A PCT Approach to the Frame Problem

Erling Jorgensen August, 2005
Post to CSGnet

I. SITUATING THE FRAME PROBLEM: *IT'S NOT WHERE YOU THINK IT IS*

As I have tried to understand *the frame problem* from sources available on the Web, it seems to betray a key problem in how the issue is formulated. It is there in how Jeff Vancouver asked the question in a recent post to CSGnet:

> That is, somehow our minds know what changes or differences to ignore (and presumably what to send up the hierarchy). How do we do this?

This gives a picture of first taking everything in, and then deciding (in some way) to disregard or inhibit a portion of it. It is in that subsequent *decision*, it seems to me, that the frame problem is thought to reside.

I think Perceptual Control Theory comes at this issue very differently. It begins with some elemental property of the environment that can be sensed—some form of environmental energy, that can be transformed into some form of neural signal. These are the *Intensities* of the proposed PCT Hierarchy. It then starts constructing various new forms of invariants, with what it has to work with.

In other words, the *framing decision* happens on the perceptual side of the loop, as perceptual input functions (PIFs) are constructed, not on the (later) output side of the loop as actions are contemplated.

In fact, that very first step of transforming light or sound (or some other form of environmental energy) into a rate of neural firing, is itself a framing step. By necessity (which means, by virtue of the architectural way we are constructed), we only perceive some subset of what is theoretically possible to perceive. *How is this determined?* Presumably, it is some kind of evolutionary process, with different species settling into different solutions (read niches) of what they will perceive and capitalize on in their environment.

II. THE COMPUTATIONAL PROBLEM: *IT IS WORSE THAN YOU THINK*

What I am now calling the framing decision is taken up again and again, as each new type of perceptual invariant is constructed. These are (in developmental order) the so-called *Sensations, Configurations, Transitions, Events, Relationships, Categories, Sequences, Programs, Principles,* and *System Concepts* of Bill Powers's heuristic proposal, for how hierarchical perceptions may be arranged in human control systems.

This suggests that *the computational aspect of the problem* (and it is a sizable one!) is located somewhat differently for PCT than for AI modeling or for philosophical epistemology. It arises at the very outset, when perceptual input functions are constructed, not at some later stage, when data structures are updated or action outputs are contemplated or generated.

It must be admitted that PCT has no magical solution to this problem of constructing realistic (or robotically useful) PIFs. This has been one of the bottlenecks for PCT research. We know that the organizational layout of a control loop controls robustly; it does what it is designed to do, and it does it very well. We also know that the transform functions can be very complicated, and it can still work extremely well.

Examples would be as follows: Tom Bourbon has experimented with environmental feedback functions being routed through other controllers before the results are monitored by the original control systems, and the perceptions still get controlled. Bill Powers has experimented with hundreds of elemental control systems sharing the same environment with limited degrees of freedom, and still being able to arrive at stable control. Both Bill Powers and Rick Marken have demonstrated that the principle of hierarchical control by systems at several levels at once is not a problem for properly designed control systems.

Richard Kennaway has experimented with applying the hierarchical control architecture to problems of a walking robot. These are all very striking demonstrations of principle for the PCT approach.

The piece that is hard in PCT research is to construct perceptual input functions that make the perceptual world of the experiment look recognizable and therefore interesting to people reading the research. Perceptions such as force, acceleration, velocity, and position do not look very compelling. Yet they have been combined into sophisticated solutions of control with difficult engineering problems such as an inverted pendulum, and an arm reaching for a moving target (called the Little Man). It is out of such efforts that more elaborate robot simulations will be devised. But so far, the results do not look jazzy enough to be called *intelligent*. Nor do they include many of those layers of perception listed above, which are proposed to comprise the richness of human control.

III. TOOLS FOR LIMITING THE DAMAGE: *LET'S NOT MAKE IT HARDER THAN WE HAVE TO*

A. Action Is Not Primal

A slightly different version of the frame problem (from *knowing what to ignore*) is implied in some of the philosophical literature. There the problem seems to be one of how to infer the consequences of action, and what information to base that upon.

Again, I believe PCT takes a very different approach.

PCT does not start with inferences and then decide on the proper action. PCT builds into its very architecture a way to monitor the consequences of action. That is the concept of *perception*, as PCT uses the term. Consequences are not inferred, they are sensed or experienced directly.

But, ah—say the philosophers—there are any number of consequences of action. Conceivably, everything is related to everything else. How does one limit the scope of what is to be monitored?

Here again, I believe PCT approaches from the other side of the equation. PCT does not begin with action.

Action is only something that is needed to get one's perceptions coming in the right way. Perception is everything, at least to a PCT first-approximation as to what matters. Action comes later, and if current actions do not alter perceptions in the right direction, different actions are tried.

This issue of *altering perceptions in the right direction* brings in the PCT second-approximation as to what matters, namely reference standards. What matters is first of all perception, and second of all a perception tracking the state of its preferred reference. Once again, action is a derivative necessity, to try to make those things happen in the right way.

B. Defining What's Relevant

Yes, but—back to the philosophers—aren't there still all kinds of consequences of attempted actions, which are causing all kinds of differences in the environment?

Here is part of the robust beauty of PCT, as I see it.

PCT avoids the computational intractability of this problem—although, it has a significant computational problem of its own; see above—by restricting what gets monitored to whatever perceptual input functions the system has so far constructed. If there is no PIF for a certain kind of difference in the environment, it is as if that difference did not exist.

Let me give an example. The operation of electrical motors generates radio wave interference, (according to my limited understanding of such matters). The vast majority of the time, that is not a problem to me whatsoever, simply because my body has not constructed sensors for picking up radio waves. Such *interference* does not exist (for my body), because I have nothing for it to interfere with. My computational problem of what I monitor is greatly simplified, thereby, and I put my efforts into controlling what I am constructed to monitor.

In other words, PCT's solution to the frame problem is not a philosophical one, but a pragmatic one. We perceive what we perceive. If we can generate actions that substantially can affect that, then we (generally) are able to also control what we perceive. Perceptions that can be controlled are *what matter* to control systems, by definition. When that definitional project is pursued with a vengeance—as it seems to have been by life forms on earth for a few billion years now—then quite elaborate symbiotic and ecological networks can be built up, as microstabilities of perceptual control are mutually exploited by collaborating organisms.

A key point here is that it is not necessary to *know* where such symbiotic stabilities arise. PCT includes the concept of *disturbances*, which include all outside influences—beneficial or detrimental—that are affecting the controlled variable of a control system.

PCT implicitly recognizes McCarthy & Hayes' (1969) classic formulation of the frame problem, namely, *the impossibility of naming every conceivable thing that may go wrong*. And PCT rolls that issue into its equations.

PCT collectively labels such eventualities *disturbances to controlled variables*. And I believe PCT has an effective way around this potential framing difficulty by restricting where it monitors all those things that may go wrong. A control system simply and solely monitors the state of its controlled variable, as constructed by its perceptual input function. That is the only world it knows or needs to know. All influences, whether disturbing ones or helpful ones, are summed together and rolled into that one measure (for each elemental control system).

There is no need to pre-specify or know ahead of time which preconditions could keep an action from having its intended result. They are present in the equations of a control system in the operationalized notion of the disturbance, and they are checked directly, in real time. That is to say, the results of any disturbances—not the disturbances themselves—are monitored directly, by how the PIF is constructed.

Furthermore, the output function of a control system provides a way to attempt to counteract the effects of any disturbances. Disturbances in themselves do not need to be known. It is sufficient to monitor the net effect of any disturbances, combined with the counteracting effect of the output function.

This vastly simplifies the problem of framing, and it avoids the need for *frame axioms* about which aspects can be presumed to remain the same when something changes. All of that is monitored in real time, at one single point in each elemental control system, namely, the current state of the perceptual input function.

IV. SUMMARY: *SOME MAPS ARE BETTER THAN OTHERS*

One way to sum up a PCT approach to the frame problem is to adapt Gregory Bateson's formulation of *a difference that makes a difference*. The two uses of *difference* in this phrase can be mapped onto PCT's notions of *perception*s and *references*, respectively.

The only differences that a control system *knows*, according to PCT, are whatever ones are embodied in its perceptual input function. This means, that different hierarchical levels (if there are such) compute different invariants (i.e., what will not be perceived as a difference), and thus perceive different *differences*. The notion of a *frame* is absorbed into the PCT concept of the perceptual input function.

Environmental energy is framed according to whatever PIFs are constructed. If something does not show up in a PIF, in effect it is *attached to a frame* that has not changed because it has not been perceived.

The input side of a control system only deals with differences it can perceive.

There is also the notion, however, of *making a difference*. I believe this maps most cleanly onto the PCT notion of *reference standards*. If there is no preferred reference for a perception, the perception may have changed but who cares? It would not *make a difference* to the organism. Again, the issue is not a philosophical one, of whether there is a rational basis for ignoring or preferring a certain difference.

The issue is a pragmatic one, of whether a given reference might contribute to bringing about a perception that might *make a difference* to that organism, with that hierarchy of reference standards currently operative.

Essentially, I believe PCT provides a better map for how to navigate through and around *the frame problem*. For one thing, it situates the issue differently, as one of perception, not one of inference or action.

Framing is inherent with how all perception is constructed, but it does not thereby render organisms overloaded and immobile. If they are constructed as perceptual control systems, they will act to control their perceptions as they are, no matter how simple or complex.

There is indeed a computational bottleneck for even a PCT approach to the frame problem. But here, too, it is not where AI researchers or epistemological philosophers have located it. It occurs early, in the perceptual input side of the design of control systems, not late in the computational output and functioning of such systems.

From a research standpoint, there may be algorithms that could construct sophisticated perceptual input functions, which could then be controlled according to the basic design of a negative feedback control loop.

The bigger problem might well be whether such *perceptions* would be recognizable as such to us with our human perceptual input functions. We know, for example, that bees and homing pigeons and migrating birds control various perceptual aspects of their environments, but how such perceptions are constructed is still beyond our technology.

The key point is that each new type of perceptual input function, whether in other species, or within an internal human hierarchy, essentially creates a new perceptual world. If anything, that compounds the computational problem, whether for AI or for PCT.

I believe the good news, from a PCT standpoint, is that the basic equations of a control system organization appear to be very robust. That would suggest that it should be possible to test them out in various situations with all sorts of *arbitrary* perceptual inputs, including those in hierarchical networks, not just with those deemed perceptually *realistic*.

In the meantime, PCT provides some compelling ways to re-frame the framing dilemma. For one thing, disturbances (i.e., things that may go wrong) do not have to be dealt with independently. They can be monitored at one point in the system, namely, in the net impact they have on the controlled variable, as calculated by a system's perceptual input function.

Similarly, the effects of actions can be monitored the exact same way. Their consequences do not need to be inferred—even though there are ways even within PCT to model an inferential planning process.

Action can simply be checked out directly, in real time, to see if it is having an effect in the preferred direction on a controlled variable. If not, something else can be attempted, even if there is no way to pre-compute its anticipated effect.

Essentially, PCT redefines what is most relevant to consider, in determining how living systems function.

While it does not solve all the computational problems, it does provide decisive place-holders for determining what is relevant.

Relevance, for a control theory task, is defined in two ways. First of all, is there something to perceive, and can it be perceived? If it cannot, it is not (yet) relevant. Actions may be taken, or contemplated, but there will be no way of determining their effects, without some perceptual consequence to monitor.

The second way that relevance is defined is in terms of reference standards. Is there a preferred reference for a given perception? If not, then that perception does not currently matter to the organism. Only perceptions for which there are corresponding reference standards constitute *differences that make a difference* to the organism.

This is a constructivist and pragmatic approach to the issues raised by *the frame problem*. I believe it is a better approach than that of rationalist philosophy.

It adopts a *satisficing* rather than an *optimizing* approach to the problem. More importantly, it suggests that optimizing is not the only game in town. Some very credible and robust findings can be obtained with a PCT approach, and so I believe it definitely deserves very close consideration.

Note: The Frame Problem was brought up by Jeff Vancouver in a post to CSGnet on August 19, 2005, and the discussion concluded on August 29, following the above essay by Erling Jorgensen.

Along the way, Bruce Nevin suggested two general references to provide context for anyone not familiar with the issues, and Erling found a third:

http://cogweb.ucla.edu/ep/Glossary.html (scroll down to *Frame Problem*)
http://cogweb.ucla.edu/ep/Sociobiology.html
http://plato.stanford.edu/archives/spr2004/entries/frame-problem/