Powers on Todorov and Optimal Control Theory – November 2011

Reading Todorov, I'm coming to realize that my concept of "optimal control theory" is out of date. Todorov notes the open-loop nature of the older ideas and sees that "intelligent use of feedback" is much preferable. He also realizes that there are uncontrolled dimensions of behavior -- he calls them "redundant" -- that can be ignored. And he approaches the idea of hierarchical control, even going so far as to say that the model-based controllers may exist at higher levels, but they work by sending "commands" to neuromuscular negative feedback servos at lower levels. At one point he proposes that reaching behavior, which is disturbance-resistant, is actually accomplished by specifying a series of intermediate targets which are reached one after the other -- he comes ***that*** close to seeing that it's just a smoothly varying reference signal that goes to the lower systems, but doesn't quite get there. At least he rejects the idea of planned and controlled trajectories.

So in the papers I've seen, he is working his way slowly toward PCT, and I suppose would get there eventually.

There are two main ideas that he would have to overcome and discard first. The first is that sensory delays make fast control impossible using only negative feedback, and the other is that the signals in the model have to be horribly noisy. I'd love to see a graduate student do a real job of tracing where those ideas came from -- Lashley is certainly not the only source. I see Todorov has an "efference copy" in one diagram, which he says is needed because of sensory delays and noise, so von Holst has to bear part of the blame (though I don't think von Holst spoke of internal models and sensory delays).

All of (what I see as) the more advanced ideas in the Todorov version of optimal control theory appeared in the first paper by Clark, Macfarland, and me in 1960. I guess the literature in which we published then and afterward was simply not read by the people who produced optimal control theory, starting with the Kalman filter. Of course I didn't know then that sensory delays could be so easily compensated by an integrator in the loop, and I hadn't developed any testable working models yet (though I used an analog computer to model control systems). Those things happened during and after the years between 1960 and 1973 when B:CP came out.

The other main idea that has to be dropped concerns the noise in the system. I always thought that was a mistake, simply because when I look around at my world (I can't comment on what others see), it's not buried in sparkly snow and jumping and wobbling around all over the place. When I lift a fork to my mouth I hardly ever miss. Tom Bourbon pointed out the millions upon millions of drivers whizzing past each other in opposite directions within a couple of feet of a collision, and hardly ever hitting each other or anything else.

I think the idea that the nervous system is noisy came from looking at signals in single axons. In the first place, not understanding what the rapid variations in firing rate represent, it's easy for an experimenter to conclude that the variations are random even if they're completely systematic. And in the second place, I doubt that any information path in the brain is served by one lone axon. In the spinal cord I know for a fact (i.e., a publication says so) that the error signal running from the spinal cord to a typical muscle is probably a bundle of around a thousand fibers coming from a thousand motor neurons. Naturally, if you just look at one of them, there will be some random variations. But their combined effect is smoothly variable and almost noise-free. With "recruitment" having its effect, the result is even almost a linear representation of the stimulus at the source. I think we can ignore neural noise for the most part, except at very low levels of signal magnitude at which only a few neurones fire at a time. I've never seen much noise in my tracking data, in which the model's behavior accounts for 95% or more of the variance of the real handle movements.

Finally, most of the apparent noise in behavior comes from disturbances, which being uncorrelated with each other do give a pretty good imitation of randomness. Most of the action of a control system is needed to counteract small (and sometimes large) disturbances. Of course only the part of the frequency spectrum of the disturbance that lies within the bandwidth of good control is counteracted.

In the big picture, Todorov's approach is not wildly different from PCT. It just hasn't been developed nearly as far. A good part of it was wasted effort because of trying to deal with nonexistent problems, but all theorists including me have been through that. Par for the course.