What PCT has to say about education

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

Date: Tue Jul 30, 1991 5:30 am PST
Subject: Basics of PCT

[From Bill Powers (910729)] --

John Maag, Polly Brown, Ed Ford (910628, 910727) et. al.

This is generally for those who wonder what control theory has to say about education. I hate to keep saying this, but it is difficult to understand explanations unless you've spent some time with the references that have been mentioned on this net. Control theory itself is a model of how all behavior works, not just about education or psychotherapy or linguistics or whatever your field happens to be. It has an inner logic and a dynamic structure that you have to understand before you can see how to apply it in a particular context. Of course I won't just leave it there, but I hope you can find the time to do a little studying of control theory as a subject in itself. A lot of studying.

First, answering a question of Polly Brown's (who asks if control is choice). Control is acting on your environment, altering it until it assumes the shape or state you want to see (or feel or hear or taste ...). In most environments, there isn't any standard motor act that will create any specific desired result. There's no one action you can just trigger off that will put the taste of chocolate in your mouth. Your relationship to the environment is always changing, and there are changes in the environment, too; also there are other forces acting on it in addition to your own actions. So the basic problem is to VARY your actions in just the way, given current circumstances, that will end up having the result you intend or desire to experience.

Control systems are a unique kind of organization that can do exactly that. They do not produce pre-programmed outputs. Instead, they compare the current external state of affairs with an internal specification or blueprint showing how things will appear when they are "right." The discrepancy between what IS perceived and what the blueprint says SHOULD be perceived (the so-called "error") drives the action, and the action affects the environment to make the error smaller -- that is, it alters the perceived external state of affairs in the direction that makes it resemble the blueprint (or "reference signal") more closely, as the control system senses it. If you keep changing the external state of affairs to bring it closer to what the blueprint says, quite soon (in some cases, in less than a tenth of a second) the match will be essentially exact. Under almost all circumstances, it's possible to set up simple relationships between error and action that will always tend to make the error smaller no matter what the direction or size of the error.

If something changes in the environment, either due to an external disturbance or a change in the relationship between the acting system and the environment, the perceived state of affairs will also change. If this perception had been matching the reference-perception, the blueprint, now it no longer matches it. An error (meaning just a difference) has appeared. This error, as always, creates a change in the behavior in the right direction to oppose whatever is causing the change. The result is that even with random, novel, and unpredictable disturbances acting, the control system's behavior will automatically change to keep the perceived state of affairs from getting very far away from the state that the reference signal specifies. In other words, the control system resists disturbances of its perceptions relative to what it intends to perceive, and does so by varying its outputs in just the way required.

But the control system, considered as part of a larger organization inside the organism, can also CAUSE changes in the perception. This is done simply by changing the reference signal -- the blueprint. By this I don't mean changing it to mean some different kind of perception. I mean just changing its magnitude, so it calls for a different AMOUNT of the SAME perception (so "blueprint" isn't really the best image). If you're holding your hands one

foot apart, you can maintain this distance at a reference-distance of one foot even while someone tugs on one or both hands. But you can change the reference distance quite easily to two feet -- and your hands instantly become two feet apart, in fact moving even as you change the reference signal. The reference signal still means distance, but its increased magnitude now specifies a larger distance. You feel the reference signal as your intention: how far apart you intend the hands to be, even before they've become that far apart.

Changing from control of one kind of perception to a qualitatively different kind means that some higher-level system has stopped sending reference signals to one control system (or has turned it off somehow) and has started sending them to a different one altogether. I think this is more like what we mean by the word "choice." Each control system perceives just one kind of thing, always. When reference signals change, they are saying "perceive more of that thing" or "perceive less of it." A very low setting of the reference signal says you want to perceive hardly any of the thing -- that is, you want to avoid it. A high setting means you want to perceive a lot of it -- you love it. This is not choice. But saying that you want to stop controlling for apples and start controlling for peanut brittle means changing the kind of perceptions, and that means changing control systems. That is a choice. It is the means by which some higher-level control system controls its own perceptions. But I hasten to add that this is only one example of higher-level systems -- "choice" is not the generic term for higher-level control. It's just one example that shows up now and then. Most behavior doesn't involve making choices.

You might think that having a separate control system for each different kind of perception is pretty wasteful. The nervous system probably doesn't really work that way (although control systems are very simple and don't use up many neurons). But this way of modeling the system is technically equivalent to other ways, and helps to distinguish between changes of KIND and changes of AMOUNT. You won't go basically wrong by assuming one control system per kind of perception.

This concept of "controlling for perceptions" is a peculiarity of control theory that surprises some people and upsets others. Technically, it just means that a control system acts to bring its own sensory or perceptual signal to a match with an internal reference signal. If the perceptual signal inside the system always has the same precise relationship to some publicly-observable variable, then of course we could say that the system controls the external variable, too. But if the perceptual system involved changes in some way, so the same external situation leads to a different state of the perceptual signal inside the system, the system will act to bring the perceptual signal back into a match with the reference signal. This means, of course, that the external situation is caused to change -- a different state of the external variable is now required to create the same state of the perceptual signal as before. So the real controlled variable is always the perceptual signal and not the external variable for which it stands. The perceptual signal is the only variable that stays under control when the parameters anywhere else in the system are altered, including the parameters of the perceptual apparatus.

This terminology also reminds us that different people see the same situation differently -- the half-full or half-empty cup, for example. Married couples, it is rumored, can look at the same living room, and one person sees it as quite clean while the other sees it as a mess. This might result from a difference in reference signals defining "clean", or it might result from desiring the same amount of cleanness but having different perceptual sensitivity to things out of place, dust-balls, and things missing that should be there or present that belong somewhere else.

You can't understand what a person is doing without knowing what that person is controlling for, and to know what that person is controlling for, you have to guess what that person is perceiving. "Controlling for" something means both the KIND of thing that is being perceived, and the AMOUNT of the perceived thing that is wanted. When you say "controlling for" you put your attention on the perceived result that is wanted, instead of on the detailed acts the person is using for control.

This is a very important aspect of catching on to control theory. In most theories of behavior the focus is on the output, the actions that a person takes. After all, you can see the other person's actions, but you can easily fail to see which of the effects of action the person is concerned with. In control theory, the actions are unimportant, because they can vary all over the place when external disturbances tend to alter their result. Those variations only tell you about disturbances; they don't reveal the effect that is being kept constant by the variations in action. When you speak of what a person is controlling for, you remind yourself to look at the *effects* of the actions instead of the actions themselves. The *effects* are what the person is controlling -- and again you have to put yourself in the other person's position and remember that it is only the person's *perception* of the effects that is really under control. What YOU see as the effects of the other's actions is quite irrelevant, unless by luck you happen to pick the same aspect of the same effect that the other perceives.

With respect to teaching, control theory can offer some powerful hints. For example, teaching a person how to thread a needle can be done by telling the person how to hold the thread, how to lick it and twirl it, and how to move the hands to make the thread go through the hole. Or you can tell the person about perceptions: the end of the thread should look like this; when you hold the thread you should hold it near enough to the end, like this, that the thread doesn't look droopy, like that. You should hold the thread and needle so you can see the hole behind the tip of the thread; then all you have to do is move the tip of the thread toward the hole, keeping the tip over the hole all the way as it appears to you. It will then go right through the hole.

A lot of teaching is oriented toward telling people what to do. It would be much more effective if it involved telling them what to perceive. Instead of telling people to line up the numbers when they are doing long division, show them how it should look when they're finished. They can make it look that way without being told each move. Furthermore, when you show them how it should look, as a control theorist you will realize that people attend to different aspects of the same situation, and you will be sure to say WHAT IT IS ABOUT THE ARRANGEMENTS that you want them to perceive. After all, if I point to a desk and tell you "Just look at that!" it might take you some time to realize that I am pointing to a pencil out of line with the others, and not to the glass of water or the open book or the toad.

Beside telling students what the right perception is, you have to tell them what the right STATE of the perception is, and how much error is tolerable. In other words, you have to tell them what they will be perceiving when they are accomplishing (not "doing") the right thing. The more precisely they understand what they are to perceive and the state in which they are to perceive it, the less you have to tell them about how to move their arms and fingers and eyes, or how to shuffle the numbers and symbols around, or how to stir the batter. Human beings are naturally organized to discover the means of control once they have a clear picture of what is to be controlled and what the reference-state is. They have to know how they can tell when they are doing it right. If they understand exactly what a right result is, how it will look to them, they won't have to ask anyone when they have achieved it. They'll just say "I did it!" The whole trick is in knowing what "it" is.

Often, teachers assume that their job is to teach students the procedures that will result in the right end-product. This is basically a stimulus- response conception, because it assumes that if you make all the correct moves exactly as you have been told, the result can only be the right result. Of course this never happens in the real world; something always interferes that you haven't anticipated. This leads to students coming up to you and saying "I did it exactly the way you told me and it didn't work!" The natural response is to assume that the student DIDN'T do it EXACTLY as you said, and to go over the moves again to make sure they are executed properly. When the real reason shows up, you can be very embarrassed; "Oh, well, you shouldn't have tried to divide by such a small number -- you have to carry more decimal places." A glitch that you hadn't anticipated, not a wrong move.

Of course students have to learn the moves -- they have to learn how to perceive when they are doing the move the way they were told. But they also have to understand precisely how to tell when they have the right result. And

there's one more thing: they often can benefit from being told how variations in the moves affect the result -- in other words, how to alter their actions to correct differences between what is happening and what ought to be happening. This is nothing more than telling them how control of this particular thing works.

Teaching students to make the moves is a way of trying to control their behavior. "Just do it the way I told you and it will come out right." In fact there is hardly anything that can be accomplished in just one way, or that can be accomplished by the same moves every time, as every teacher of cake-baking really knows. Teaching the moves gives students the impression that the slightest mistake will be fatal. That, too, is almost never true. Children weep and moan when Daddy solves an equation using manipulations that the teacher didn't use -- that isn't RIGHT! This tells Daddy immediately that the teacher is demanding obedience, not teaching how to control for perceptions. A child taught how to control for perceptions in algebra doesn't worry about different ways of accomplishing the same result. Or so the control theorist would insist.

Mary reminds me that there is one more important thing to say about this subject. John Maag, you point out that some recommendations that David Goldstein came up with are nothing but familiar things that other people have said. It's true that this often happens: control theory leads to recommendations or interpretations that others have discovered for themselves, empirically. As you approach this from inside your particular discipline, this may make it seem that control theory isn't adding much. But from my vantage point, I can see that people in practically every discipline in the life sciences come up with similar remarks -- you're not saying anything new, so-and-so said almost the same thing in 1927.

When you hear this sort of remark from many disparate sources, you have to begin to think that there really is something to control theory: its predictions in terms of empirical phenomenology appear to be borne out by competent observers in a great number of completely different disciplines. If you're in a nasty mood, you can reply to the statement that everyone knows that giving children control of their learning is beneficial by asking "Yeah, but WHY is it beneficial?" Very few workers in any of the empirical disciplines have an answer to a question like that other than "Well, it just is." In most disciplines in the life sciences there really isn't any theory, any scientific justification for the observations the workers make. All they know is what they observe and how they interpret it. There isn't even any basis for distinguishing a correct interpretation from an incorrect one -- except another statistical study.

Don't get me wrong -- control theory *does* have some new things to say about education as well as most other fields, as I tried to indicate above. But it's equally important to know that control theory is vindicated by things that people have already discovered. There is a necessary interaction between theory and phenomenology; theory helps us interpret phenomena, but it also predicts phenomena that have to be observed. If the phenomena are not observed, or don't match the prediction, the manner of failure tells us how to adjust the model. In control theory we don't just toss off an hypothesis, give it a statistical whirl with a bunch of subjects, and if there's no result try on another hypothesis at random. We require the model to make specific predictions, and if the observations don't match EXCEEDINGLY well, we ask why not, and look for the aspect of the model that caused the mistaken prediction. We sniff and poke and adjust and try again until the model does predict as exactly as we can measure the phenomenon. Then we think up some variation that the model also has to predict, try that, and so on -- it never ends.

Control theory is already at a stage where it can save people a lot of trouble. After all, if it can predict from general principles that "experiential and discovery" learning would work better than the old kind, one wouldn't have to do so much random experimentation to discover this fact by accident. How long did it take educators to discover this fact? Maybe, once you grasp the fundamentals of control theory, there will be other such predictions and recommendations that could be derived from it without waiting a hundred years for someone to get lucky.

There's nothing wrong with discovering things that work through pure empiricism. But pure empiricism will never tell you that you have found the thing that will work BEST, and it will never help you refine something that works some of the time into something that works all of the time. For such refinements you need a good theory, an accurate model of the system you're working with. Even a television repairman has to reason out what is wrong with your set using an underlying theory, a quantitative model of the workings of a television set. We can't expect to understand a system as complex as a human being without a model of at least comparable depth and predictivity. Just fishing around at random using possibilities that pop into your head is the pre-Galilean way of understanding nature. Control theory is a post-Galilean way (although from reading Koestler I understand that we should say "post-Keplerian").

I hope that all of this helps a little.

Best to all, Bill P.