

Experience, Reality, and HPCT

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

Date: Sat, 17 Sep 1994 08:31:54 -0600
Subject: Re: Experience, Reality, and HPCT

[From Bill Powers (940917.0600 MDT)]

Hugh Petrie (940916.1000 EDT) --

> Well, you sucked me in at least a bit. I hope the very limited time I was able to give this is of some help.

As I expected. Yes, we put some experiences in the role of evidence, and others in the role of theory. That's the distinction I wanted, but couldn't say.

I'm going to ramble through some thoughts about theory and observation, the two kinds of experiences we're been talking about. Skip to the next post if you're getting bored with this subject.

Theory, as I see it, purports to be about what we can't experience but can only imagine (neural signals, functions like input, comparison, output functions, mathematical properties of closed loops), while evidence is about what we can experience. Both theory and evidence are perceptions, but the way we use these perceptions in relation to each other puts them in different roles.

In the behavioral/social sciences, the word "theory" seems to mean something else: a theory is a proposition to the effect that if we look carefully, we will be able to experience something. A social scientist can say "I have a theory that people over 40 tend to suffer anxiety about their careers more than people under 20 do." The theory itself describes a potentially observable phenomenon. The test is conducted by using measures of anxiety and applying them to populations of the appropriate ages. If we observe that indeed the older population measures higher on the anxiety scale than the younger, we say that the theory is supported -- or, as some would put it, the hypothesis can now be granted the status of a theory that is consistent with observation.

This meaning of theory leads to the popular statement that a theory is simply a concise summary of, or generalization from, observations. That definition has been offered by quite a few scientists past and present. I think it misses an essential aspect of science, the creative part that proposes unseen worlds underlying experience. Before the "unseen worlds" definition can make any sense, however, it is necessary to understand, or be willing to admit, that there is more to reality than we can experience.

If reality is exactly what we can experience, then there are no unseen worlds and in ways obvious or subtle every theory is just a way of describing experience. Our senses and measuring instruments indicate to us the state of the real world. A properly-constructed and tested theory, therefore, cannot be false. The only way it might be false is for some error of observation or description to be made, or for the test to contain some internal error or inconsistency.

It is this view that leads some scientists to take a rather self-congratulatory view of science. A scientist is simply someone who has learned to describe and generalize correctly. If no mistakes have been made in observation, description, or method of generalization, then the theory that summarizes these results must be correct. The personality or the wishes of the scientist play no part in this process; truth is independent of the observer.

It is this view, I think, that leads to the Gibsonian approach to perception. To maintain this view, it is necessary that what we perceive of the world be a true representation of the world. So by hook or by crook, we must find a way to show that we, as observers, look through our perceptual systems at the real world. The existence and the functions of human neural perceptual systems cannot be denied. But to accept what seems to be the case at face value would

mean that we perceive only an interpreted world, a partial view of the world, or a projection of the world through unknown transformations into the space of experience. This, in turn, would mean that all descriptions of the world are functions of human nature, and thus that all theories about the world are human theories, not ultimate truths. And it would mean that the phenomena we experience are related to the properties of the real world in ways that we can't directly perceive. This is exactly the conclusion that the Gibsonian approach is intended to deny.

More to the point, the implication would be that some elements of our theories are not really, in some subtle way, reducible to reports of observations, but are made up by human imagination. It would mean that the concept of "an electron," for example, amounts to an imagined observation, with no justification other than that assuming its existence leads to consistent explanations of experience. If this were admitted, the result would be to make science much less secure in its claims to logically-derived knowledge about the real world.

Some scientists know this; others vehemently deny it. Richard Feynman, for example, knew it. When he was asked how he arrived at his diagrams showing particle interactions, he said "I made them up." There were physicists who considered this a flippant answer, consistent with Feynman's reputation as a joker. But Feynman was quite serious. Particle physics, he said, is a game we play. It takes a sense of humor to admit that.

This same dispute underlies the controversy over whether the Heisenberg uncertainty principle describes a true uncertainty in nature itself, or a limitation on our methods of observing nature. If you assume that reality consists exactly of what we can observe about it, then uncertainty is an aspect of reality. If you assume that there is a reality independent of, and perhaps quite different from, our observations of it, then you leave open the possibility that nature is regular but our observations of it are uncertain. This was Einstein's view. I say you "leave open the possibility" because in the latter view, there can be no question of verifying the causes of the uncertainty; all we can do is make up possible properties of the world which, if they existed, would account for our observations. There is nothing to prevent our imagining that the world itself is uncertain, but that does not prove that it is. It proves only, at best, that making that assumption leads to a consistent view of the observations, an ability to predict particular observations with some degree of accuracy.

In PCT there are observations and there are theories. When I attempted to describe levels of perception, I was trying to describe observations, how the world seems to come apart when analyzed and how these parts seem to be related to each other. There is no theory intended in these proposals. It seems to me that when I see a relationship, I also see the things that are related, which themselves are not relationships. I could not see any relationship if there were not things to be related, yet I could see any of those things (events, transitions, configurations, sensations, intensities) individually, not in relationship to anything else. The only question I have is whether anyone else in the universe experiences the world in the same way. Either they do or they don't; we're talking observation here, not theory. If these are truly universal classes of perception, then every undamaged adult human being should report the same elements of experience, and the same dependencies. Again: either they do or they don't. That is a question of observation, not theory.

The theoretical aspect of PCT comes in when we try to explain why it is that the world of experience is organized in this way (if, in fact, my experiences are like anyone else's). That's when we start talking about input functions and signal pathways and control systems, none of which has a direct experiential counterpart. Of course in theorizing one tries to imagine hidden aspects of the system that might, one day, actually be observed. But today, at the time the theory is proposed, we do not observe them. We can only imagine them. And no matter how much verification the theory receives from future observations, there will always be a level of description at which we can only imagine the level that underlies it.

The same interplay between theory and observation is involved in experiencing control. You do not need a theory in order to hold your hand in front of your face and deliberately will the hand to assume various configurations. Nor do

you need a theory to tell you that what you will is very closely followed by what you then experience your hand doing. You don't need a theory to tell you that when you grasp the knob on a door, your intention is for the door to take on an appearance other than the one you are now experiencing. These are the facts, the phenomena, that we need a theory to explain.

The theory of control offers an explanation in terms of perceptual signals, closed causal loops, and mathematical properties of such systems. These entities, while perfectly experienceable in the mind, are not the experiences to be explained. We are saying that IF such an organization existed in the nervous system, THEN the experiences we are trying to explain would follow. The theory proposes the existence of entities in the world hidden from direct experience; perhaps not all of them hidden forever, but certainly hidden now.

The most important part of such theories is that they not only account for what we do experience, they predict experiences we have not yet had. The models of PCT are adjusted so that in simulation they behave in the same way as the particular instance of control behavior we're trying to explain. But once the model is constructed, we can vary the conditions that, hypothetically, affect it, and strictly from the properties of the model make predictions about how the real system would behave under those changed conditions. This is where the power of modeling shows up; not in its ability to fit the behaviors we observe, but in its ability to predict how behavior will change when we alter the conditions presented to the real system. We can fit a model to the hand motions involved in tracking a target moving in a triangular pattern, and then using the best-fit parameters predict very closely the hand motions that will occur when the target moves in a random pattern, and when a second random disturbance is applied directly to the cursor in parallel with the effects of hand motion.

I think that one main reason for the misunderstandings that occur in the life sciences about control theory is that this kind of modeling is essentially unknown to most practitioners. The idea of proposing a model that is more detailed than our observations, and then using this model to predict new observations under new conditions, does not appear in textbooks of psychology, sociology, psychotherapy, or related sciences. It is an idea with which engineers are familiar from their earliest days in college, but only where engineering has encroached on the life sciences does it appear in relation to the behavior of organisms. This method is almost the diametric opposite of generalization; instead of deriving general classes of observation that include actual observations, the method of modeling proposes the existence of more detailed variables and relationships below the level of observation, from which observations can be deduced. I have heard the term "hyothetico-deductive" used in situations that make me think of modeling, although I'm not sure that is what was intended.

Honestly, I'm almost finished.

Now think about what happens when a person who has never heard of the method of modeling comes up against PCT. To this person, the diagrams of PCT are simply diagrams of observations. The arrows show how one event leads to the next event. If this diagram describes any particular behavior, then it can be accepted as a theory (or not -- people quite often draw different diagrams, because they "have a different theory"). But such a person does not see what we see: a diagram of a specific physical system, connected in a certain way, which we can't directly observe. This person doesn't realize that what we can see is supposed to arise from the operation of the diagrammed system, not that it is supposed to be represented by the diagrammed system.

When, some day, the Center for the Study of Living Control Systems goes into operation, one of the introductory classes that must be taught there will be an introduction to modeling. It is obviously possible to teach what modeling means; all engineering students learn it, although nobody ever tells them what they are learning. They pick it up from seeing it done and learning the mechanics of doing it. They learn by osmosis the difference between describing the behavior of a system and describing the organization of a system that can produce that kind of behavior (as well as many other kinds). I think this can be taught explicitly, and that by learning it, students will not only come to grasp the meaning of PCT as it applies to human behavior, but will discover

that they can probably come up with better models than their mentors managed to build.

Wordily, Bill P.

Date: Fri, 14 Apr 1995 19:32:10 -0600
Subject: Re: physical reality

[From Bill Powers (950414.1830 MDT)]

Bruce Buchanan (950414.1800 EDT)--

A great many of my arguments will make more sense to you if you simply assume that I am NEVER talking about real reality the way it really is. When I speak of physical variables, I mean variables defined according to the models developed by the science of physics, not the variables in the real world of which physical variables are presumed to be representations. When I speak of how something "really works" I mean the way it has been found to work through physical experimentation, and in terms of the underlying physical theories. I use this as a benchmark way of discussing our shared world, because it comes the closest to being sharable.

Another reason for using the terms of physics in speaking of the world outside the human nervous system is that they are minimally influenced by "psychological" considerations. It is fairly easy to see when variables in the physics-model world, or physical world for short, are spoken of as if they carry unwarranted properties which don't belong in that model. Predictiveness is such a property: you won't find it in any textbook of physics. However we are to account for predictiveness, we can't do it by reference to the physical world, and this tells us that predictiveness is something we are adding to experience outside the physics model.

But this way of speaking is deeply ingrained and letting go of it is a struggle. In the same way it would be difficult to stop speaking of sunsets as pretty or candy as sweet or chocolate. To human beings, perceptions ARE the world. It is difficult to think of a world in which most of the entities we can experience are missing. The language and formalisms of the physics model take us partway toward that realization, by stripping away at least the higher levels of interpretation from our observations. We can never go all the way, but the more clearly we can differentiate low-level perceptions from high-level interpretations, the clearer will be our understanding -- at least by analogy -- of our relationship to That Which Is.

The hardest thing to realize is that the "external world," the "perceptual world" and "the intellectual world" are ALL IN THE SAME PLACE, in our heads. They are all ideas that exist in human experience; they are simply different subdivisions of experience. They are ways of classifying experiences, but otherwise they are all of a piece.

You say

> This statement by Bill seems to me to assume that the external (unknowable) world is identical to or indistinguishable from the concepts of physical science. I think that this assumption of identity is untenable.

I hope you will now see that this is a misinterpretation of my view. When I speak of the concepts of physics, I am treating them in their own terms, not in relation to the Boss Reality. I don't have any idea what the Boss Reality is like, and I don't know anyone who does. So there's not much point in speculating about it. All we can do is use our best models.

Best, Bill P.