

Feedforward Revisited – Bill Powers December 2010

Something is coming together that is making sense of some ideas I have resisted for a long time. It has to do with the brain's models of the external world. From the way I have seen those models proposed by others such as Ashby and Modern Control Theory adherents, I have thought they were simply impractical, calling for far too much knowledge, computing power, and precision of action -- as indeed they are and they do, as they have been presented.

But those ideas may nevertheless be right. Some of those other blind men standing around the elephant are perhaps only a little nearsighted, and are seeing something going on that looks fuzzily like modeling, but there's something funny about it so it isn't quite how it seems from this angle or that. This particular blind or nearsighted man writing these sentences has not seen models; he has seen a hierarchy of perceptions that somehow represents an external world, and a large collection of Complex Environmental Variables (as Martin Taylor calls them) that is mirrored inside the brain in the form of perceptions.

Briefly, then: what I call the hierarchy of perceptions is the model. When you open your eyes and look around, what you see -- and feel, smell, hear, and taste -- is the model. In fact we never experience ANYTHING BUT the model. The model is composed of perceptions of all kinds from intensities on up.

Warren Mansell asked some questions about feedback and feedforward that stirred a few thoughts up. I think his ambition to integrate different ideas people have had about control theory suddenly looked more appealing than before. I've been working on and thinking about how to get a better fit of the current tracking model to the real behavior, and that has stirred up a lot more thoughts. I was thinking about how to add a two-level controller in which the upper level controls position and the level below controls rate of change (yes, I know that's backward). I realized that I would need a sensor that senses rate of change of position, and that, in turn, called to mind the neat analog-computing technique that computes first derivatives by putting an integrator in the feedback path of a little control system -- it's actually described in LCS3, chapter 5.

I considered using that method to implement a new model for the TrackAnalyze program and for some reason didn't like the idea of doing it that way. Then the reason dawned on me: I was actually proposing to put a model of the physical environment into my PCT model, and I'm not supposed to be in favor of doing that. But it happens that if you integrate the force applied to a mass, the value of the integral represents the velocity, which keeps changing in proportion to the force. The velocity is the first derivative of position. The factor applied to the force as it is being integrated represents the reciprocal of the mass of the object being pushed upon. So I was proposing to put a model of the mass of an arm, together with Newton's laws of motion, into my sacred PCT model.

So: I was thinking of sticking a model into my model, between the output and the input, as a convenient way of getting a signal that would represent velocity. It would be generated by applying a force to a simulated mass. So the arm controller would sense the force its muscles were producing and integrate the force to create a

synthesized perception of the velocity, and then it would have a controller for controlling that integrated perception and we would have one level of control.

But wait. Where did that model come from? Don't we need to control through the real world outside? It came from applying perceived forces to perceived things and -- for one example -- seeing them move. A kinesthetically detected output force becomes a perceptual signal representing force; the force signal is integrated to produce a visual perception of changing velocity; a visual perception of velocity is integrated to produce a visual perception of position, and a changing velocity produces a perception of acceleration. And this is all happening inside the nervous system. In a model.

The modern control theorists came closest to seeing how this works. They said that the internal model was carefully constructed to have the same properties as the external "plant" that was to be controlled. Then the brain could work out, internally, what signal it had to send into the model to make it behave in a certain way, and when it had that working, it could send the same signal to the external "plant" and it would behave the same way. They admit that to make this work the model of the plant has to be rather dauntingly accurate, and every disturbance has to be accurately anticipated as to size, direction, and time of occurrence.

So the picture I got was that the brain supposedly had the ability to examine the plant and measure its properties, and then constructed a computed model inside itself based on the data thus obtained. But of course I knew that the brain can do no such thing: all it knows are the perceptions it gets, and it has no way to compare them with the real plant Out There to see if it got the measurements right. Everything it does has to be done with the perceptions, not with the real plant.

That is where I had always stopped before, just prior to discarding the model-based control idea once again. But for some reason, this time I kept going.

We can sense output force because the tendons have sensors that report how hard the muscles are pulling, and we have pressure sensors all over that detect how hard a hand or foot is pressing against something else. We have sensors to tell us if a joint angle is changing as a result of the force, and of course we have vision to give a different spatial view of the result. So by experimenting with output forces, we can build up a set of control systems for controlling the immediate consequences of applying forces. We can get to know how much consequence a given amount of force produces. Years later we will learn that the ratio of force to consequence is called "mass." But if we integrate the force to produce a velocity, we can discover empirically what the value of this ratio is for different objects, without calling it anything.

That is all we need to do to build up a model of the external world. It's not even that; it's just a model of the world. The idea that there's also an external world that we don't experience takes a while to develop. At first it's just the only world there is.

So that is the model that Ashby and the Modern Control Theorists are talking about. It's the world we experience. When we examine that external plant in order to model it, we're already looking at the brain's model. It lacks detail, but as we probe and push and peer and twiddle and otherwise act on these rudimentary perceptions, new

perceptions form that begin to add features and properties -- like mass -- to the model. We say we are analyzing the plant. What we are doing is building up perceptions of properties and features that can be affected by sending signals outward, learning how to control the perceptions. Why we have to act one way instead of another to get a particular effect is unknown, but we learn the rules. When we don't get the effect we want, we alter what we are doing until we do get it.

We never do actually, knowingly, interact with the plant itself.

It seems very risky to be operating entirely on an internal model without any ability to know what is really going on that we can't see, but really, it's not. Before you step into the bathtub you feel the water, so if you've made a mistake you're not going to scald your whole body. We detect errors very quickly and make adjustments almost as quickly to limit the errors, and eventually to keep them from ever getting very large. We're always interacting with whatever is Out There, and we learn fast. Most of us, most of the time, don't even think about the invisible universe Out There. The visible one is sufficient to keep us busy and interested. The idea that there's another bigger one that actually determines what the rules are doesn't usually arise.

I'm beginning to get an idea now about how to model perceptions, at least at the lower levels. All we have to do is make a model of the environment, just like that analog-computing trick for calculating rates of change by using integrators, which turns out to embody Newton's laws of motion. This whole idea is still very new and I don't see very far along the path ahead, but I have a feeling that what looked very difficult before may start getting a little less difficult.

I'd better get to bed; it's very strange to look around at this room and think "This is my model. I, or something in me, constructed every detail in it, all the things I recognize and know about it and can do to it. Help, is this solipsism?"

But no, it's not. Solipsism says there really isn't anything else. We can freely assume that there is a huge lawful universe full of regularities, as long as we realize that all we will ever experience of it is the model that we build in our brains. When it does what we call raining we get what we call wet, but we can only assume that those experiences occurring in our models correspond in some unknowable way to whatever else there is.

I hope all of this doesn't evaporate overnight.

Best,

Bill P.