A point I have been trying to make is that the operations of the brain are neither analog operations (as we conceive them) nor digital operations (as we conceive them). Our conceptions of computers of any type exist -- the conceptions themselves exist -- as activities in our brains. Our models are made from formalized, stylized, idealized rules and symbols which we can write down and manipulate in an explicit and agreed-upon way. The ability to make such models, to manipulate symbols according to rules, is a fundamental

ability of the brain. Whatever this ability is, without it there would be no models, no symbols, no rules, no mathematics.

Whether we are talking about computer science or circuit design or mathematical analysis, what we are doing is most significant (in the realm of modeling) not in itself but as evidence of capacities of the brain. We are looking at high-level aspects of brain function and observing their products. That fact is at least as interesting as the actual details of any particular system for manipulating symbols.

Unfortunately, very few people who do circuit design or mathematics are capable of standing back from their normal occupations and describing the basic processes involved -- what goes on in their heads as they prove theorems and make models. They are much too involved in the results they are getting, in the complex structures of logic they are building, to look on such things as _phenomena_ rather than being concerned with their correctness or incorrectness or how they will develop next. This is especially true when these mental processes purport to be concerned with modeling the brain itself, the very brain that carries out these processes.

When we're considering these levels of brain function, it doesn't matter what model we're talking about, what mathematical system, what theorems or algorithms. It doesn't even matter if the manipulations are mutually consistent or inconsistent, congruent with observation or incongruent. What matters is identifying the processes that are going on. We have to get outside any particular system of rules and manipulations, outside mathematics and physics and neurology and chemistry, and try to see _what is going on here_. Only by grasping this kind of level or order of brain functioning as a phenomenon of experience can we begin to see the role it plays in the whole design of the living human system.

What I am suggesting is extremely difficult to do. As adults, we have all come to specialize in certain levels of brain function, different for different people. As the artist develops unusual skill in the manipulation of images and visual patterns, so does the logical thinker become especially adept at the constructions of complex organizations of symbol-manipulations, the politician become expert at the management of interpersonal relations, and the athlete at the control of physical movement through space and time. All people, of course, do all of these things all of the time, but some are more focused on one kind of activity than others, and have carried the processes further than others have done.

The inevitable result of this specialization is that we tend to interpret all of experience in terms of the level of perception and control that we have developed to the highest degree. The world of experience takes on an appearance that expresses the activities of one level of brain function to the exclusion of all others. I remember a story about a chess player (was it Queen's Gambit?) who at one stage of development saw the entire world as positions and moves connecting them. Even in a crowd of people at a party, this person would be thinking "If she moved left and forward, she could get into the kitchen; but if that man took a step to his right, he could block the opening."

I proposed some time ago a principle which says that you can't be aware _of_ the level of brain function that you are being aware _from_. If your awareness is operating from a base in the relationship-perceiving level, you will experience a world full of relationships, but you will not be aware _that you are constructing relationships_. Closer to the point, if your awareness is operating from a base in the level concerned with logical symbol manipulations, you will experience a world that is fundamentally logical and symbolic in nature, but you will not be aware that this is a product of your point of view. The properties of the level from which you are currently aware will be projected into the world that is the object of attention, as those properties existed independently of the perceiver.

In order to become free of this process of projection, it is necessary to find a higher-level point of view from which the characteristics of the lower level become visible. The critical insight is simply the realization that what

formerly appeared to be properties of the world are actually properties of a former point of view. The obsessive chess- player, as it were, awakens, and realizes that the people at the party are just moving normally around, and are not carrying out any strategies. He sees that he had been _imposing_ a structure on the perceived world, not _recognizing_ one.

I think it is very important to do such exercises, to become free to adopt various levels of viewpoints without becoming helplessly identified with any one of them. There are useful levels both lower and higher than the level of logical symbol manipulation. From lower levels you see aspects of experience that are omitted from the abstract logical manipulations; from higher levels you begin to see the manipulations as means rather than ends. I think that awareness of having a viewpoint is as valuable an experience as any that can be had.

Best to all, Bill P.