Causation and Negative Feedback Control Loops

Erling Jorgensen July, 1999; December, 2005 Posts to CSGnet

Discussion on determinism

[From Erling Jorgensen]

> Bjorn Simonsen
> After re-studying texts about Teleology and parts of PCT, I have the understanding that PCT is quite independent of any causal relations.

I don't think I would express it the way you are above, but I agree that control processes necessitate a reconceptualization of cause-effect relations as traditionally expressed in lineal causality. The whole analysis changes when you close the loop with feedback processes, because it introduces a different relationship to time. I would almost say that this feature is constitutive of living processes. They do not just exist with external causes; because of the circular causality embedded in negative feedback arrangements, they in some sense cause themselves.

Back in 1999, I posted to CSGNet an essay that tried to partition the various notions of "cause" as applied to control loops. I think it is relevant to what you are raising, and I'd like to bring it into the discussion.
Here is a reposting of that essay:

Causation and Negative Feedback Control Loops.

From a systemic standpoint, I'm not sure "cause" is a very helpful word in talking about control systems. It all depends on which portion of the system you are considering at the moment.

A) Every snippet of the control loop can be thought of as propagating a signal, and in that sense it has a (causal) input and a (resulting) output. To the extent that we in CSG analyze in this black box fashion, we usually focus on the "nodes" of the control loop, i.e., the comparator function, the output function, the perceptual input function (PIF), and occasionally the environmental feedback function.

Aside: Regardless of how many actual neurons, evoked potentials, graded potentials, membrane permeability, etc. may actually be involved, and whether the signals are meeting in ganglia or other types of neural tissue, the Comparator has been elegantly modeled as having two net inputs with inverted signs—reference and perceptual signals respectively—and one net output, the error signal. In a sense, the interaction of reference and perception "cause" the error, but that's not the best way to think about it.

2nd Aside: Some have spoken of the error signal "causing" the output or action of a control system, but that again seems to cut the loop into snippets. The output function has been powerfully modeled (in the tracking demos) as an integrating function, requiring not only the reference-minus-perception input, but a multiplier constant representing gain, a multiplier constant representing a slowing factor (I think that's the "leaky" part of the integrator), and the previous output of the function as a new input! What's the "cause" in all of this, or is that not the right concept to impose on control systems?

3rd Aside: Apart from some weighted sums, and applying some logical operators, I have seen almost no modeling of different types of perceptual input functions. If Bill's suggestion about hierarchical levels of perception is a useful launching point (and I think it is), then theoretically there should be ten or eleven qualitatively distinct ways of modeling PIF's.

The actual neural computations of perceptions are undoubtedly incredibly more complex, but for a model all we would need to begin empirically testing of its concepts is to reproduce some essential feature of a postulated level of perception. For instance, the essence of a Transition, to my way of thinking, is the simultaneous experience of variance mapped against invariance. An Event is a series of transitions framed—one might almost say arbitrarily—with a beginning and an end. So an event control system (again, as I conceive it) is the one that does "framing", but to test whether such perceptions can be constructed and stabilized against disturbances, we first need measurable models of variants and invariants mapped against each other. Such complexities are beyond my current modeling abilities, (not to mention the point Bill has raised about getting the right dynamic equations to model the environmental forces on the computer.)

B) We can also consider a single control loop "in isolation," and ask whether causality is a helpful concept there. The rules change (and so should the concepts) when you close the loop. In one sense, every part of the loop is a cause of every other part.

The corollary to this is that every part of a closed loop is a cause of itself! Some theories like to respond with the idea of "circular causality," but I think Rick is right that it often just amounts to linear causality chasing itself incrementally around the circle. The fundamental idea of accumulating integrating functions (with all their ramifications) doesn't seem to enter the picture. It seems better to think of the organization of components itself, not some event occurring within it, as the effective cause.

C) We can move the zoom focus slightly farther out and consider a single control loop together with its inputs. As the basic model now stands, every loop has only two inputs from outside itself—one from inside the organism, the reference signal (which, again, can be modeled as the net effect of whatever neural and chemical processes actually bring it about), and one from outside the organism, the (net) disturbance. A traditional view of causality would say that the reference and the disturbance are the only two candidates for being a "cause," and in a sense we in CSG accept that.

But by quantifying the relations in the loop into equations, Bill et al. have been able to say something much more precise about these external causes. Only the reference is an effective cause of the stabilized state of the perceptual input quantity.

Any causal effect from the disturbance on that quantity is neutralized by the negative feedback action of the loop. The cost is that the disturbance becomes an effective cause (in inverted form) of the behavioral output. [This latter point seems to be what Herbert Simon was referring to in his quote about the behavior of ants, that was hotly debated awhile back on CSGNet.]

D) We can move even further back and look at more of the hierarchy, as it's currently proposed to operate. Here almost every control loop is embedded in a network of control loops "above" it and "below" it. So in one sense, higher loops "cause" it to operate by providing changing reference signals, and it "causes" lower level loops to control by the same mechanism. I deliberately say "higher loops", plural, and I mean it in two senses. For one thing, many loops at the next higher level can be contributing to the net reference signal at a loop at the next lower level, so perhaps all those loops are causal. But we can also speak of proximal and distal causes, and include each relevant loop all the way up the hierarchy as a "cause" of a given low-level loops' operation. This is why I have no problem considering "attending a meeting" as one (distal) cause of contracting a given muscle on the way to the garage.

Just as closing the loop changes the notion of causality, so does embedding everything in a network.

E) Sticking with this hierarchical vantage point for one more iteration (if you've stuck with me this far!), it needs to be emphasized that the interaction between levels does not occur by intact loops sending signals to other intact loops below them. Rather, those lower level loops are *part of the structure*, part of the loop itself, of the higher level.

Remember, all loops are closed through the environment—(other than the "imagination switch," if we can figure out a way to get it to function!)—which means that higher loops have the longest (and slowest) path to travel to achieve their control. And they only achieve it if the lower level loops to which they contribute are achieving sufficient control of their own variables.

So maybe this reflection has come full circle (sorry about the pun, but it fits!), in that when higher levels "cause" lower level perceptions to become stabilized, they are simply causing their own control to happen. Basically, I think we have two choices for using causality in a way that reminds us (instead of deflecting us) about how control loops operate.

1) Either we allow this reflexive notion of "self-causality" to be part and parcel of how we use the term—which means processes in the loop are always in a time relationship within themselves, as well as always functioning and embedded in higher and lower loops.

Or ***2)*** we say causality cannot be determined apart from the organizational structure that one is considering. In essence, it is not a relationship among events that pass through the loop, but rather a property of the organization itself. The answer to "what's causing this action?" is the same as to "what's causing this perception?" It is the fact that these components are organized into the functional form of a control system. So to speak about causes, you can't stick with the events. You have to address the question, how does the system (specifically as a system) bring about its own functioning.

All the best, *Erling*

Bruce Nevin July 1999 commenting on
Erling Jorgensen's post earlier in the day

An excellent synopsis! A tour de force, tracing the theme of causation coherently through the Gordian all-at-onceness of the control hierarchy, and showing how a model can help us distinguish the multiple ambiguities of a simple word like "cause."

Another step: the relationship between a neural cell (an autonomous control system) and a multicelled control system in which it participates. This is a special case of the relationship between the cellular order and multicellular orders of organisms. Perhaps there are parallels between virus-host and parasite-host, prokaryote-cell and symbiote-organism, primitive multicellular structures and observable social structures. But your theme is causation. Here, the causal cord is cut, or anyway has a more accidental cast, as a side effect that is not only unintended but even beyond the perceptual capacities of the organism effecting it. It seems clear that cells are indeed autonomous control systems, that they do not control any variables that are controlled by the organisms that they constitute (cells do not control neural signals as such), and vice versa (humans do not control rate of flow of ions across a cell membrane, or whatever it is that the cell is controlling). And this however much the control of variables at each level or "order" of organization (e.g. cellular and human) may *influence* the state of variables controlled at the other. We may speak of a cancerous tumor in the brain as causing a loss of eyesight, but that is not a result that is controlled by the cells or by the tumor; it seems rather that cancer is a side effect of the failure of some cells to control their "social" relations with other cells in an organism.

Organisms that collectively stabilize their shared environment (which includes especially each other) survive better than those that don't.

Consequently, on an evolutionary time scale, in each type of organism that survives there must arise innately some controlled variables, or more likely some way of controlling variables (input and output functions), that has as a side effect a tendency of the individuals to stabilize in higher-order systems. But this is not a discussion of perceptual input control; it is a discussion of the evolution of innate properties (variables, values, input functions, output functions) in populations of control systems. There is causation here too, but even more indirect than what you have discussed.

Bruce Nevin