Closed Loop

Threads from CSGnet

Fall 1993

Volume 3

Number 4

Front cover

Closed Loop

Threads from CSGNet

Fall 1993 Volume 3 Number 4

Edited by Greg Williams, 460 Black Lick Rd., Gravel Switch, KY 40328

CONTENTS

Members of the Control Systems Group receive *Closed Loop* quarterly. For more information, contact Ed Ford, 10209 F. 56th St., Scottsdale, AZ 85253; phone (602) 991-1860.

CSGnet, the electronic mail network for individuals interested in control theory as applied to living systems, is a lively forum for sharing ideas, asking questions, and learning more about the theory, its implications, and its problems. The "threads" in each *Closed Loop,* stitched together from some of the net's many conversations, exemplify the rich interchanges among netters. Some issues of *Closed Loop* also feature research reports by netters, in hopes of initiating new conversations.

There are no sign-up or connect-time charges for participation on CSGnet. The Internet address is CSGL@VMD.CSO.UIUC.EDU while CSGL@UIUCVMD is the Bitnet address. Messages sent to CSGnet via these addresses are automatically forwarded to over 120 participants on five continents, as well as to hundreds of NetNews (Usenet) sites where CSGnet can be found as the newsgroup bit.fistserv.csg-l. CSGnet also can be accessed via CompuServe, AT&T Mail, MCI Mail, or any other computer communication service with a gateway to Internet or Bitnet. For more information about subscribing to CSGnet, contact Gary Cziko, the network manager, at G-CZIKO@UIUC.EDU, phone him at (217) 333-8527, or send a FAX to (217) 244-7620.

Two Views of Control-System Models

Hans Blom: What is the fastest way to get a spaceship to Mars? The solution is well-known, although impractical: apply full thrust until you are at the exact midpoint of the trip, turn your ship around and apply full thrust again, braking until you arrive at Mars with zero speed. This is an example of what is called "bang-bang" or "minimum time" control, a control paradigm quite different from the "stabilizing control" that is usually discussed on CSGnet. Features of bang-bang control are these: (1) outputs are either zero or at their maximal limits, (2) the only important parameters are the times at which outputs go from zero to maxima, or from maxima to zero, (3) in general, it is quite difficult to find optimal values for those times, and (4) for long periods of time (between the decision points), it might seem to an outside observer that control is absent, because nothing changes—because there is no modulation of the outputs.

This is discussed in "The Neural Control of Limb Movement," by William S. Levine and Gerald E. Loeb, in the December 1992 issue of *IEEE Control Systems*. Does the organism use bang-bang control? No. "The experimental data... show a substantial deviation from the optimal control model." Why is that? Partly in order to protect the organism: "the feedback from the joint sensors, while certainly present, would be too late to prevent injury if a human jumper tried to perform a mathematically optimal [i.e., top-performance] jump." And partly because "it is important for both biologists and control engineers to remember that the control systems that have been invented to date are almost certainly a meager subset of all possible types of control and even of all control methods used in biological systems. Thus, the study of biological systems should not be confined to testing whether their performance is compatible with control schemes invented to date but must include detailed examination of their inner workings to discover new types of control."

Some type of stabilizing control is needed in all cases where full-time control relative to a setpoint cannot be relinquished even for a moment. But stabilizing control is incompatible with top performance, such as in sports. In high jumping, only the maximum height of the jump is important, not the full trajectory. In the Mars rocket, the output resources are used at 100% capacity during 100% of the time; the only control decision is to find the exact point in space-time of the turnaround. Mathematically, due to the nonlinearity of the problem, finding this point is generally intractable and therefore usually a matter of trial and error (search) or creative insight. In humans, finding the optimal decision points requires a considerable period of tuning and fine-tuning (training).

The authors pose more questions than they provide solutions: “much more work needs to be done before the above suggestions can be called a theory.” Yet, in my opinion, this paper provides some insights into why stabilizing control, which works so well in ordinary circumstances, breaks down when maximum performance is required.

Bill Powers: Hans, human control systems are pretty close to the design limits set by the materials used. It’s possible, for example, for an arm muscle to pull itself loose from its attachments to the bones, if feedback is lost and an energetic movement is attempted. Even with an intact set of control systems, tendons and muscles can be ripped loose if an emergency situation results in sending abnormally large reference signals to the spinal motor neurons.

The “substantial deviation from the optimal control model” that Levine and Loeb mention might not be a deviation from what is optimal for the whole human being using the control system. Control models of an arm usually propose the application of torques at each joint, but in the human system there are no motors at the joints. Instead, there are nonlinear muscles attached in clever ways that produce many kinds of torques, some through clever linkages (as in the two bones of the forearm), and some by having the muscle wrap around the joint in a strange way (like the biceps).

Even the muscles work differently from the servo motors that engineers use. They don’t apply forces directly, but by shortening the contractile elements in the muscle to alter the resting length of the series spring component. In principle, a movement could be carried out by suddenly shortening all of the contractile elements in a muscle and storing energy in the spring components, then letting the spring components execute most of the movement without any further expenditure of muscle energy until time for deceleration. Actual movements work somewhat in this way. This is something like the solution for maximum rocket efficiency given a finite fuel supply. In fact, the human system is far more efficient than any robot so far invented; it moves 100 to 200 pounds of weight around all day expending only two or three kilocalories of energy and using less than 0.1 horsepower of total muscle output power. And the fuel supply has to support not just the muscles, but the brain and the general metabolic requirements.

The reason a human being can’t perform a mathematically optimal jump is simply the rocket problem: you would need to produce an impulse of muscle force of zero duration and infinite amplitude. That would hardly be a feasible solution for a servomechanism, either.

The “feedback too slow” argument turns up even here, doesn’t it? Actually, the speed of feedback in a human control system is just right—to explain the behavior we see.

And you also say: “But stabilizing control is incompatible with top performance, such as in sports. In high jumping, only the maximum height of the jump is important, not the full trajectory.” Human beings hardly ever control the “full trajectory.” They control the variables that matter to them. Rodney Brooks has the right idea here: don’t plan trajectories, avoid obstacles. It isn’t necessary to know where obstacles will be, if the system has sensors that can detect proximity to an obstacle.

“Stabilizing control” is something of a misnomer, suggesting that all that a control system does is to keep something constant. More generally, it makes the perceptual signal track the reference signal. This means that a control system for producing a directed force (as in throwing a ball or launching a high jump) can make the sensed acceleration have the right magnitude and direction right up to the moment of release. When we learn how these perceptions must change in order to have a desired result remotely or later, we vary the reference signals to repeat the experienced thrust as nearly as possible, and we get pretty dose. Of course, if we got too dose, people would stop doing such things—or they’d set the bar higher, or put the target farther away, until errors in control once again made the game interesting.

I think that when normal human movements such as walking are finally modeled fully, we will find that the system uses as little energy as possible, letting momentum and spring effects carry most of the movement through, with muscle contraction being used primarily to trim the result into a useful form. When we walk, we choose a pattern of walking to control that is as dose to the zero-energy pattern as possible, given the higher-level goals of actually getting somewhere in a reasonable time. Only when we have some reason to get there faster, as in running a rare, do the control systems try to produce patterns that cost a lot of energy. And even then, the patterns finally chosen are pretty efficient—after all, the fuel supply and distribution have to suffice to get the body to the finish line.

Hans Blom: Bill, you say: “Human beings hardly ever control the ‘full trajectory.”’ If that is the case, “new types of control,” which do not try to maintain minimum error between reference values and perceptions at all times, might provide superior performance in some cases. Or greater ease. When I fly to New York, I (attempt to) control my destination, but in the plane I have to trust the pilot. Part of my trajectory will be, as far as I am concerned, ballistic.

What makes control in organisms so difficult to study is the simultaneity of a great many different ongoing goals, whose importance might, moreover, fluctuate from moment to moment due to influences beyond our control and usually beyond our knowledge. Only in the simplest of experiments one variable can be considered to be controlled, if at all. "Keep your finger pointed at the knot." But the subject also has to control the upright position of his or her body and otherwise keep all sensory channels open, if only to hear you say, "You can stop now."

Still, a high jumper wants to jump as high as possible, period. An objective measure is provided to test that performance. All else is unimportant (within limits, see below). What more can you ask for? There is no prescribed trajectory to be followed; a new world record often is an unprecedented experience for the jumper.

Human control systems normally function well within design limits. We have very little experience with operation near those limits: pain effectively causes us to stay away from them. But pain is carried by slow nerve fibers; in emergencies, the experience of pain can arrive too late to prevent harm. Is a case where "tendons and muscles can be ripped loose" really an indication of "an intact set of control systems"? I consider that to be pathology, a control system gone haywire, operating beyond its design limits. I would maintain that one of the most important of an organism's objectives is, at all times, not to seriously damage itself. But that cannot be formalized by control in the usual sense of the word, that is, a perception following a reference signal. The control system is operating under constraints, i.e., it must stay away front certain experiences with a high probability of success. Short-term goals are rarely important enough to jeopardize long-term goals, which need an intact organism.

You say: "The reason a human being can't perform a mathematically optimal jump is simply the rocket problem: you would need to produce an impulse of muscle force of zero duration and infinite amplitude. That would hardly be a feasible solution for a servomechanism, either." Impulses are not required, step functions will do nicely. After all, a trainer just wants to study the peak performance that a real individual is capable of, given his / her motor equipment, and search for whatever means there are to teach him / her to fire his / her nerves in such a way that this peak performance is reached.

Also, Levine and Loeb do not say that feedback is too slow; bang-bang control requires very accurate timing. They say that when the need for performance becomes extreme, protection mechanisms are required to prevent muscles and tendons from being torn loose. Feedback from those protective sensors would probably be too slow if training did not slowly familiarize the high jumper with the sensations that they provide. (Much of psychotherapy seems to serve the same function:

trying to get the client "in contact" with his / her feelings without being overwhelmed by them.) This is much like walking as dose to the abyss as you dare without risking the damage that a fall would cause. The fall would provide you with feedback, of course, but you wouldn't want that feedback, would you? (In psychotherapy, one of the frequent goals is to show the client that much of his / her "fear of falling" is imaginary, and that the abyss is much farther away than he / she thinks. This, too, is a difficult and often fearful type of exercise.)

You say: "Human beings hardly ever control the 'full trajectory.' They control the variables that matter to them." Yes. And bodily (and mental) integrity matters a great deal.

You also say: "'Stabilizing control' is something of a misnomer, suggesting that all that a control system does is to keep something constant. More generally, it makes the perceptual signal track the reference signal." Exactly how would you know that the jumper follows a reference signal when for the very first time he / she jumps higher than he / she ever did before? How does the reference signal get established in the first place? I do not allow the answer that it is an "imagined" reference signal; that would be impossible to either prove or refute, and it would therefore be unscientific (following Popper). I do allow the answer that the reference signal is discovered "by accident," through trial-and-error learning. But that would mean that the very first time there was no reference that could be followed, i.e., that not all behavior (here: peak performance) is control of perception.

Perception is not the only human capability that we depend on to control our behavior. Sometimes memory will do: a child will stay away from a hot stove after having been burned by it only once. Sometimes "knowledge," such as from a newspaper, will do: stay away from Chernobyl for a while. In neither case do you control for an exact distance from the feared location, you just want to keep at least a minimum distance away from it.

Maybe we have a different conception of what perception is. For me, perception is everything that my senses register and what can be derived from that. You might include memory as some type of "observation" through "inner senses." Is that what you mean?

That leaves the discrepancy of wanting something and not wanting something. More philosophically, I think that this distinction explains what gives us freedom. *There is* not one optimal location that is dictated by a match between our inner drives (reference levels) and our perceptions of the outside world. I do not dispute that we have reference levels and that we use our perceptions to get us dose to them. I just want to add something like "negative reference levels," things to stay away from. Freedom is a name for ranges in n-dimensional objective space where you can move about "at will," because the objective function is flat. It is as if you try to find the highest peak in a mountain range, and once you get there, you discover a wide, high-altitude table-land.

An example: you get conflict when the heater is set to 22 degrees Celsius and the air conditioner to 20 degrees. You get a region of "freedom" if the heater is set to 20 degrees and the air conditioner to 22 degrees.

As a control systems designer, I do not create control systems in the hope that they function correctly; hope has no place in the model. I do not rely on things going right only usually. I specify an objective function that I know will lead to a correct design. And if I cannot guarantee correctness, I will at least strive for optimality in some sense, such as longest mean time between failures or longest time before first failure. Would evolution be sloppier, given its billions of years of experimentation?

I assume that evolution, through a harsh billion-year-long struggle for survival, could have come up with some pretty clever solutions to the control problems that have arisen. *E. coli* has a funny (partly random) but clever control law that results in what is called a biased random walk This "primitive" control law serves it quite well; E. coli is far more numerous than *Homo sapiens*. Higher organisms have other (better?) control laws, some of which we seem to have more or less uncovered (control of voluntary muscles in humans) and which resemble linear quadratic control, at least as long as muscles function well within their force limits. Linear quadratic control works well in stabilization, i.e., stand-still and slow movements. In other cases, we know that there are better control laws. An example of that is when peak performance is required and the forces that muscles can deliver come to their limits. In that case, the nonlinearities of the actuators cannot be neglected any more, and linear quadratic control becomes sub-optimal. Intuitively, I agree with Bill Powers when he supposes that there is only one control law that governs the control of muscles. Linear quadratic control is, in my opinion, its more readily understandable "special case," just like Newtonian physics is a more readily understandable special case of general relativity.

People are very good (but often highly nonlinear) controllers. Moreover, it is my perception that people have a whole range of control schemes and frequently even apply the appropriate one at the appropriate time. This is a continual source of amazement (and envy) for control engineers who generally do much worse.

I know by now what perceptual control theorists mean by the mantra "organisms control perception." As so often with jargon, it is an abbreviation for a whole philosophy and only understandable for those who have gotten to know that philosophy. It is right, from a certain perspective. From another perspective, organisms control their outputs. I find it hard, in a control loop, to see one apart from the other. But, of course, sometimes you concentrate on the one, sometimes on the other. Very often, the output is controlled as well, for example in cases where different actions are possible (steak or salmon?), all leading to similar perceptions (great food!). Then you actively have to choose between outputs ("I would like...").

Perception is controlled by actions; actions are controlled by perception Remember the loop!

I agree with Bill Powers' "it's all perception" in the sense that perceptions (of the outside world and of our inner physical and mental mechanisms) are the only sources of information available to us. But perceptions are built upon and result in higher-level things that I would not call perceptions any more. Beliefs, superstitions, the "facts" of our lives. All those together constitute what I call a model (of the world, ourselves included). A model is, technically, always a simplification, and always has a purpose. That it is a simplification is due to the facts that we have experienced only a limited set of perceptions, and that our processing of those perceptions must be done by a mere three pounds of flesh. Models are never unique; it is always possible to translate one model into another, equivalent one. Sometimes a simple, approximate model works well enough, sometimes only a very complex and very accurate one will do, depending upon the goal that it serves. The highest purpose of the biological model is, in my opinion, best described by Dawkins: transmission of genes. Everything serves that supreme goal. The evolutionary process has weeded out every organism that did not serve its purpose well enough. A high degree of optimization has taken place during billions of years, and in that sense all currently existing organisms can surely be called well-designed control systems. Control systems, because they need to achieve a goal. There are numerous ways to achieve that goal. Viruses, bacteria, cats, and humans do it differently, thus far equally successfully. All other goals are sub-goals, designed through evolution to serve the one supreme goal. The sub-goals of each organism are uniquely related to its potential for actions, i.e., its body. A virus needs very few perceptions to achieve its goal; it mainly relies on the forces of nature ("free" energy) to work for it. A human, on the other hand, cannot survive without a great many perceptions.

In short, I think that the PCT perspective is extremely valuable when you study human behavior. A different perspective might be better for me, because I study very simple things like control systems. Let's by all means keep exchanging perspectives! Sometimes it seems less limiting to have two different perspectives on the same reality at the same time. Could that be why binocular vision proved to be successful?

Control engineers have a broader conception of control than you seem to have. Control does not necessarily imply feedback. In fact, engineers prefer non-feedback systems if at all possible, because they cannot possibly have stability problems. Regrettably, non-feedback control is possible only if the system to be controlled in invariable and not significantly subject to disturbances.

I think that by now I understand what perceptual control theory is about. I have followed and enjoyed the discussions for more than a year now, mostly quietly. Once in a while I grab the chance to vent some of my ideas, which are more or less related, hoping for a useful reply—usually not in vain. Reconciliation is not what I look for; I find that friction—clashing points of view—generates much more creative energy. Engineers and psychologists are not dose neighbors. They speak different languages, have a different culture, and work on different problems, although it is fascinating to discover similarities. I believe that engineers can learn as much from psychologists as the other way around. Doesn't this net show it?

Bill Powers: Hans, you say that "'new types of control', which do not try to maintain minimum error between reference values and perceptions at all times, might provide superior performance in some cases. Or greater ease. When I fly to New York, I (attempt to) control my destination, but in the plane I have to trust the pilot. Part of my trajectory will be, as far as I am concerned, ballistic." I think you're going about this backward. When we study human behavior, we aren't comparing it with some "optimal" or "best" way of controlling. We're just trying to understand what people are actually controlling under various circumstances. In some regards, people control things very well indeed, by clever means that surpass what any engineer knows how to build. In other ways, people control stupidly and poorly, and suffer the consequences.

More to the point, people use the means available to achieve whatever degree of control is possible. When I buy a ticket on an airplane, show up for the flight, and strap myself in, I have done all that is possible to get myself to the destination by that means of transport. So that's all of the control I have; if the plane is hijacked to another destination, that disturbance is beyond my ability to resist. All I can do is wait until the plane lands and I can get off it, and then start controlling again for getting to the destination by some other means. It could easily be that I would have arrived at the destination sooner, even without the hijacking, by taking a bus. But I didn't think of that. People are not optimal controllers; they just do the best they can.

You say: "What makes control in organisms so difficult to study is the simultaneity of a great many different ongoing goals, whose importance might, moreover, fluctuate from moment to moment due to influences beyond our control and usually beyond our knowledge." The hierarchical model helps here, because higher-level goals change more slowly than lower-level goals. Many of the fluctuations in conditions are just disturbances, which lower-level systems automatically compensate for by adjusting lower-level goals. Much of the apparently chaotic nature of behavior becomes more understandable when we ask about higher-level goals. We can then understand many external events as disturbances, and we can see how the changes in detailed behavior oppose their effects. This reveals regularity where formerly we couldn't see any. I think that most behavior is actually quite regular, once we understand what's being controlled at many levels.

You're right about the fact that more variables are under control than we can measure in any one experiment. But it's interesting that without much trouble we can get those other variables to remain constant enough to get good repeatable data.

You say: "Still, a high jumper wants to jump as high as possible, period. An objective measure is provided to test that performance. All else is unimportant..." The highest-level goal is to win the contest, not to jump as high as possible. There is strategy involved, as well as just trying to produce maximum effort. Some jumpers will pass at a certain height, saving their strength for later. they don't try to jump at all. Also, if you assume that every time you see a high jumper, the objective is to jump as high as possible, you will usually be wrong; most of the time, the high jumper is just trying to go high enough to clear the bar. On many other occasions, the jumper might not be concerned at all with controlling for height. The jumper might be working on the approach or the takeoff, or the form at the peak of the trajectory, or the flip that raises the legs at the critical instant, and not be worrying at all about maximum height. You can't tell what a person is doing just by looking at what the person is doing. The Test for the Controlled Variable helps you to understand what is actually being controlled (as opposed to what you logically assume is being controlled).

You say: "Is a case where 'tendons and muscles can be ripped looser really an indication of 'an intact set of control systems'? I consider that to be pathology, a control system gone haywire, operating beyond its design limits." Certainly it is. If pathology is involved, it is a higher-level system that is misusing its lower-level control systems. Is it pathological for a father to lift a Volkswagen off his child, suffering born muscles and ligaments (and a lot of pain) as a result? When a person shoots himself in the head, all of the control systems for grasping the gun, aiming it, and pulling the trigger are working perfectly well until the last moment; all that's haywire is the higher-level system that has chosen this outcome. And even that choice might not be pathological, if the person is facing torture or the pain and humiliation of a vicious disease by staying alive.

And: "I would maintain that one of the most important of an organism's objectives is, at all times, not to seriously damage itself." Normally, perhaps. Not always.

You say that "bodily (and mental) integrity matters a great deal." I disagree. This is like saying that organisms control for "survival." Organisms control specific variables relative to specific adjustable reference levels. An outcome of doing so might be that the organism "survives" or preserves "physical and mental integrity," but that is not a concern of the organism. It's an opinion of a third-party observer. I don't think that there is any reference signal specifying survival or integrity. Organisms don't survive or preserve their integrity, anyway. They all die.

You ask: "Exactly how would you know that the jumper follows a reference signal when for the very first time he / she jumps higher than he / she ever did before? How does the reference signal get established in the first place?" The trajectory is a side-effect of controlling variables that the jumper can control. It is not itself a controlled variable. Once the jumper has left the ground, there is no action that can alter the trajectory of the center of gravity. There are, of course, many variables that can be controlled during the trajectory, such as the relative configuration of the parts of the body. These can make quite a difference in whether the bar falls or not, but they have no effect over the path followed by the center of gravity. One of the tricks of high jumping is to control the body's configuration so that the center of gravity passes under the bar while the body itself passes over it. That process is under continuous control all during the trajectory.

I think that competitors control what they can control: the approach, the takeoff, and the body configurations. The outcome depends on how well they are able to control those variables.

The peak height of the trajectory, perceived over dozens or hundreds of occasions, might be a controlled variable if there are things the jumper can do to affect this average peak height. The associated control system would be very slow, and would operate by adjusting many lower-order reference signals for such things as practice time, amount of effort, adjustments of form, and so forth. During any one jump, of course, this averaged perception can't be controlled. But over time, the jumper can gradually raise the reference signal for height jumped, as long as this is consistent with maintaining the necessary elements of the jump in the right forms. On the initial jump of a competition, no jumper strives for maximum height. The reference height is set comfortably above the bar, but no higher than necessary.

I think you would have a dearer picture of the PCT approach if you kept the hierarchy in mind. The first time anything is accomplished, there can be no reference signal derived from experience of accomplishing it. At worst, one can have reference signals only for the lower-order components of perceived behavior that are to be put together in a new way. There are many possible ways for that to happen, including instruction followed by imagining the meaning of the instructions. At best, you've studied movies of someone else doing it and have some concept of the coordinations required.

On the first attempt, one seldom achieves perfect control. But the first attempt provides a perception of doing the control action, and from that experience, more realistic reference signals can be selected. Also, the new control system's parameters are probably not set to the best possible values; reorganizing them takes many trials, too.

To speak of "the" reference signal being "discovered" doesn't sound right to me. A reference signal is variable; it can be set to high or low levels. In any complex behavior, reference signals must be varied during the behavior if high-level perceptions are to be controlled at their given reference levels. Even when a behavior is well-practiced, the reference signals can be set to different states within the possible range. As I said, a jumper doesn't set a reference signal for the maximum possible jump early in the competition; you don't see champion pole vaulters clearing a 15-foot bar by five feet. I don't think that "maxima" have anything to do with it, anyway. The jumper simply sets a target height that is enough above the bar to clear it. When the bar is set too high, the target is still set above the bar, but now the jumper can't produce lower-level control actions sufficient to dear the bar, and fails.

If a jumper really set a reference signal for "maximum height'' (say, one kilometer), there would be an enormous error signal, and the output function would saturate, destroying control. To achieve maximum performance, one should set the reference signal just slightly above the level that the maximum possible efforts can achieve.

You say: "Maybe we have a different conception of what perception is. For me, perception is everything that my senses register and what can be derived from that. You might include memory as some type of 'observation' through 'inner senses'. Is that what you mean?" That all sounds OK to me. Perception is what we know of the world and ourselves. It exists physically as signals in a brain.

And: "I do not dispute that we have reference levels and that we use our perceptions to get us close to them. I just want to add something like 'negative reference levels,' things to stay away from." There are many reference settings that result in staying away from something. The simplest kind is a reference setting of zero. If you set your reference level for the perception of a loose tiger to zero, then any perception of a loose tiger constitutes an error, and you will act to reduce the perception of the tiger to zero by moving it away or yourself away from it.

And: “Freedom is a name for ranges in n-dimensional objective spare where you can move about ‘at will’, because the objective function is flat.” You get the same result from an inverse-square function If you keep the perception of the tiger at zero, you still have all of the other degrees of freedom of movement, the only restriction (which you set yourself) being that the perception of the tiger should not depart significantly from zero. Actually, by the way, you would probably not set the reference signal to zero, but to some small nonzero amount. If there’s a tiger on the loose, you want to see a very small image of a tiger, but you definitely want to see some image of the tiger. It would not be wise to lose track of where it is.

You say: “As a control systems designer, I do not create control systems in the hope that they function correctly; hope has no place in the model.” Well, you hope that somebody doesn’t pull the power plug, or that the motor doesn’t burn out a bearing, or that the environment doesn’t become so nonlinear that your design becomes unstable, and so on. Every system, however carefully designed, has failure modes, doesn’t it?

In fact, designed control systems live in an environment that’s almost totally predictable, so you can be pretty sure that nothing disastrously unexpected will happen. But human beings roam free through an undisciplined environment that is far more complex than any of them can understand. That environment is also full of disturbances that can’t be predicted (weather, for example) or even be sensed before they occur. Most of our “predictions” are statistical in nature; sometimes they work, and sometimes they don’t So there’s no way that living systems could evolve to anticipate every circumstance or act correctly every time.

There’s another factor that the designer has considerable control over. the forms of the analytical functions involved in the design. Most control systems are deliberately designed with linear components for the simple reason that we can’t solve the equations with nonlinear functions—not because nature doesn’t present us with nonlinear situations. In most real control problems, if you actually use the mathematical forms that fit the behavior of the environment most accurately, you find that you can’t solve the equations and can’t complete the design without trial and error. So we all use approximations; we fit a quadratic to the curve, instead of using a power of 2.113, which would fit better.

The human control systems have to work with the components that are given. They can’t approximate.

My job is actually easier than yours. I’m not trying to optimize anything—just to match the behavior of a model with that of a real human subject. I’m just trying to produce a model that controls as well as people do, not to produce engineering miracles.

Of course, real control engineers know a lot more than I do about the design of complex control systems, and some day they will take PCT much further than I possibly could. My job is not to compete with them or tell them their business. It's to get them to look at control in novel ways, ways that are not part of the customary approach—and not to improve the control systems they design, but to help us understand the behavior of organisms, most of which are not control engineers, either.

Pure reason isn't going to identify the actual variable under control by a given person in a given circumstance. A guess about what someone is controlling for could be quite right, or quite wrong. The only real way to find out is to apply a disturbance to the proposed controlled variable and see whether it's resisted in the way a control system would resist it. An even better way is to match a model to the behavior and find the parameters that give the best fit, and that predict future behavior in detail. This is why we refer to the Test for the Controlled Variable—because it provides a formal way of determining what is in fact being controlled, as opposed to what seems reasonable. People are not always reasonable. They don't all control for the same things in the same way. Sometimes they seem positively determined to do things the hard way. All we can do as theoreticians and experimenters is to find out what's really going on in a given person

You say: "I know by now what perceptual control theorists mean by the mantra "organisms control perception." As so often with jargon, it is an abbreviation for a whole philosophy and only understandable for those who have gotten to know that philosophy. It is right, from a certain perspective. From another perspective, organisms control their outputs." This isn't really jargon or "in" talk, but it is a problem with word usage. When I think of the "output" of a system, I mean the physical effect on the environment that is due to the actions of the behaving system *alone.* In the human system, this means muscle tensions, because that's that last place in the chain of outgoing effects where environmental disturbances can't get into the process and alter the consequences. Measuring the consequences any farther from the nervous system can give a false impression of what the nervous system is actually doing.

In a servo system, with this understanding of "output," I would not call the output of a motor the shaft position or speed, but the torque applied to the armature of the motor (at low speeds, anyway). Only that torque can be varied by the active system without regard to what the environment is doing. Only the torque output gives an accurate indication of the electrical output of the control system. The shaft position or speed will depend on the torque *and* on external loads and disturbances, so can't be used to indicate the output activities of the control system by itself (especially if the loads and disturbances aren't predictable).

So this is more a matter of labeling than ideology. I'm sure you would agree that a servomechanism doesn't control the torque applied to the armature of its motor, but only some consequence of that torque measured farther downstream in the causal chain. As disturbances come and go, the servo system varies its output torque, but it doesn't try to maintain any particular torque (unless torque itself is being sensed and controlled, which isn't the most common case). The torque has to be free to vary if disturbances of position or speed are to be counteracted.

The "control of perception" part is also a matter of labeling. I think you'll agree that in order to control an effect of a system's actuator output (to distinguish it from "outputs" farther along the chain), that effect must be monitored by a sensor and accurately represented as a signal. The more accurate the representation, the more accurate the control can be.

Furthermore, if the sensor characteristics change, the signal will still be brought to a match with the reference signal, but the variable it represents will no longer be maintained in the same condition. If the temperature-sensing element of a thermostat goes out of calibration, the thermostat will still think it is controlling the same temperature and will keep its movable contact nearly at the same position as before, but the room temperature will be controlled at a different level.

The only aspect of a control loop that is under reliable control, therefore, is the sensor signal. The external counterpart of that signal remains under reliable control only as long as the sensor keeps its calibration accurately. So, if we had to pin down any one aspect of the loop to be "the" controlled aspect of the situation, we would have to choose the sensor signal. Sensor signal = perceptual signal; hence, control of perception.

I think that my way of defining output and control is the least ambiguous. After all, if you define output at a place where disturbances can have an effect, you can't reason backward to the power or force output of the control system just from knowing the state of the variable called "output," because disturbances are contributing an unknown amount to the state of that variable. It seems strange to me to define output in such a way that by knowing the output you can't deduce what the control system is putting out. I don't object to looser usages for the sake of convenience, but when we want to avoid misunderstandings, I think my usage is the least ambiguous.

You say: "Perception is controlled by actions; actions are controlled by perception. Remember the loop!" Let's not confuse "control" with "affect." Control entails bringing a variable to a specified state and keeping it there. Perceptions don't bring actions to specified states and keep them there. It's the variations in the actions that bring perceptions to specified states, despite disturbances that bend to change their states. If you add a disturbance to the actuator output of a control system, the control system will alter its own output effects, not keep them the same.

In ordinary environments, the loop is asymmetrical. There is power gain going through the organism, power loss going through the environment. The part of the loop with the power gain does all of the controlling.

Hans Blom: Bill, you say: "When we study human behavior, we aren't comparing it with some 'optimal' or 'best' way of controlling. We're just trying to understand what people are actually controlling under various circumstances. In some regards, people control things very well indeed, by clever means that surpass what any engineer knows how to build. In other ways, people control stupidly and poorly, and suffer the consequences." That is not my impression. In my opinion, in the billions of years of experimentation through evolution, people (and organisms in general) have found superb ways to realize their goals. If we think that they are stupid, then we are in error. we just have not properly identified their (many!) goals. This is in line with your remark: "Much of the apparently chaotic nature of behavior becomes more understandable when we ask about higher-level goals." In my world view, an organism's behavior is perfectly in line with its top-level goals. Reaching idiosyncratic goals can, of course, be hindered by the laws of nature and of society. Every organism is always at its own local optimum. Of course, we might not agree with its definition of optimum and think that it is just plain stupid. We might even have convinced the organism of that "fact." I realize that this is a personal world view that can in no way be proven. Nevertheless, it is one of my basic life rules, until a better-working one appears. By the way, your use of "suffer the consequences" applies in any case. Behavior has unforeseeable short- and long-range side-effects, always. Our perception is limited, although training might improve things slightly.

You say: 'The highest-level goal is to win the contest, not to jump as high as possible." How do you know? The rules of the game are usually considered to be as follows: when I invent a hypothetical situation, I know what goes on in that situation, because I invented it. You go against the rules here. I say, in effect, "assume that X," and you reply "no, I cannot assume X, I assume Y." You do not play according to what I think the rules are. When I think of a reason, I can only come up with the suggestion that high jumping looks different to you than to me. Your high jumper wants to win the contest. My high jumper really wants to jump as high as possible; he is not interested in winning the contest since he already knows that he is by far the best of those he meets today. No, he is setting his sights much higher. he is training for the next Olympics. He has to compete not with his direct competitors this day, he has to compete with the figures in the world records book that he studies every day. But not even that is enough. He knows that a world record holds only for six years on average. He wants to do better than that and hold the record for many years to come. He will just give this jump his very best effort.

Are these extra perceptions helpful in seeing the situation differently? You could have been right. Your understanding might have explained somebody else's behavior. But in different persons identically looking actions can result from completely different motives. A few lines later you do seem to take that position: "You can't tell what a person is doing just by looking at what the person is doing." And later again: "Pure reason isn't going to identify the actual variable under control by a given person in a given circumstance. A guess about what someone is controlling for could be quite right, or quite wrong." Yes.

You say: "My job is actually easier than yours. I'm not trying to optimize anything—just to match the behavior of a model with that of a real human subject." I have to be precise here: our jobs are very similar. You *are* trying to optimize something: you are trying to find an optional match between a model and a real human subject.

You say: "Of course, real control engineers know a lot more than I do about the design of complex control systems..." Maybe, maybe not. Anyway, that extra knowledge might not account for much when it comes down to designing good control systems. After all, there is not much good theory around to travel by. "Feeling" and "intuition" are required as substitutes for knowledge. I don't think you lack those.

The question of "control" versus "affect" seems to have to do with either intended versus unintended or full versus partial correlation. In either case, it has to do with our limited predictive powers. The first raises the question of what it means to "intend" or to have "goals." The second raises the problem that actions will always have effects in addition to those "intended. " My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very "noisy" perceptions. Control must always be limited; the world is just too complex for our three pounds of brains to model it and our 50 pounds or so of muscles to subdue it.

In engineering, we take great liberty in defining inputs, outputs, and systems. I can take for an input anything that I can manipulate, and for an output anything that I can measure. A system is anything in between. One person's choice might differ from another one's.

Bill Powers: Hans, I don't think many evolutionists would agree with your statement that “in the billions of years of experimentation through evolution, people (and organisms in general) have found superb ways to realize their goals.” Evolution doesn’t optimize anything; it just weeds out unworkable organisms. What’s left is just barely good enough to survive—for a while.

I would have to agree with your implication that organisms control as well as they can. That’s a matter of definition. But in looking at the state of our world, I am not greatly impressed with the way people control for social harmony, economic viability, or maintenance of an environment fit to live in.

You say: “In my world view, an organism’s behavior is perfectly in line with its top-level goals.” I think you’re defining top-level goals from outside of the organism. When I speak of goal-seeking, I’m not normally dunking of “goals” like maintaining the life-support system and combating invasive microorganisms, or even “surviving”—the unlearned goals that I assume to drive reorganization. I’m thinking more in terms of the learned goals, things like being a good person, making a decent living, and so forth. I don’t think that people are particularly adept at constructing systems of goals that hang together, are consistent with each other. Most of the people in the world live in poverty, hunger, and illness. I don’t see how you can claim that they are optimal control systems.

In offering alternatives to the highest-level goal that you suggested (jumping as high as possible), I wasn’t denying that some people might actually have the goal of jumping as high as possible. I was only pointing out that other goals are equally plausible, and, in my experience, more common (particularly when you ask what the *immediate* goal is). In explaining to me that in different persons identical actions might come from different motives, you’re simply echoing my point.

You say: “You *are* trying to optimize something: you are trying to find an optimal match between a model and a real human subject.” You’re a pretty slippery customer. What you say is true: I’m controlling for the best fit between the model and the real behavior. Achieving this requires the same sort of trial and error that tuning a radio or focusing a lens requires, because the amount of error doesn’t tell you which way to move, and there’s no a priori way to specify the magnitude of the effect at the minimum (or maximum). This sort of control does happen. It’s not very common. And it’s not very tight.

Same subject: “My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very ‘noisy’ perceptions.” They do that only as well as the statistics and the accuracy-time tradeoff permit. I don’t worry much about extracting signal from noise; most of the behaviors we observe work at signal levels where noise can be neglected.

Then you say: "Control must always be limited; the world is just boo complex for our three pounds of brains to model it and our 50 pounds or so of muscles to subdue it." Well, I won't be nasty and remind you of how wonderful our evolved control systems are supposed to be. What's really wrong with your statement is the implication that it's hard to find instances of good control. Control is, to be sure, limited—but it's hard to find examples of behavior in which control isn't pretty good by anyone's standards. "Limited" is one of those qualitative terms; the importance of the limits depends on quantitative definitions. Human motor behavior works with a bandwidth of only about 25 Hz—certainly too limited to enable us to balance on end a stick one inch long. On the other hand, this bandwidth seems to be just sufficient to handle most of the disturbances that actually occur on scales that matter to us. On those scales, the limitations are irrelevant.

You say: "In engineering, we take great liberty in defining inputs, outputs, and systems." I think this is one of the reasons that engineers failed to come up with PCT. When you're focused on producing some outcome in the environment, there's no organizing principle for laying out the control system. You can put your stabilizing filters in the input function or add little loops anywhere you like that will do the job. The result is that there are no real principles of design in control engineering (that I know of). There are plenty of principles, but none having to do with how to design the functions of a control system in some systematic way. Basically, you kludge up a design that looks as if it will work, then buckle down to analyzing what you designed.

The PCT approach is to define the problem in terms of sensed variables: it is the sensed variable that will ultimately be controlled, so it should represent something specific in the environment to be controlled. The engineer can violate this principle, because the engineer knows what is to be controlled. But if the control system is in an organism, its perceptions have to be useful in a variety of higher-level systems, and they can't have haphazard relationships to the outside world. This forces the modeler to propose a consistent set of definitions of input, output, system, and environment.

I think that a little more systematicity would also help control engineers, but that's their business.

Hans Blom: Bill, you confuse "optimal" (an engineering word with an exact meaning) with "good" (a moral categorization) in both of these remarks: "I would have to agree with your implication that organisms control as well as they can. That's a matter of definition. But in looking at the state of our world, I am not greatly impressed with the way people control for social harmony, economic viability, or maintenance of an environment fit to live in." "Most of the people in the world live in poverty, hunger, and illness. I don't see how you can claim that they are optimal control systems." The "optimal" of engineering means only that some system reaches its grand overall goal as closely as possible, by definition. Engineering is not concerned with the question of whether that goal is "good." Engineers are, though. In my own personal, idiosyncratic world model, I tend to equate "optimal" with "good" (subjectively, for that person, given his / her opportunities, limitations, and life plan). Maybe that provoked your remarks.

You go on to say: "I don't think that people are particularly adept at constructing systems of goals that hang together, are consistent with each other." In optimal control theory, there is only one "supergoal" that can be controlled. There can be subgoals, however. It would be possible to declare the two (seemingly conflicting) goals "drive in the middle of the road" and also "drive one yard to the right of the middle." But then you would have to combine them into one goal. This can be done, for instance, by stating that the first goal is twice as important as the second goal, or that the first goal is 100% important during the first leg of the journey and 0% thereafter. No conflicts here. Again, I think, "conflict" is a uniquely human word with a moral implication.

I had remarked: "My point is that the human perceptual and conceptual systems are so beautifully designed that they even extract information from very 'noisy' perceptions." You commented: "They do that only as well as the statistics and the accuracy-time tradeoff permit I don't worry much about extracting signal from noise; most of the behaviors we observe work at signal levels where noise can be neglected. " This is certainly true in the domain of muscle control. But is it also true in the other domains which concern you like "being a good person; " "making a decent living," and so forth?

You say: "The PCT approach is to define the problem in terms of sensed variables: it is the sensed variable that will ultimately be controlled, so it should represent something specific in the environment to be controlled." Modern control theory thinks differently. It is, of course, the sensed variables that are our only source of information about how our actions affect the objects that we want to control. But the control problem is not necessarily to bring some variables to some prescribed values and keep them there. That is, of curse, a legitimate field for study, but control theory is far broader. By the way, I think that your use of the notion "reference level" confuses some psychologists and their ilk into having to think about "homeostasis. " Recognition of this confusion might make the PCT approach more acceptable to journal editors and referees.

Bill Powers: Hans, if you're trying to wrap up an entire organism as a single hypercomplex control system, I suppose you would have to look for some grand overall system and a single overall purpose. That isn't the approach in hierarchical perceptual control theory. There might be many highest-level control systems acting in parallel, with relative independence. Of course there is an overall control system in the HPCT model, too, a reorganizing system, but it isn't concerned with learned behavior. Its reference levels and perceptual signals are built-in, and its mode of action is to reorganize the rest of the system. It isn't really a single entity, but a collection of control systems concerned with maintaining the life support systems, each one being concerned with a specific variable.

You say that I "confuse 'optimal' (an engineering word with an exact meaning) with 'good' (a moral categorization)... The 'optimal' of engineering means only that some system reaches its grand overall goal as closely as possible, by definition." I'm sort of between these meanings. If there are two control systems with incompatible goals inside the organism, clearly they are going to expend a lot of energy canceling each other's efforts. This is suboptimal under certain assumptions: that energy expenditure is probably a cost to the whole organism and that reduction of the control range resulting from conflict reduces the ability of both control systems to counteract disturbances. These losses of ability aren't "morally" bad, but the organism would be able to control over a wider range and for a longer time if they were not present. Of course, given the conflict, the control systems are in fact coming as dose as possible to reaching their goals. But with a suitable adjustment of the system organization, they could come a lot closer. A great deal of psychotherapy is aimed at helping people resolve conflicts; I suppose you could say that helping them is a moral choice, but it does have engineering overtones.

"In optimal control theory, there is only one 'supergoal' that can be controlled." Can you explain why this has to be true? What if there is more than one control system operating at the highest level of organization? Of course, you could make up some "supergoal" having to do with an optimal balance between these systems, but in that case the criterion of optimality would be in the eye of the beholder—there would be nothing in the system itself trying to achieve that optimality.

I think that one of the legacies of traditional psychology is a general impression that human behavior is complex and chaotic, with regularities appearing only as statistical averages, and with the future being a matter of rather shaky predictions. PCT, once you get used to seeing the things it calls to attention, shows a very different picture. Most behavior is highly regular and closely controlled; there is very little left to chance.

If this were not true, the world we experience would be very different. People would keep getting lost on the way to work; buildings and houses would constantly be falling down, or fail to have doors or windows, or be located in inaccessible places. Cars, if they ran at all, would always be crashing into each other or wandering off across fields. Nobody would know how to grow crops, or harvest them, or transport the food to some regular destination, or how to cook the food or keep it from spoiling. Most of the things that we use, encounter, or rely upon wouldn't even exist.

What astounds me is the way in which psychologists could have looked at the endless regularities of human existence, mostly maintained by and completely products of human efforts, and failed to recognize them. It is terribly naive just to take the world the way you find it without asking how it could possibly be that way. Psychology has focused on unusual side-effects, on tiny irregularities, and has failed to see the massive regularity that characterizes all living systems and the environments they have shaped to fit their wants. The signal-to-noise ratio in most aspects if life is very, very high. That has not prevented scientists from concentrating on the noise and ignoring the signal.

"It is, of course, the sensed variables that are our only source of information about how our actions affect the objects that we want to control. But the control problem is not necessarily to bring some variables to some prescribed values and keep them there." No, I have never said it was. PCT leads to HPCT, in which higher levels of control act by varying the reference signals for lower systems. They do so as their way of controlling derived perceptions, more generalized perceptions. Those systems, in turn, have their reference levels adjusted by still higher systems, concerned with still more abstract perceptual variables. The only dissonance between this view and your ideas of optimal control has to do with your assertion that at some level there is a single highest control system with a single highest goal.

As to your criteria of optimality, they are completely discretionary. I don't see any reason to suppose that organisms have adopted such criteria or seek to realize them. You're talking about engineers building control systems, not the processes by which living control systems evolve. The engineer can, by choice, combine all lower goals into supergoals, but there is nothing that compels us to suppose that organisms do the same thing—except when they're trained as engineers.

All that the brain knows about the external world comes to it in the form of perceptual signals in the afferent neural pathways. There is no other way for that information to get into the brain. If the brain wants to control the position of a real glass on a real table, it's out of luck: it doesn't have any way to know about the real glass and the real table. It can, however, adjust its output signals so that a neural signal representing the glass can be manipulated to achieve a certain relationship with a neural signal representing the table. *That* the brain *can* do.

I should think that all of this would be self-evident to any engineer who has ever actually built a working control system. A real hardware control system can't interact directly with the physical plant it is controlling. All it can do is alter its electronic output signals and see what happens to the signals being generated by its sensors. That's all it knows about what is happening outside it. If the sensors jump out of calibration, the control system will happily continue controlling the miscalibrated perception, while the technician in charge rushes to hit the *Stop* button What is controlled is *only* what is perceived. One hopes that what is perceived has some relationship to what is, but that is something that has to be determined indirectly.

This simple concept which should cause no problems for any control engineer causes immense problems for conventional sciences of life. The reason is that these conventional sciences ignore the difference between what is perceived and what is—at least when they're trying to explain behavior. And not having any experience with real system design, it seems perfectly reasonable to such conventional scientists that a stimulus input from real objects in the environment should be able to cause motor outputs that steer the organism through a variable environment along a path to the cheese or the mate or whatever. What's the problem? You can see them doing it, so it must be easy.

If you're an engineer watching an organism behave, you will have a hard time making your mental model behave in this simple cause-effect way. You will notice that the eyes keep moving around, that the head moves and bobs up and down, that the steps are a little imprecise and slightly wobbly, that things in the environment are shifting around. Being a person who is charged with making systems actually work, you will wonder how the organism gets away with such imprecision of action—where are all the stimuli coming from that cause the corrections of the little mistakes and overshoots and hesitations? How does the environment know that it should stimulate the organism just in the right way to correct for a previous stumble? How does that little unevenness in the path send just the right stimulus up the spine to make just the right muscles change their tension to keep the leg from jamming into the ground or flailing in empty air on the next step? Any engineer who pays attention in a professional way to the claims of S-R theorists would soon walk away shaking his or her head. No way!

Unfortunately, engineers seem to abandon their normal professional attitudes when they start trying to explain behavior. They start listening to the psychologists and physiologists and neurologists who think that behavior can just be "generated," open-loop. Perhaps they're just being polite because they're on another scientist's turf. They say, "Oh, is that how it works? OK, you must know what you're talking about; I'll see if I can make that work." And, of course, they can make it work. Good engineers can make any damn fool idea work. They can build an arm that's as solid as the front end of a Mack truck, equip it with precision bearings and gears and stepper motors, compute the driving signals using 80-bit floating point arithmetic, and make the arm move exactly as wanted. The smart ones must surely realize that this is *nothing* like the way a human arm works. But the psychologists see what they've done, and nod wisely. It works just the way they expected.

PCT is all about the realization that human systems simply can't work that way. Their outputs are rubbery and imprecise; their neural computers are good to maybe 1% at best; they don't sense everything in the environment that might interfere with the action Yet they work precisely and well, for four score years and six. A person with his little 1% analog computers can get out of bed in the morning and perform one action after another all day long, each action starting where the last one left off, and 16 hours later end up exactly at the side of the same bed, with no cumulative errors at all. Only one kind of system can accomplish that sort of behavior: a negative-feedback control system.

Hans Blom: I enjoy reading / scanning CSGnet a lot; I have derived many eurekas from it (not in the sense of discovering new "truths," but in the sense of gaining new perspectives), and I have come to respect Bill Powers' view of reality. The following remarks are probably more meta-science than science. But many of these discussions are, aren't they?

In systems science, we have the notion that any model accomplishes a particular end. You develop a model with a certain goal in mind; the goals might be different for different modelers. Models can be viewed as theories: you want to summarize all findings within a limited scientific domain in a certain form, e.g., a block diagram. Models can be viewed as tools: you want to encapsulate all properties of a system that you deem important into a simplified form, so that you can control the important aspects of an otherwise too-complex reality. Models can be viewed as predictors or extrapolators: if something happened in the past, it might happen again in a similar way. In *all* cases, we have to understand that each and every model is a simplification of reality, in which we leave out those aspects that we deem unimportant. Therefore, each model is a personal choice: what is unimportant to you might be the most important thing in the world to another person. Or, as the saying goes amongst control engineers: one person's noise is another person's signal.

Of course, such a personal choice might be picked up by others and become part of culture—but only if those others agree with how you split the world into "important" and "unimportant." Sometimes, agreeing is easy: color does not contribute to a body's mass. In other cases, it's not that easy: do people have free will? You might protest that "free will" is a badly defined notion. That is true. But so are "color" and "mass"; no two people or measuring devices will perceive exactly the same color or mass. You might complain again and say that the mass that two well-calibrated scales measure when exposed to the same body is *practically* the same. But that depends upon the practice at hand. In real life, we frequently (always?) seem to have to deal with fuzzy notions. In many cases, this fuzziness does not matter, but in others it can matter a great deal.

We each have a personal, emotional investment in our models. They encapsulate what *we* think is important and leave out what we believe is unimportant. Models are personal creations, much like works of art, that we experience as the best that we can produce. On this net, Bill defends what he sees as important. Of course. But so does everyone else. Isn't that one of the central tenants of your theory, Bill?

This brings me to the issue that, in my opinion, is expressed too little in PCT philosophy. Control is about *control*. You focus on *perceptions* as the important things—and, as a concomitant, on which perceptions are controlled. I have a different ordering of things important. Prime is that we have *goals* (reference levels, as you call them); a control system is a device that allows us to reach or approach those goals in the best possible ways, given our biological and mental limitations. This is also the orthodox control engineering vision. You have a goal, so go design a system that makes it come true. Use the information that the available sensors provide in the best possible way, using any type of processing and data storage that is available or can be newly designed. Control engineers do it this way, and evolution as well, I think. In control engineering, theory has its part; it provides a number of well-proven (partial) solutions. Hunches, trial and error, too, have their parts. No new design is exactly the same as a previous one, alas.

Does this difference in focus matter, you might ask? Yes, I think so. In science, it seems as if we have left all "grand unified theories" behind—although physics is still searching. It seems as if there are no "first principles"; you can go deeper and deeper all the time, if you have the resources. First principles seem to be theories as well. And they are practically useless to explain the world in all its complexity. The formulas of quantum mechanics are barely able to "explain" the movement of *one* electron around *one* proton (the simplest atom that exists), but anything more complex is beyond its powers of synthesis. The synthesis problem is much older, of course: the classical three-body problem of classical mechanics does not allow precise long-term predictions. We are now mentally just coming to grips with these strange facts: that even if first principles are given, a synthesis based on those first principles might be too complex computationally (and mentally) to derive higher-order laws and "explain" more complex systems. That's what chaos theory is all about. Ask any practical control engineer: the existing theories do not suffice when you design a new control system. Always, some extra creativity is required. It is not that those theories are useless; they are not sufficient. Ask any AI-type who works with expert systems: it is not the "reasoning process" that provides the performance of a knowledge-based system, but the knowledge incorporated into it; the more, the better. But then we start to encounter the „complexity problem": a system with a large number of basically independent "knowledge chunks" starts to show unpredictable and uncomprehensible behavior because of the unforeseen ways in which those chunks (sometimes) interact. The result is that the paper model cannot explain or predict anymore. You actually have to *build* it and *run* it to see how it behaves. Philosophers who study culture start to recognize *the same* thing: post-modernists say that the time of the "grand stories," of the ideologies, is over. It is the 'little stories," the personal, subjective accounts, that are the important things that build up the world (and, if generally accepted, might grow into "grand stories").

In my view, no model is wrong, unless it is internally inconsistent. Of course, *any* model is wrong in the sense that it must necessarily be incomplete. In another sense, a different model might be right *as well:* it just has a different purpose (focus) and is based on different notions of what is important. This is true for all models, even PCT models—unless you talk in abstractions that can neither be proven or disproven. It follows from the basic notion that *every model is an approximation*.

If you can accept that different models reflect different goals and therefore incorporate and / or explain different observations, what is a fact to one modeler can be noise to another. A concomitant of this is that a model is (approximately) valid only within some restricted domain. It might "explain" a certain set of observations, but it is without value, or simply wrong, outside its domain. Einstein's E = mc2 certainly does not relate someone's "psychic energy" to his or her body weight.

Don't underestimate statistics. Astronomical data that remain from the days of Kepler show small and large measurement errors. Newton's laws could never have been derived without discarding quite a lot of outliers and assuming that the theory need not *exactly* fit the measurements. Yet, Newton's laws have shown their value. But they, too, are approximations, as Einstein showed. And, undoubtedly, Einstein's relativity theory is an approximation as well.

Bill's hierarchical control model consists of a multitude of simple, functionally identical blocks. The model is an elegant simplification, but we know that the brain isn't quite that homogeneous, neither at the cell level nor at the level of configurations of cells (wiring). Bill, you can marvel at the beauty of your model (it *is* elegant!), yet acknowledge that even in its very basics it cannot possibly be correct.

But that often does not matter much. One system can be modelled in a great many different ways, yet these models can *functionally* show (approximately) the same behavior. This I consider a basic conflict in your model: on the one hand, you want your model to represent physiology as accurately as possible; on the other hand, you want it to show the same *functional* behavior as a human. We are, I think, still very far from the point where we can link the lowest levels (cells, synapses) with the highest levels. In my opinion, and based on the arguments that I presented above, establishing such a link might be impossible in theory, as well.

As has often been noted on this net, things that "actually exist in nature" will forever remain outside our grasp. The best thing we can do is build *models* of what is out there. You know this, Bill, yet it seems that you cannot really accept it. What we require of a model is (a) that it is internally consistent, and (b) that it is consistent with our observations of the "real world." The problem lies in the latter, where we encounter the limitations. We cannot take into account every observational detail. We have to select. And *how* we select depends upon both what we deem important and what we have as capabilities, i.e., we make a personal choice based on our personal goals but within our personal limitations when we build our model. I strive for what I want, building upon what I already know. This is true in mathematics, in control engineering, in life.

Model or theory building is basically a creative process, in which you suddenly have this eureka-feeling of "yes, that's it!" But then science expects you to "prove" your model or theory, and you suddenly find that the theory does not explain all of the data or does not explain with full accuracy. That is when we have to introduce notions like "noise" (small discrepancies that we choose to disregard), "outliers" (large discrepancies that we choose to disregard), "statistics" (can I get an impression of how well my new theory fits the observations despite the fact that I disregard so much?) and things like that. Finally, a theory might start to lead its own life and be taken more seriously than the data. Bill, I assume that you, too, take Newton's laws more seriously than the data that they were originally based on, and more seriously as well than a great deal of more recent measurements.

All of the notions that you use are high-level abstractions, much like "force," "pressure," and "temperature," which have no objective existence but are cultural notions, ways of looking at what surrounds us. In every case, philosophers will tell you, we could have arrived at different but equally valid notions. To use a simple example: you use feet and Fahrenheit, while I use meters and Celsius.

As Rick Marken can tell me so eloquently: “It’s all perception.” Translate this into: “It’s all your own personal subjective theory / model of what’s out there,” and you are dose to what I want to say.

As you can see, Bill, my ‘life model” is, in many ways, different from yours. Why? Our models are based on different data, on different perceptions of what is important, and on different goals. My model has been built up through my experiences that have gradually taught me (a) how to perceive (what to notice, what to disregard), (b) which goals to set (the things that I have come to consider important) and (c) how to act (through the goal-reaching skills that have worked for me).

Everybody has one goal in common, however personal that goal looks: to make the world more controllable / understandable. Every trick in the book—as well as every new one that you can think of—is used to reach that goal. One trick is to observe others and see how they control; maybe (who knows?) their methods will work for me, too. Let’s be inclusive, not exclusive. Let’s find the best tricks and use all of them combined in our personal repertoire. Please take this contribution in that vein. As you might have noticed, I take your “life model” seriously. It provides a much needed additional perceptive. Yet, allow me to think that I, given different perceptions, might have discovered a “life model” that might have some value as well, even if it does not coincide with yours. In works of art, I often find it difficult to say which painting or sculpture is “better” than another. I am slowly discovering that I have a similar problem with scientific theories.

Bill Powers: Hans, you say: “In systems science, we have the notion that any model accomplishes a particular end.” Yes, in the sense that any model that actually works does *something.* But there are two kinds of ends-achievement going on in PCT modeling. One is the modeler’s goal of constructing a model that behaves like the real system. The other is the model’s goal of bringing some perceptual representation of its environment to a reference-state endogenous to the model. If you construct a food-seeking model that depends on balancing smell intensities in a bug’s antennae, but get the sign of the perceptual computation wrong (a - b instead of b - a), the bug will seek a goal, all right, but it will be the goal of traveling away from the food. So the model, while achieving its own goal, will not achieve the modeler’s goal. The modeler wants the bug to want to get near the food and so will reverse that sign, altering what the bug-model perceives to make the outcome the same as what the modeler wants.

And you say: “You develop a model with a certain goal in mind; the goals might be different for different modelers.” What I see missing in systems science is the concept of systems that have their *own* goals. That is, the system is designed to accomplish what the modeler wants done, but the idea that the system itself might want something doesn't seem to be addressed. Am I wrong about that? I admit that artificial devices aren't asked very often what they want, nor does it matter, but when we're modeling the modeler, we have to put the goals into the model.

You say: "Models can be viewed as tools: you want to encapsulate all properties of a system that you deem important into a simplified form, so that you can control the important aspects of an otherwise too-complex reality." The question remains, who does the controlling toward whose ends? Your statement seems to imply that the model's behavior is there only to satisfy the modeller's goals. This says that the model is not a model of the modeler, but of some device to be used for achieving the modeler's purposes. How, then, do we model the modeler, whose goals aren't being given by some other person to suit that other person?

You also say: "Models can be viewed as predictors or extrapolators: if something happened in the past, it might happen again in a similar way." I can agree to this in a very broad sense, but I wonder if it's the same sense you mean. Models in PCT aren't designed to produce particular behaviors under circumstances that led to those behaviors in the past. The *components* of these models could be seen that way—if a comparator has always produced a certain error signal given a particular reference and perceptual signal, we expect it to go on behaving that way. This is what we mean when we describe each function box with a mathematical form. We observe or propose that this function has been performed by that box in the past, and we predict that it will continue to perform the function.

A control-system model can be designed, on the other hand, to produce behavior like that of the real system, quite accurately, in the presence of conditions that have never occurred before. We can measure the control parameters for simple pursuit tracking, for example, and predict how a teal person will perform in a new task with a new pattern of movements of the target, and *with a second disturbance applied directly to the cursor,* which was not present when the parameters were evaluated. Now the model is presented with new conditions (as is the human subject), and the model still behaves just like the subject. This is not exactly extrapolating from past performance, is it? At least it's a kind of extrapolation that is very different from just observing disturbances and the behaviors that follow them, and predicting that recurrence of the same disturbances will produce the same behavior.

"In *all* cases, we have to understand that each and every model is a simplification of reality, in which we leave out those aspects that we deem unimportant." Yes, indeed. The trick is to know when you're leaving out or simplifying something vital. You find this out when you match the model's behavior to the real behavior, or when you change conditions in a way that brings the omitted parts into play. But this is the whole modeling game, isn't it? You get the model to work in as simple a form as possible, then change the conditions until the model stops behaving like the real system. The way in which it fails can sometimes be traced to simplifications or omissions, in which case you go back and use a more detailed model. Other times, the model fails completely, and you have to reconsider it from scratch. The PCT models we use in tracking experiments today represent a long history of wrong guesses, although they're still so simple that it might seem impossible that they were overlooked in the beginning.

"Therefore, each model is a personal choice: what is unimportant to you may be the most important thing in the world to another person." In principle, maybe. In practice, it doesn't feel that way. Some models just don't work no chatter how hard you try to make them work. I suppose you could invoke psychoanalysis and say that if a model fails, its inventor really didn't *want* it to work. But it's hard to believe that when you can see a model designed exactly as you wanted it to be designed that behaves in a way completely different from the real behavior you thought you were modeling. No matter how much you like the model, no matter how many of your private beliefs or prejudices it expresses, if it doesn't work, it doesn't work, and there's no way but self-delusion to make it seem to work.

While I don't think that any models are the last true words about how nature works, I think that some models are definitely better than others. This isn't self-evident if you just construct conceptual models and never test them experimentally. It isn't self-evident if the models are simply descriptions of observations (there are countless ways of describing the same observations). The relative worth of models can be seen only when they're expressed as working simulations that can generate behavior out of their own properties. When you've committed yourself to the point of constructing a working model, there is no way you can make the model work other than the way you designed it to work—and if the way it works doesn't resemble the way the real system you're modeling behaves, you've just shown that your model is wrong.

"Of course, such a personal choice may be picked up by others and become part of culture. But only if those others agree with how you split the world into 'important' and 'unimportant.'" This is a different subject: not which model is best, but what aspect of experience you want to model. In PCT we generally agree that we want to model ordinary behavior: what people do in daily life, at many levels. We're not trying to model chakras or satori or survival after death or ghosts or metabolism or lots of things like that. Just plain vanilla behavior. Generally we took at the same things that other theories have looked at: environmental events near organisms, actions and their consequences produced by the muscles of organisms, perceptions of various kinds, nervous systems and their possible functions. We aren't emphasizing or de-emphasizing any of these phenomena; we're just asking what makes them work the way they seem to work.

"You might protest that 'free will' is a badly defined notion. That is true. But so are 'color' and 'mass'; no two people or measuring devices will perceive exactly the same color or mass." That's a bit qualitative for a valid comparison. We can characterize color and mass well enough to reproduce them within a few parts per thousand and agree on perceptions of them within a few percent, but I defy anyone to reproduce "free will" in any way that can be quantified. No two people perceive color or mass *exactly* the same, but no two people perceive free will even *approximately* the same: many claim they don't even perceive it. Let's at least compare apples with round things.

"In real life, we frequently (always?) seem to have to deal with fuzzy notions. In many cases, this fuzziness does not matter, but in others it can matter a great deal." The quality of our lives is vitally affected by fuzzy notions we would be better off without, or at least with, but in sharper form. The point of science, in my mind, is to clarify fuzzy notions or to get rid of them if they are intractably blurred.

"We each have a personal, emotional investment in our models. They encapsulate what tae think is important and leave out what we believe is unimportant." I have some investment in a model of tracking behavior in which the model's simulated handle position follows a course through time that deviates from the handle position created by a person in the same experiment only three to five percent, RMS. I think it is important for the behavior of a model to be as close to the behavior it supposedly models as possible. I'd like it to be closer, but so far can't accomplish that Someone else might consider this sort of match unimportant, preferring ire cream or skiing. Someone else might think that tracking behavior isn't very interesting, considering the problems in Somalia. But anyone who thinks that models of overt physical behavior should reproduce and predict behavior accurately has this model to contend with.

I doubt that the behavior of this model has much to do with my personal emotional investments.

"Models are personal creations, much like works of art, that we experience as the best that we can produce." There's a bit more than that to models that I respect. A model should deal with data that's publicly observable by means on which we can agree and reproduce independently. The reasoning that leads to the model should be laid out in public view in sufficient detail that anyone who understands basic logic and mathematics could recreate the model from scratch if necessary, and come up with the same model. The model should behave the same way in anyone’s hands and should fit behavior correctly as evaluated by any user of the model. I don’t think that very many of these considerations apply to works of art.

“On this net, Bill defends what he sees as important. Of course. But so does everyone else. Isn’t that one of the central tenants of your theory?” Certainly, and I’m glad that you see the theory as correctly describing human behavior.

‘This brings me to the issue that, in my opinion, is expressed too little in PCT philosophy. Control is about *control*. You focus on *perceptions* as the important things—and, as a concomitant, on which perceptions are controlled.” It would be pretty hard to focus on perceptions as the important things without the rest of the control loop. Perceptions aren’t just sort of vaguely “important.” It just happens that when you try to find the variable in a control loop that is the most reliably controlled under the most changes of conditions, it proves to be the perceptual signal. We didn’t pick perceptions as pivotal for private or silly reasons, or just because we’re perception freaks. Perceptions are all that an organism can know about the world outside it. That means you, too. It follows that goals have to be defined in terms of perceptions. You can’t compare an internal goal with an external unperceived object; the object must appear as a perception in the same place where the goal is before any comparison can take place. PCT is about goals, too, and about error signals and input functions and actuators and all of the parts of a control system.

“In my view, no model is wrong, unless it is internally inconsistent.” I guess our views differ. I demand that a model behave *like* the world it is supposed to describe or explain. A model can be internally consistent yet totally at variance with experimental observations. What is “important” has nothing *to* do with this. If a model predicts something unimportant incorrectly, it is still wrong. Models that don’t have *anything to* do with observation and that produce no predictions of behavior *to* be compared with observation don’t even count as models in my world. There’s no reason to take them seriously unless the math grabs you.

“Don’t underestimate statistics. Astronomical data that remain from the days of Kepler show small and large measurement errors. Newton’s laws could never have been derived without discarding quite a lot of outliers and assuming that the theory need not exactly fit the measurements.” “Measurement error” is something very different, quantitatively, from “variance” in psychological observations. You can measure a rat’s running speed in a maze with a measurement error of perhaps 0.1 percent, if you use instrumentation. But the supposed effects of stimulus conditions on that running speed will have a variance of hundreds to thousands of percent. Newton and Kepler were trying to formulate models of celestial mechanics that would predict the positions of planets within the existing measurement errors. If the kinds of statistical methods used in psychology had been brought to bear on this problem, celestial mechanics would consist of the firm statement that the planets are up there, not down here. It is very hard to underestimate the power of statistics as used in the behavioral sciences.

"Model or theory building is basically a creative process, in which you suddenly have this eureka-feeling of 'yes, that's it!' But then science expects you to 'prove' your model or theory, and you suddenly find that the theory does not explain all the data or does not explain with full accuracy. That is where we have to introduce notions like 'noise' (small discrepancies that we choose to disregard), 'outliers' (large discrepancies that we choose to disregard), 'statistics' (can I get an impression of how well my new theory fits the observations despite the fact that I disregard so much?) and things like that." This is a rather remarkable statement, in that in summarizes exactly what I think is wrong in the behavioral sciences. Concepts like noise, outliers, statistics, variance, and so forth were invoked by psychologists as a way of explaining why their theories of behavior didn't predict worth a damn. Instead of blaming the poor results on a mismatch of theory to the organism, they blamed it on the organism. In PCT, any time we get results like the *best* statistical results in conventional behavioral experiments, we look for what is wrong with the model. And we usually find it. Behavior, I strongly suspect with some smattering of data in support, is nowhere near as variable as it has seemed to psychologists viewing it through their theories.

"All of the notions that you use are high-level abstractions, much like 'force; 'pressure; and 'temperature; which have no objective existence but are cultural notions, ways of looking at what surrounds us. In every case, philosophers will tell you, we could have arrived at different but equally valid notions." True, but high-level abstractions are grounded in lower-level ones, down to the level normally accepted in science as "observational"—the level where you can report just how much. How much of *what is* determined theoretically, but the relationships among observations are predicted at a low level of abstraction: how far one trace on a chart deviates from another.

As to the philosophers, it's easy to say that you *could* arrive at a different but equally valid notion. Actually doing that is a bit harder. What I hope for is a model for which *nobody* can think of an *equally valid* alternative. The fact that one might hypothetically exist doesn't bother me much. I'm concerned with the model we do have today, not one that might show up later.

The Hierarchical Behavior of Perception

Richard S. Marken
(Life Learning Associates, 10459 Holman Ave., Los Angeles, CA 90024)

Abstract

This paper argues that the coincidental development of hierarchical models of perception and behavior is not a coincidence. Perception and behavior are two sides of the same phenomenon—control. A hierarchical-control-system model shows that evidence of hierarchical organization in behavior is also evidence of hierarchical organization in perception. Studies of the temporal limitations of behavior, for example, are shown to be consistent with studies of temporal limitations of perception. The control model shows that the perceptual limits are the basis of the behavioral limits; action systems that are capable of rapid response cannot produce controlled behavioral results faster than the rate at which these results can be perceived. Behavioral skill turns on the ability to control a hierarchy of perceptions, not actions.

Introduction

Psychologists have developed hierarchical models of both perception (e.g., Bryan & Harter, 1899; Palmer, 1977; Povel, 1981) and behavior (e.g., Albus, 1981; Arbib, 1972; Greeno & Simon, 1974; Lashley, 1951; Martin, 1972; Rosenbaum, 1987). This could be a coincidence, a case of similar models being applied to two very different kinds of phenomena. On the other hand, it could reflect the existence of a common basis for both perception and behavior. This paper argues for the latter possibility, suggesting that perception and behavior are two sides of the same phenomenon: control (Marker, 1988). Control is the means by which agents keep perceived aspects of their external environment in goal states (Powers, 1973). It is argued that the existence of hierarchical models of both perception and behavior is a result of looking at control from two different perspectives: that of the agent doing the controlling (the actor), and that of the agent watching control (the observer). Depending on the perspective, control can be seen as a perceptual or a behavioral phenomenon.

From the actor's perspective, control is a perceptual phenomenon. The actor is controlling his or her own perceptual experience, making it behave as desired. However, from the observer's perspective, control is a behavioral phenomenon The actor appears to be controlling variable aspects of his or her behavior in relation to the environment. For example, from the perspective of a typist (the actor), typing involves the control of a dynamically changing set of kinesthetic, auditory, and, perhaps, visual perceptions. If there were no perceptions, there would be no typing. However, from the perspective of someone watching the typist (the observer), perception is irrelevant; the typist appears to be controlling the movements of his or her fingers in relation to the keys on a keyboard.

These two views of control have one thing in common; in both cases, control is seen in the behavior of perception. For the actor, control is seen in the behavior of his or her own perceptions. For the observer, control is seen in the behavior of his or her own perceptions of the actor's actions. (The observer can see the means of control but can only infer the perceptual consequences as experienced by the actor). If control is hierarchical, then it can be described as the behavior of a hierarchy of perceptions. Hierarchical models of perception and behavior can then be seen as attempts to describe control from two different perspectives, those of the actor and observer, respectively. This paper presents evidence that hierarchical models of perception and behavior reflect the hierarchical structure of control.

A Perceptual Control Hierarchy

The concept of control as the behavior of perception can be understood in the context of a hierarchical-control-system model of behavioral organization (Powers, 1973, 1989). The model is shown in Figure 1. It consists of several levels of control systems (the figure shows six levels), with many control systems at each level (the figure shows 11). Each control system consists of an input transducer (I), a comparator (C), and an output transducer (O). The input transducer converts inputs from the environment or from systems lower in the hierarchy into a perceptual signal, p. The comparator computes the difference, e, between the perceptual signal and a reference signal, r. The output transducer amplifies and converts this difference into actions which affect the environment or become reference signals for lower-level systems.

The control systems at each level of the hierarchy control perceptions of different aspects of their sensory input, but all of the systems control perceptions in the same way: by producing actions that reduce the discrepancy between actual and intended perceptions. Intended percep-

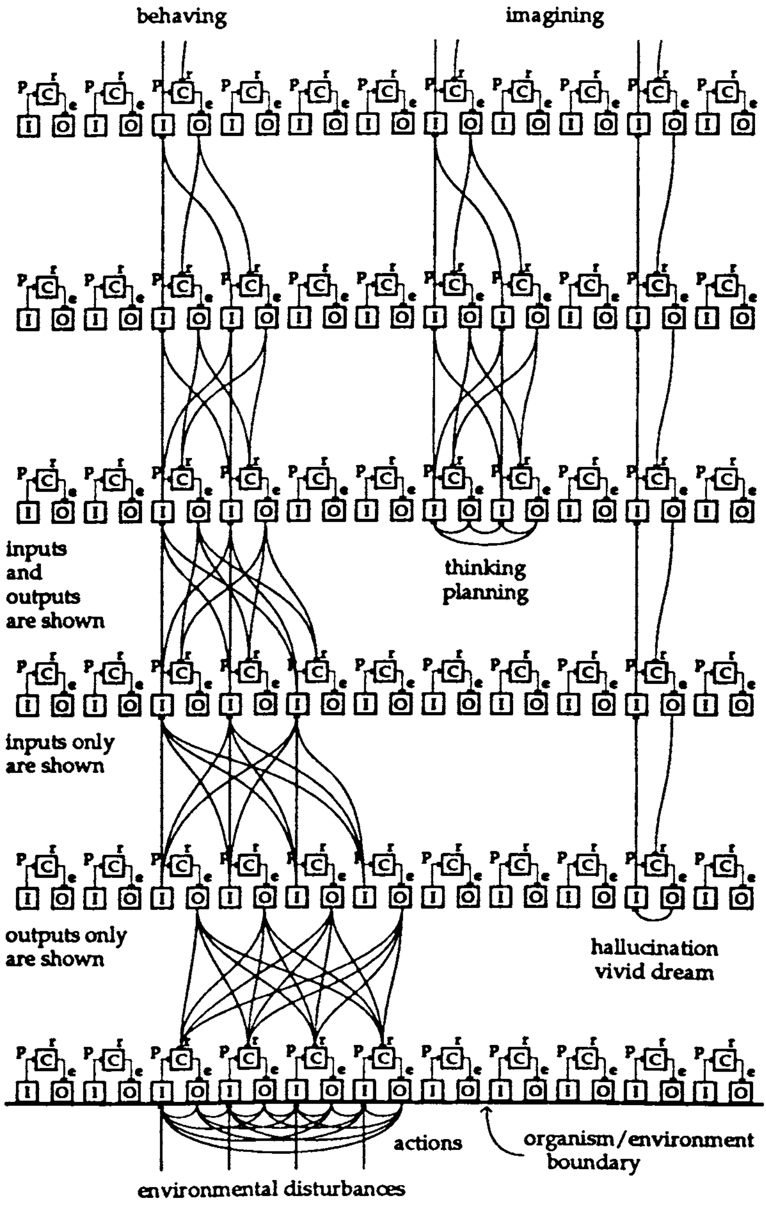

Figure 1. Perceptual control hierarchy (after Powers, 1989, p. 278).

tions are specified by the reference signals of the control systems. The actions of the control systems coax perceptual signals into a match with reference signals via direct or indirect effects on the external environment. The actions of the lowest-level control systems affect perceptions directly through the environment. The actions of higher-level control systems affect perceptions indirectly by adjusting the reference inputs to lower-level systems.

The hierarchy of control systems is a working model of purposeful behavior (Marker, 1986, 1990). The behavior of the hierarchy is purposeful, inasmuch as each control system in the hierarchy works against any opposing forces in order to produce intended results. Opposing forces come from disturbances created by the environment, as well as from interfering effects caused by the actions of other control systems. The existence of disturbances means that a control system cannot reliably produce an intended result by selecting a particular action. Actions must vary to compensate for varying disturbances. Control systems solve this problem by specifying what results are to be perceived, not how these results are to be achieved. Control systems control perceptions, not actions. When set up correctly, the control systems in the hierarchy vary their actions as necessary, compensating for unpredictable disturbances, in order to produce intended perceptions. Indeed, the term "control" refers to the process of producing intended perceptions in a disturbance-prone environment.

Levels of Perception

Powers (1990) has proposed that each level of the hierarchy of control systems controls a different class of perception. Moving up the hierarchy, these classes represent progressively more abstract aspects of sensory input. The lowest-level systems control perceptions that represent the intensity of sensory input. At the next level, the systems control sensations (such as colors), which are functions of several different intensities. Going up from sensations, there is control of configurations (combinations of sensations), transitions (temporal changes in configurations), events (sequences of changing configurations), relationships (logical, statistical, or causal co-variations among independent events), categories (class memberships), sequences (unique orderings of lower-order perceptions), programs (if-then contingencies among lower-level perceptions), principles (general rules perceptible in the behaviors of lower-level perceptions), and system concepts (particular sets of principles exemplified by the states of many lower-level perceptions; see Powers, 1989, pp. 190-208). These 11 classes of perception correspond to 11 levels of control systems in the hierarchical-control model. All control systems at a particular level of the hierarchy control the same class of perception, though each system controls a different exemplar of the class. Thus, all systems at the configuration level control configuration perceptions, but each system at that level controls a different configuration.

The rationale for hierarchical classes of perceptual control is based on the observation that certain types of perception depend on the existence of others. Higher-level perceptions depend on (and, thus, are functions of) lower-level perceptions. For example, the perception of a configuration, such as a face, depends on the existence of sensation and intensity perceptions. The fare is a function of these sensations and intensities. The lower-level perceptions are the independent variables in the function that computes the higher-level perception. Their status as independent variables is confirmed by the fact that lower-level perceptions can exist in the absence of the higher-level perceptions, but not vice versa. Sensation and intensity perceptions can exist without the perception of a fare (or any other configuration, for that matter), but there is no fare without perceptions of sensation and intensity.

The Behavior of Perceptions

From the point of view of the hierarchical-control model, "behaving" is a process of controlling perceptual experience. Any reasonably complex behavior involves the control of several levels of perception simultaneously. For example, when typing the word "hello," one controlled perception is the sequence of letters "h," "e; " l, " "l," and "o." The perception of this sequence is controlled by producing a sequence of keypress-event perceptions. Each keypress event is controlled by producing a particular set of transitions between finger-configuration perceptions. Each finger configuration is controlled by a different set of force sensations, which are themselves controlled by producing different combinations of intensities of tensions in a set of muscles.

The perceptions involved in typing "hello" are all being controlled simultaneously. Transitions between finger configurations are being controlled while the force sensations that produce the configuration perceptions are being controlled. However, the typist is usually not aware of the behavior of all these levels of perception. People ordinarily attend to the behavior of their perceptions at a high level of abstraction, ignoring the details. We attend to the fact that we are driving down the road and ignore the changing muscle tensions, arm configurations, and steering wheel movements that produce this result. Paying attention to the details leads to a deterioration of performance; it is the opposite of Zen behavior, where one attends only to the (perceptual) results that one intends to produce and lets the required lower-level perceptions take care of themselves (Herrigel, 1971). However, while it violates the principles of Zen, attention to the detailed perceptions involved in the production of behavioral results can provide interesting hints about the nature of the perceptual control hierarchy.

The Perception of Behavior

The behavior of an actor organized like the hierarchical-control model consists of changes in the values of variables in the actors environment. An observer cannot see what is going on inside the actor; he or she can only see the actor's actions and the effect of these actions on the external environment. The effect of these actions is to cause purposeful behavior of certain variables in the environment: the variables that correspond to perceptions that the actor is actually controlling. The purposefulness of the behavior of these variables is evidenced by the fact that consistent behaviors are produced in the context of randomly changing environmental disturbances. Thus, a typist can consistently type the word "hello," despite changes in the position of the fingers relative to the keyboard, variations in the push-back force of the keys, or even a shift from one keyboard arrangement to another (from QWERTY to Dvorak, for example).

Since the actor controls his or her own perceptions, the observer cannot actually see what the actor is "doing"; the acts "doings" consist of changing the intended states of his or her own perceptions. The observer sees only the variable results of the actors actions-results that might or might not be under control. For example, the observer might notice that a click occurs each time the typist presses a key. The click is a result produced by the typist, and the observer is likely to conclude that the typist is controlling the occurrence of the click. In fact, the click might be nothing more than a side-effect of the typist's efforts to make the *key* feel like it has hit bottom. There are methods that make it possible for the observer to tell whether or not his or her perceptions of the actor's behavior correspond to the perceptions that are being controlled by the actor (Marker, 1989). These methods make it possible for the observer to determine what the actor is actually doing (i.e., controlling).

Hierarchical Control

The hierarchical nature of the processes that generate behavior would not be obvious to the observer of a hierarchical control system. The observer could tell that the system is controlling many variables simultaneously, but he or she would find it difficult to demonstrate that some of these variables are being controlled in order to control others. For example, the observer could tell that a typist is controlling letter sequences, keypress events, finger movements, and finger configurations. But the observer would have a hard time showing that these variables are hierarchically related. The observer could make up a plausible hierarchical description of these behaviors; for example, finger positions seem to be used to produce finger movements which are used to produce keypresses which are used to produce letter sequences. But finding a hierarchical description of behavior does not prove that the behavior is actually produced by a hierarchical process (Davis, 1976; Kline, 1983).

Hierarchical Invariance

Hierarchical production of behavior implies that the commands required to produce a lower-level behavior are nested within the commands required to produce a higher-level behavior. For example, the commands that produce a particular finger configuration would be nested within the commands that produce a movement from one configuration to another. Sternberg, Knoll, & Turlock (1990) refer to this nesting as an invariance property of hierarchical control. Lower-level commands are like subprograms invoked by programs of higher-level commands. The invariance of hierarchical control refers to the assumption that the course of such a subprogram does not depend on how it was invoked from the program (low-level invariance); similarly, the course of the program does not depend on the nature of the commands carried out by the subprograms (high-level invariance).

Convergent and Divergent Control

The hierarchical-control model satisfies both the low- and high-level invariance properties of hierarchical control. The commands issued by higher-level systems have no effects on the commands issued by lower-level systems, and vice versa. It is important to remember, however, that the commands in the control hierarchy are requests for input, not output. Higher-level systems tell lower-level systems what to perceive, not what to do. This aspect of control-system operation solves a problem that is either ignored or glossed over in most hierarchical models of behavior. how does a high-level command get turned into the lower-level commands producing results that satisfy the high-level command? If commands specify outputs, then the result of the same command is different when there are varying environmental disturbances. The high-level command to press a key, for example, cannot know which lower-level outputs will produce this result on different occasions. This problem is solved by the hierarchical-control model because intended results are represented as a convergent function, which produces a single perceptual signal, rather than as a divergent network, which produces multiple behavioral outputs.

Most hierarchical models of behavior require that a high-level command be decomposed into many lower-level commands to produce an intended result. In the hierarchical-control model, both the high-level command and the intended result of the command are represented by a single, unidimensional signal. The signal that represents the intended result is a function of results produced by many lower-level commands. But the high-level command does not need to be decomposed into all of the appropriate lower-level commands (Powers, 1979). The difference between the high-level command and the perceptual result of that command is sufficient to produce the lower-level commands that keep the perceptual result at the commanded value (Marken, 1990).

Levels of Behavior

The hierarchical invariance properties of the control hierarchy provide a basis for determining whether its behavior is actually generated by hierarchical processes. Hierarchical control can be seen in the relative timing of control actions. In a control hierarchy, lower-level systems must operate faster than higher-level systems. Higher-level systems cannot produce a complex perceptual result before the lower-level systems have produced the component perceptions on which it depends. This nesting of control actions can be seen in the differential speed of operation of control systems at different levels of the control hierarchy. Lower-level systems not only cornea for disturbances faster than higher-level ones; they carry out this correction process during the higher-level correction process. The lower-level control process is temporally nested within the higher-level control process.

Arm Movement

Powers, Clark, & McFarland (1960) describe a simple demonstration of nested control based on relative timing of control system operation. A subject holds one hand extended straight ahead while the experimenter maintains a light downward pressure on it. The subject is to move his or her arm downward as quickly as possible when the experimenter signals with a brief, downward push on the subject's extended hand. The result of this simple experiment is always the same: the subject responds to the downward signal push with a brief upward push followed by downward movement of the arm. An electromyograph shows that the initial upward push is an active response and not the result of muscle elasticity.

The arm movement demonstration reveals one level of control nested within another. The subject's initial upward push (which cannot be suppressed) is the fast response of a lower-level control system that is maintaining the perception of arm position in a particular reference state (extended forward). The behavior of this system is nested within the response time of a higher-level system that moves the arm downward. The higher-level system operates by changing the reference for the arm-position control system. The downward signal push causes the brief upward reaction because the signal is treated as a disturbance to arm position. This is particularly interesting because the signal is pushing the arm in the direction it should move; the lower-level reaction is "counterproductive" with respect to the goal of the higher-level system (which wants to perceive the arm down at the side). The reaction occurs because the lower-level system starts pushing against the disturbance to arm position before the higher-level system can start changing the reference for this position.

Polarity Reversal

More precise tests of nested control were carried out in a series of experiments by Marken & Powers (1989). In one of these experiments, subjects performed a standard pursuit tracking task, using a mouse controller to keep a cursor aligned with a moving target. At intervals during the experiment, the polarity of the connection between mouse and cursor movement was reversed in a way that did not disturb the cursor position. Mouse movements that had moved the cursor to the right now moved it to the left; mouse movements that had moved the cursor to the left now moved it to the right.

A sample of the behavior that occurs in the vicinity of a polarity reversal is shown in Figure 2. The upper traces show the behavior of a control-system model, and the lower traces show the behavior of a human subject. When the reversal occurs, both the model and the subject respond to error (the deviation of the cursor from the target) in the wrong direction, making it larger instead of smaller (any deviation of the error trace from the zero line represents an increase in error). The larger error leads to faster mouse movement, which causes the error to increase still more rapidly. A runaway condition ensues, with error increasing exponentially.

About 1/2 second after the polarity reversal, the subject's behavior departs abruptly from that of the model. The subject adjusts to the polarity reversal, and the error returns to a small value. The model cannot alter its characteristics, and so the error trace quickly goes off the graph. These results provide evidence of two nested levels of control operating at different speeds. The faster, lower-level system controls the distance between cursor and target. This system continues to operate as

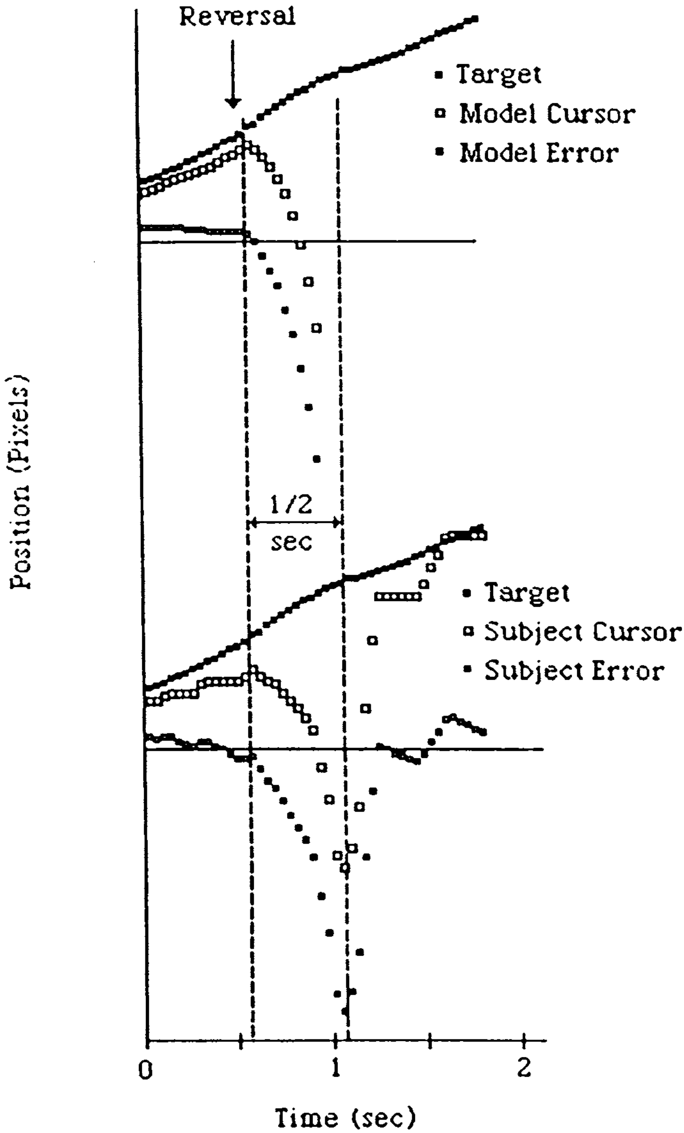

Figure 2. Low-level runaway response to mouse-cursor polarity reversal (after Marken & Powers, 1989, p. 415).

usual, even when, due to the polarity reversal, this causes an increase in perceptual error. Normal operation is restored only after a slower, higher-level system has time to control the relationship between mouse and cursor movement.

Levels of *Perception*

The arm movement and polarity shift experiments reveal the hierarchical organization of control from the point of view of the observer. The hierarchical-control model suggests that it should also be possible to view hierarchical organization from the point of view of the actor. From the actor's point of view, hierarchical control would be seen as a hierarchy of changing perceptions. One way to look at this hierarchy is again in terms of relative timing—in this case, however, in terms of the relative timing of the perceptual results of control actions, instead of the actions themselves.

Computation Time Window

The hierarchical-control model represents the results of control actions as unidimensional perceptual signals. A configuration, such as the letter "h; " is a possible result of control actions, as is a sequence of letters, such as the word "hello." The model represents these results as perceptual input signals, the intensity of a signal being proportional to the degree to which a particular result is produced. This concept is consistent with the physiological work of Hubei & Wiesel (1979), who found that the firing rate of an afferent neuron is proportional to the degree to which a particular environmental event occurs in the "receptive field" of the neuron.

Many of the higher-level classes of perception in the control hierarchy depend on environmental events that vary over time. Examples are transitions, events, and sequences. The neural signals that represent these variables must integrate several lower-level perceptual signals that occur at different times. Hubei and Weisel found evidence of a computation time window for integrating perceptual signals. Certain cells respond maximally to configurations (such as "lines") that move across a particular area of the retina at a particular rate. These are "motion detector" neurons. The neurons respond maximally to movements of configurations that occur within particular time windows. Movements that occur outside of these time windows are not included in the computations of perceptual signals representing motion.

Levels by Time

The hierarchical-control model implies that the duration of the computation time window increases at higher levels in the hierarchy. The minimum computation time window for the perception of configurations should be shorter than the minimum computation time window for the perception of transitions, which should be shorter than the minimum computation time window for the perception of sequences. I have developed a version of the psychophysical method of adjustment that makes it possible to see at least four distinct levels of perception by varying the rate at which items occur on a computer display. A computer program presents a sequence of numbers at two different positions on the display. The presentation positions are vertically adjacent and horizontally separated by two centimeters. The numbers are presented alternately in the two positions. The subject can adjust the rate at which the numbers occur in each position by varying the position of a mouse controller.

The results of this study are shown schematically in Figure 3. At the fastest rate of number presentation, subjects report that the numbers appear to occur in two simultaneous streams; the fact that the numbers are presented to the two positions alternately is completely undetectable. However, even at the fastest rate of number presentation, subjects can make out the individual numbers in each stream. At the fastest rate, there are approximately 20 numbers per second in each stream. This means that there is a 50-millisecond period available for detecting each number. This duration is apparently sufficient for number recognition, suggesting that the computation time window for perception of configuration is less than 50 milliseconds. Studies of the “span of apprehension” for sets of letters suggest that the duration of the computation time window for perception of visual configuration might be even less than 50 milliseconds, possibly as short as 15 milliseconds (Sperling, 1960).

As the rate of number presentation slows, the alternation between numbers in the two positions becomes apparent. Subjects report perception of alternation or movement between numbers in the two positions when the numbers in each stream are presented at the rate of about seven per second. At this rate, an alternation from a number in one stream to a number in another occurs in 160 milliseconds. This duration is sufficient for perception of the alternation as a transition or movement from one position to the other, suggesting that the computation time window for transition perception is on the order of 160 milliseconds. This duration is compatible with estimates of the time to experience optimal apparent motion when configurations are alternately presented in two different positions (Kolers,1972).

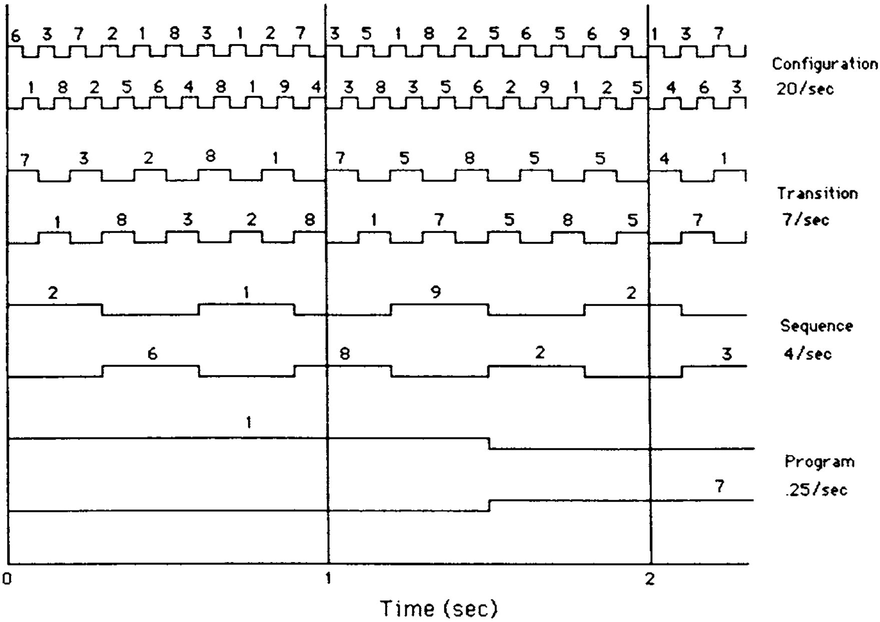

Figure 3. Schematic representation of the results of the number rate adjustment study.

The numbers presented in each stream are always changing. However, subjects find it impossible to perceive the order of the numbers as they alternate from one position to another, even though it is possible to clearly perceive the individual numbers and the fact that they are alternating and changing across positions. The rate of number presentation must be slowed considerably, so that each stream of numbers is presented at the rate of about two per second, before it is possible to perceive the order in which the numbers occur. At this rate, numbers in the sequence occur at the rate of four per second. These results suggest that the duration of the computation time window for the perception of sequence is about 05 seconds. This is the time it takes for two elements of the sequence to occur—the minimum number that can constitute a sequence.

The numbers in the rate-adjustment study did not occur in a fixed, repeating sequence. Rather, they were generated by a set of rules—a program. The sequence of numbers was unpredictable unless the subject could perceive the rule underlying the sequence. The rule was as follows: if the number on the right was even, then the number on the left was greater than five; otherwise, the number on the left was less than five.

(Numbers in the sequence were also constrained to be between zero and nine). Subjects could not perceive the program underlying the sequence of numbers until the speed of the two streams of numbers was about 025 numbers per second, so that the numbers in the program occurred once every two seconds. The perception of a program in a sequence of numbers requires considerably more time then it takes to perceive the order of numbers in the same sequence.

The perception of a sequence or a program seems to involve more mental effort than the perception of a configuration or a transition. Higher-level perceptions, like programs, seem to represent subjective rather than objective aspects of external reality; they seem more like interpretations than representations. These higher-level perceptions are typically called "cognitions." Of course, all perceptions represent subjective aspects of whatever is "out there"; from the point of view of the hierarchical-control model, the location of the line separating perceptual from cognitive representations of reality is rather arbitrary. Behavior is the control of perceptions which range from the simple (intensities) to the complex (programs).

Perceptual Speed Limits

The hierarchical-control model says that all perceptions of a particular type are controlled by systems at the same level in the hierarchy. This implies that the speed limit for a particular type of perception should be about the same for all perceptions of that type. The 160 millisecond computation time window for perception of transition, for example, should apply to both visual and auditory transition There is evidence that supports this proposition Miller & Heise (1950) studied the ability to perceive an auditory transition called a "trill." A trill is the perception of a temporal alternation from one sound sensation or configuration to another. The speed limit for trill perception is nearly the same as the speed limit for visual transition perception found in the number rate adjustment study—about 15 per second. As in the visual case, when the rate of alternation of the elements of the auditory trill exceeds the computation time window, the elements "break" into two simultaneous streams of sound; the perception of transition (trill) disappears, even though the sounds continue to alternate.

There is also evidence that the four-per-second speed limit for sequence perception found in the number-rate adjustment study applies across sensory modalities. Warren, Obusek, Farmer, & Warren (1969) studied subjects' ability to determine the order of the component sounds in a sound sequence. They found that subjects could not perceive the order of the components until the rate of presentation of the sequence was less than or equal to four per second. This was a surprising result, because it is well known that people can discriminate sequences of sounds that occur at rates much faster than four per second. In words, for example, the duration of the typical phoneme is 80 milliseconds, so people can discriminate sequences of phoneme sounds that occur at the rate of about 10 phonemes per second. But there is reason to believe that the phonemes in a word are not heard as a sequence; that is, the order of the phonemes cannot be perceived. Warren (1974) showed that subjects can learn to tell the difference between sequences of unrelated sounds that occur at rates of 10 per second. However, the subjects could not report the order of the sounds in each sequence, only that one sound event differed from another. A word seems to be a lower-order perception—an event perception—that is recognized on the basis of its overall sound pattern There is no need to perceive the order in which the phonemes occur, just that the temporal pattern of phonemes (sound configurations) for one word differs from that for other words.

The Relationship Between Behavior and Perception

Configurations, transitions, events, sequences, and programs are potentially controllable perceptions. An actor can produce a desired sequence of sounds, for example, by speaking sound events (phonemes) in some order. An observer will see the production of this sequence as a behavior of the actor. The hierarchical-control model suggests that the actor's ability to produce this behavior turns on his or her ability to perceive the intended result. Since perception depends on speed, it should be impossible for the actor to produce an intended result faster than the result can be perceived. The observer will see this speed limit as a behavioral limit. An example can be seen in the arm-movement experiment described above. In that experiment, it appears that the time to respond to the signal push is the result of a behavioral speed limit: the inability to generate an output faster than a certain rate. But a closer look indicates that the neuromuscular "output" system is perfectly capable of responding to a signal push almost immediately, as evidenced by the immediate upward response to the downward signal push. The same muscles that produce this immediate reaction must wait to produce the perception of the arm moving downward. The speed limit is not in the muscles. It is in the results that the muscles are asked to produce; a static position of the arm (a configuration perception) or a movement of the arm in response to the signal push (a relationship perception).

Sequence Production and Perception

Some of the most interesting things people do involve the production of a sequence of behaviors. Some recent studies of temporal aspects of sequence production are directly relevant to the hierarchical-control model. In one study, Rosenbaum (1987) asked subjects to speak the first letters of the alphabet as quickly as possible. When speed of letter production exceeded four per second, the number of errors (producing letters out of sequence) increased dramatically, indicating a loss of control of the sequence. The speed limit for sequence production corresponds to the speed limit for sequence perception—four per second.

The letter-sequence study does not prove that the speed limit for letter-sequence production is caused by the speed limit for letter-sequence perception. It could be that the speed limit is imposed by characteristics of the vocal apparatus. However, in another study, Rosenbaum (1987) found the same four-per-second speed limit for production of errorless finger-tap sequences. The speed limit for finger-tap sequence production is likely to be a perceptual rather than a motor limit, because we know that people can produce finger taps at rates much higher than four per second. Pianists, for example, can do trills (alternating finger taps) at rates which are far faster than four per second. Further evidence of the perceptual basis of the finger-tap sequence speed limit would be provided by studies of finger-tap sequence perception. When a subject produces a sequence of finger taps, he or she is producing a sequence of perceptions of pressure at the finger tips. A perceptual experiment where pressure is applied to the tips of different fingers in sequence should show the four-per-second speed limit. Subjects should have difficulty identifying the order of finger-tip pressures when the sequence occurs at a rate faster than four per second.

Confounding Levels

It is not always easy to find clear-cut cases of behavioral speed limits that correspond to equivalent perceptual speed limits. Most behavior involves the control of many levels of perception simultaneously. People control higher-level perceptions (like sequences) while they are controlling lower-level perceptions (like transitions). This can lead to problems when interpreting behavioral speed limits. For example, Rosenbaum (1983) presents some finger tapping results that seem to violate the four-per-second speed limit for sequence perception. When subjects tap with two hands, they can produce a sequence of at least eight finger taps per second. But each tap is not necessarily a separate event in a sequence. Some pairs of taps seem to occur at the rate at which sequences are experienced as events. A sequence of finger taps is an event in the same sense that the sequence of muscle tensions that produce a finger tap is an event; the order of the components of the sequence cannot be perceived. These finger-tap events are then unitary components of the sequence of finger-tap perceptions.

The fact that cetain pairs of finger taps are produced as events rather than ordered sequences is suggested by the errors made at each point in the finger-tap sequence. Errors occur most frequently at the point in the sequence at which a fast pair is being initiated. Errors rarely occur for the second element of a fast pair. This suggests that the errors occur at the sequence level rather than the event level. The subject's attempts to produce a key-press sequence too rapidly apparently interfere with sequence rather than event production. Events are already produced at a fast enough rate, and an increase in the speed of sequence production has little effect on the ability to control the component events.

Changing Perception Can Change Behavior: Going Up A Level

The relationship between perception and behavior can be seen when a person learns to perform a task by controlling a new perceptual variable. An example of this can be seen in simple pursuit-tracking tasks. In the typical tracking task, the target moves randomly. When, however, a segment of target movement is repeated regularly, the subject's tracking performance improves markedly with respect to that segment (Pew, 1966). According to the hierarchical-control model, control is improved because the repeated segment of target movement can be perceived as a predictable event. With the random target, the subject must wait to determine target position at each instant in order to keep the cursor on target. With the repeated target, the subject controls at a higher level, keeping a cursor-movement event matching a target-movement event. The fact that the subject is now controlling a higher-level perception (an event, rather than a configuration) is evidenced by the longer reaction time when responding to a change in target movement. When controlling the target-cursor configuration, the subject responds almost immediately to changes in target position. When controlling target-cursor movement events, it takes nearly 1/2 second to respond to a change in target-movement pattern.

An experiment by Robertson & Clines (1985) also shows improved performance resulting from changed perception. Subjects in the Robertson and Clines study performed a learning task where the solution to a computerized game could be perceived at several different levels. Subjects who were able to solve the game showed three distinct plateaus in their performance. The level of performance, as indicated by reaction-time measurements, improved at each succeeding plateau. Because the same outputs (key presses) were produced at each level of performance, each performance plateau was taken as evidence that the subject was controlling a different perceptual variable.

Behavior/Perception Correlations

Few psychologists would be surprised by the main contention of this paper: that there is an intimate relationship between perception and behavior. However, most models of behavior assume that the nature of this relationship is causal: that behavior is guided by perception. This causal model provides no reason to expect a relationship between the *structure* of perception and behavior. For example, the causal model provides no reason to expect a relationship between the ability to identify a sequence of sounds (perception) and the ability to produce a sequence of actions (behavior). This does not mean that the model rules out such relationships; it just does not demand them.

The control model integrates perception and behavior. Behavior is no longer an output, but instead a perceptual input created by the combined effects of the actor and the environment. Behavior is perception in action. From this point of view, behavioral skills are perceptual skills. Thus, it is not surprising to find some indication of a correlation between behavioral and perceptual ability. For example, Keele and his colleagues (Keele, Pokorny, Corcos, & Ivry, 1985) have found that the ability to produce regular time intervals between actions is correlated with the ability to perceive these intervals. These correlations are fairly low by control-theory standards, but they are expected if the production of regular time intervals involves control of the perception of these intervals.

Conclusion

This report has presented evidence that human behavior involves control of a hierarchy of perceptual variables. There is evidence that the behavior of non-human agents, such as chimpanzees, also involves the control of a similar hierarchy of perceptions (Plooij & van de Rijt-Plooij, 1990). A model of hierarchical control shows how studies of perception and behavior provide evidence about the nature of control from two different perspectives. Perceptual studies provide information about the ability to perceive potentially controllable consequences of actions. Behavioral studies provide information about the ability to produce desired consequences. The factors that influence the ability to perceive the consequences of action should also influence the ability to produce them. In both cases, we learn something about how agents control their own perceptions.

The hierarchical-control model implies that limitations on the ability to produce behavior reflect limitations on the ability to perceive intended results. The speed at which a person can produce an errorless sequence of events, for example, is limited by the speed at which the order of these events can be perceived. But not all skill limitations are perceptual limitations. Controlled (perceived) results are produced, in part, by the outputs of the behaving agent. The ability to produce certain outputs can limit the ability to control certain perceptions. For example, it is impossible to perceive oneself lifting a 300-pound barbell until the muscles have been developed to the point that they are able to generate the output forces necessary to control this perception.

Perception and behavior are typically treated as two completely different types of phenomena. Perception is a sensory phenomenon; behavior is a physical phenomenon. But the concept of control as the behavior of perception suggests that this separation is artificial. Perception and behavior are the same phenomenon seen from two different perspectives. In order to understand how this phenomenon works, it will be necessary to understand how agents perceive (perception) and how they act to affect their perceptions (behavior). Studies of perception and behavior should become an integral part of the study of a single phenomenon: control.

Availability of Software

A HyperCard version of the number-rate-adjustment program can be obtained from the author. Send a formatted 3.5-inch double-density or high-density diskette in a reusable mailer with return postage.

References

Albus, J. (1981). *Brains, behavior, and robotics.* Petersborough, NH: Byte Books.

Arbib, M. (1972). *The metaphorical brain.* New York: Wiley.

Bryan, W. L. & Harter, N. (1899). Studies on the telegraphic language: The acquisition of a hierarchy of habits. *Psychological Review, 6, 345*-375.

Davis, W. J. (1976). Organizational concepts in the central motor network of invertebrates. In R. M. Herman, S. Grillner, P. S. G. Stein, & D. Stuart (Eds.), *Neural control of locomotion (p. 265).* New York: Plenum.

Greeno, J. G., & Simon, H. A. (1974). Processes for sequence production. *Psychological Review,* 81,187-197.

Herrigel, E. (1971). *Zen in the art of archery.* New York: Vintage.

Hubel, D. H. & Wiesel, T. N. (1979). Brain mechanisms of vision. In J. M. Wolfe (Ed.), *The mind's eye: Readings from Scientific American* (pp. 40-52). New York: Freeman.

Keele, S. W., Pokorny, R. A., Corcos, D. M., & Ivry, R I. (1985). Do perception and production share common timing mechanisms: A correlational analysis. *Acta Psychologica,* 60,173-191.

Kline, R. (1983). Comment on Rosenbaum, et al. Hierarchical control of rapid movement sequence. *Journal of Experimental Psychology.. Human Perception and Performance,* 9, 834-36.

Kolers, P. (1972). The illusion of movement. In R. Held & W. Richards (Eds.), *Perception: Mechanisms and models (pp.* 316-323). San Francisco: Freeman.

Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior: The Hixon Symposium* (pp.112-136). New York: Wiley.

Marken, R S. (1986). Perceptual organization of behavior: A hierarchical control model of coordinated action. *Journal of Experimental Psychology: Human Perception and Performance,* 12, 267-76.

Marken, R. S. (1988). The nature of behavior: Control as fact and theory. *Behavioral Science,* 33, 196-206.

Marken, R. S. (1989). Behavior in the first degree. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 299-314). Amsterdam: North-Holland.

Marken, R. S. (1990). Spreadsheet analysis of a hierarchical control system model of behavior. *Behavior Research Methods, Instruments & Computers,* 22, 349-359.

Marken, R. S., & Powers, W. T. (1989). Levels of intention in behavior. In W. A. Hershberger (Ed.), *Volitional action: Conation and control* (pp. 409-430). Amsterdam: North-Holland.

Martin, J. G. (1972). Rhythmic (hierarchical) versus serial structure in speech and other behavior. *Psychological Review,* 79, 487-509.

Miller, G. A., & Heise, G. A. (1950). The trill threshold. *Journal of the Acoustical Society of America,* 22,637-638.

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology,* 9, 441-474.

Pew, R. W. (1966). Acquisition of hierarchical control over the temporal organization of a skill. *Journal of Experimental Psychology,* 71, 764-771.

Plooij, F. X., & van de Rift-Plooij, H. H. C. (1990). Developmental transitions as successive reorganizations of a control hierarchy. *American Behavioral Scientist,* 34, 67-80.

Povel, D-J. (1981). Internal representation of simple temporal patterns.

Journal of Experimental Psychology: Human Perception and Performance, 7, 3-18.

Powers, W. T. (1973). *Behavior. The control of perception.* Chicago: Aldine.

Powers, W. T. (1979, August). The nature of robots: Part 4: A closer look at human behavior. *Byte, pp.* 309-323.

Powers, W. T. (1989). *Living control systems.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1990). A hierarchy of control. In R. J. Robertson & W. T. Powers (Eds.), *Introduction to modern psychology: The control-theory view* (pp. 59-62). Gravel Switch, KY: Control Systems Group.

Powers, W. T., Clark, R. K., & McFarland, R. L. (1960). A general feedback theory of human behavior: Part II. *Perceptual and Motor Skills, 11,* 309-323.

Robertson, R J., & Glines, L.A. (1985). The phantom plateau returns. *Perceptual and Motor Skills,* 61, 55-64.

Rosenbaum, D. A. (1983). Hierarchical control of rapid movement sequences. *Journal of Experimental Psychology: Human Perception and Performance,* 9, 86-102.

Rosenbaum, D. A. (1987). Hierarchical organization of motor programs. In S. P. Wise (Ed.), *Higher brain functions: Recent explorations of the brain's emergent properties.* New York: Wiley.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74(11).*

Sternberg, S., Knoll, R. L., Monsell, S., & Wright, C. E. (1989). Motor programs and hierarchical organization in the control of rapid speech. In M. Jeannerod (Ed.), *Attention and performance XIII: Motor representation and control (pp.* 3-55). New Jersey: Erlbaum.

Stemberg, S., Knoll, R. L., & Turlock, D. L. (1990). Hierarchical control in the execution of action sequences. *Phonetica,* 45,175-197.

Warren, R M. (1974). Auditory temporal discrimination by trained listeners. *Cognitive Psychology, 6,* 237-256.

Warren, R M., Obusek, C. J., Farmer, R. M., & Warren, R P. (1969). Auditory sequence: Confusion of patterns other than speech or music. *Science, 164,* 586-587.

Mimicry, Repetition, and Perceptual Control

W. Thomas Bourbon
(Research Division, Department of Neurosurgery, University of Texas Medical School - Houston, 6431 Fannin, Suite 7.148, Houston, TX 77030)

Abstract

In their attempts to explain, predict, and control human behavior, behavioral scientists typically overlook controlling done by themselves and by the people they study. The literature on perceptual control theory (PCT) describes several reasons for that omission, and in this paper I show another. when they mimic events in the environment, or when they repeat actions that they imagine or remember, people "act like" the kinds of lineal causal systems portrayed in most behaviorist and neuro-cognitive theories of behavior. A PCT model can emulate the behavior both of the person who acts like the lineal causal models and of the lineal models themselves. The results described in this paper show that the lineal causal models used in the behavioral sciences produce behavior that is a special limiting case of the behavior exhibited by the control-system model in PCT.

Mimicry and Repetition are Limiting Cases of Perceptual Control

People are living control systems who control many of their own perceptions. This paper is about two circumstances that have led scientists to think people are not living control systems: (1) when people try to mimic the actions of variables in the environment, and (2) when people try to repeat remembered or imagined patterns of actions. In these cases, the behavior of a control system can be mistaken for that of a lineal causal system whose actions are caused by antecedent events. To show that observers can mistake people for cause-effect systems, I use a demonstration that builds on work described in a previous paper, "Models and Their Worlds" (Bourbon & Powers, 1993), hereafter referred to as "Worlds." In the present demonstration, a person does variations on a simple pursuit-tracking task. In the process, the person unintentionally imitates the performance of two popular cause-effect models of people. Then I show a PCT model that duplicates the person's performance, as well as that of each lineal causal model.

The *Experimental Setting: Pursuit Tracking*

Figure 1A shows the experimental setting from "Worlds." A person uses a control handle to affect a cursor (a short horizontal mark on a computer screen) while two target marks unaffected by the handle move in unison up and down on the screen. Figure 1B shows the environmental variables that affect the cursor and target. For each of 1800 moments sampled during a one-minute run and modeled during a simulation, the following program statement determines the position of the cursor:

c: = h + d,

where c is cursor position, h is handle position, t is the momentary value of the target function generated by the computer, and d is the momentary value of a computer-generated disturbance (zero for some runs).

For the first part of the demonstration, the task was the same as the one described in "Worlds":

> The person's task in all phases of the experiment is to keep the cursor exactly between the target lines. (There is nothing special about that relationship between cursor and target; the person could easily select any other.) This task is known as "tracking" (Bourbon & Powers, 1993, p. 55).

Seen Cursor Position Minus Seen Target Position Equals Zero

Figure 1C shows the results when the person kept the cursor aligned with a moving target. The target moved up and down at a constant velocity, and no disturbance affected the cursor (d = 0). The person moved the handle in a pattern that necessarily, but unintentionally, resembled the pattern for the target.

Perceptual control theorists often use the PCT model to reproduce and predict results like these. Correlations between predicted and actual handle positions often exceed .995, even when the predictions precede the person's data by one year (Bourbon, Copeland, Dyer, Harman, & Mosley, 1990) or five years (Bourbon, 1993a). In those studies, people kept the cursor aligned with the target, but a person could easily select any other relationship to control, as I show next.

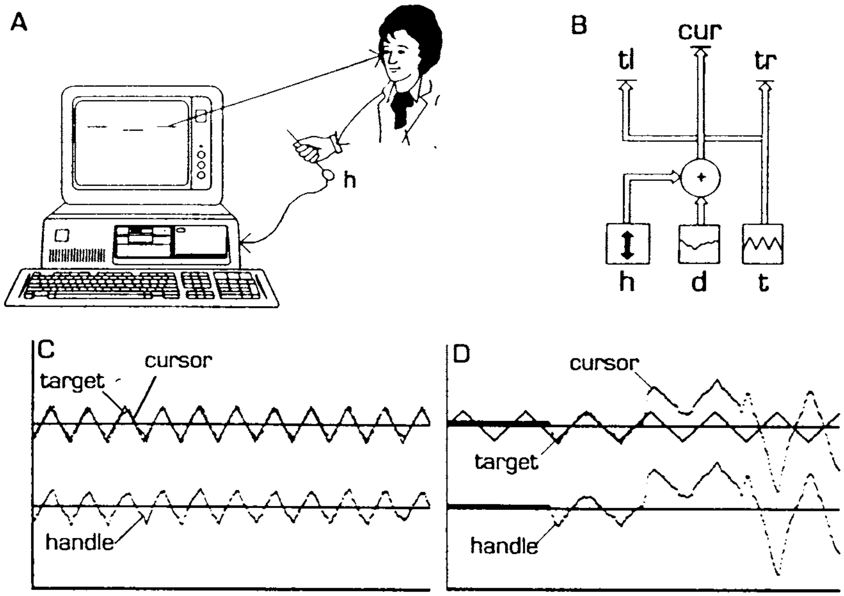

Figure 1. A. The experimental setup, in which a person uses a control handle to keep a cursor in a desired relationship with a moving target on a computer screen. B. Environmental connections among handle position (h), the disturbance function (d), target function (t), target marks (tl and tr), and cursor (cur). C. Results when a person did the tracking task and d was zero. D. Results when a person did the tracking task and, during successive 15-second periods, (a) did not move the handle, (b) kept the cursor aligned with the target, (c) kept the cursor one inch above the target, and (d) moved the cursor to twice the inverse of the target position. (In the plots, "up" represents the handle moving away from the person, and the cursor and target moving upward on the screen. The horizontal axis of each plot represents time, from 0 to 60 seconds.)

Seen Cursor Position Minus Seen Target Position Equals Various Values

Figure 1D shows the results with the person in the same setting as before, but with the target moving at a slower velocity. During successive 15-second intervals, the person (a) did not move the handle, (b) kept the cursor even with the target, (c) kept the cursor an inch above the target, and (d) moved the cursor to positions twice as great as the inverse of the target. The person did not need practice to produce these results. Bourbon (1993b) showed that a simple PCT model can duplicate results like these. When a person and a PCT model adopt and create different intended perceptions, they disprove the common misconception that control systems cannot change their "goals" or intended results.

Predictions by the Three Models from "Worlds"

In "Worlds," after a run with the conditions shown in Figure 1C, we tested two popular lineal causal models and the model from perceptual control theory. We compared the models' predictions of what the person would do when the experimental conditions changed. We described the models in detail in "Worlds"; I summarize them in this paper's Appendix.

Running the Models

In 'Worlds," we described the procedures for running (simulating) each model. We used the person's data from an initial experiment to estimate the parameters for each model, then ran the models under altered conditions. The present demonstration followed the same procedure: I used data from Figure 1C to estimate the parameters of the models, then ran them in simulation. The top row of Figure 2 shows the results of the simulations, which are the same as those in Phase 3 of "Worlds" (p. 65). Each result is a quantitative prediction by a model (described in the Appendix) of what would happen if the person functioned like that particular model.

The PCT model. The PCT model tests the idea that when the person produced the results in Figure 1C, he compared his momentary perceptions against what he intended to perceive. When there was a mismatch between present and intended perceptions, his actions changed to create and maintain a match. If the person acted that way during the first task, then he could probably keep the cursor aligned with the target, even when it followed a new and variable pattern and a random disturbance affected the cursor. His handle positions, which would vary as necessary to oppose the random disturbance, would be unintended side-effects of control and would no longer duplicate the positions of the target or the cursor they control.

In the present simulation of the PCT model (Figure 2A), the reference signal specified the perceptual signal, and any discrepancy between the signals drove the handle to positions that canceled the effects of the disturbance to the cursor. The cursor remained aligned with the target, as was intended, and the position of the handle was an unintended side-effect of control.

The S-R *model.* A stimulus-driven (stimulus-response, S-R) model tests the idea that for the results in Figure 1C, the position of the target

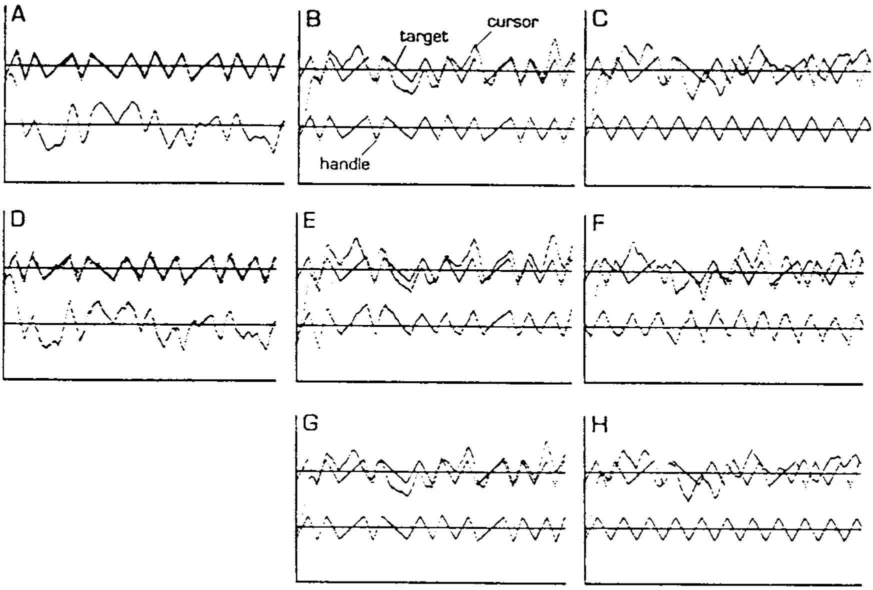

Figure 2. Top row, predictions by (A) a PCT model (cursor position - target position = zero), (B) an S-R model (handle position = target position), and (C) a plan-driven model (handle position = planned handle position). (These models are described in the Appendix.) Middle row, data when the person controlled (D) to keep the seen cursor even with the seen target, (E) to keep felt (unseen) handle position = seen target position, and (F) to keep felt (unseen) handle position = planned handle position. Bottom row, results when the PCT model impersonated the other models, with reference signals for (G) handle position = target position, and (H) handle position = planned handle position. In each run or simulation, the target path was the same, and the same random disturbance affected the cursor. (In the plots, "up" represents the handle moving away from the person, and the cursor and target moving up on the screen. The horizontal axis of each plot represents time, from 0 to 60 seconds.)

reflexively determined the position of the person's control handle. If the person acted that way during the first run, then his handle could still follow the target when it traced a new and variable pattern and a random disturbance affected the cursor. In that case, the position of the cursor would become an unintended side-effect of control. In the present simulation of the S-R model (Figure 2B), the target determined the position of the handle, and their momentary positions were nearly identical. The cursor "wandered" away from the target and its position was an unintended side-effect.

The plan-driven model. The plan-driven neuro-cognitive model tests the idea that, for the results in Figure 1C, the person's memory of momentary handle positions from earlier practice sessions determined the position of his control handle. If the person acted that way during the first run, then his handle positions should duplicate the ones in Figure 1C, even when the target followed a new and variable pattern, and a random disturbance affected the cursor. Neither the handle nor the cursor would duplicate the pattern of movement traced by the target. The position of the cursor would be an unintended side-effect of control.

In the present simulation of the plan-driven neuro-cognitive model (Figure 2C), the plan for a pattern of target movements ("remembered" from the data in Figure 10 determined the position of the handle. The uncontrolled cursor wandered independently of the target. Its position was a side-effect of control.

The Person Performs Under New Conditions

In the simulations I just described, the three models predicted different results for the person running under the environmental conditions from Phase 3 of "Worlds" (pp. 64-67). Now I report what happened when the person repeated the tracking task three times under those conditions. A random disturbance affected the position of the cursor, and, from one excursion to the next up or down the screen, the probability was 2/3 that the velocity of the target would change to another of three possible values. During the first repetition, the person again kept the cursor aligned with the target; in the other two, he created results in which the position of the cursor became an uncontrolled side effect.

Seen Cursor Position Minus Seen Target Position Equals Zero

First, the person kept the cursor aligned with the target. Figure 2D shows the results. The patterns of positions for the target and cursor were similar ($r = .91$, $n = 1800$ data pairs). The pattern of the person's handle movements necessarily differed from that in Figure 1C. He controlled the relationship of the cursor and target, which was a consequence of actions, but did not control his actions. The relationship between his actions and the movements of the target necessarily varied to eliminate effects of the disturbance on the cursor's position. It is impossible for a person to specify and plan the required actions before a condition as variable and disturbed as this one.

Failed models? In "Worlds," we compared the person's data in this condition against the predictions by the two lineal causal models, which in the present demonstration is the same as comparing Figure 2D with Figure 2B and 2C. Obviously, the results when the person used the handle to keep the cursor aligned with the target were different from the results of the lineal causal models. Those models controlled the position of the simulated handle, but not the cursor. The lineal models, which can accurately explain the results of the undisturbed condition shown in Figure 1C, failed to predict the results in Figure 2D. The person did not act like a lineal causal system, but *can* he?

Mimicry and repetition. So far, the person has controlled the position of a cursor compared with a target, and the positions of his control handle were unintended and uncontrolled side-effects. What would happen if he did not control the position of the cursor at all, and instead controlled his actions? The position of the cursor would become an uncontrolled side-effect. Next, I show the results when the person ran under the same conditions as shown in Figure 2D but controlled his felt perceptions of hand movements. First he made them match the seen movements of the target, then he made them match a remembered pattern of felt movements. In the first case, his movements mimicked a present perception in another sensory modality. They appeared to fit the S-R model, where "stimuli" (movements of the target) cause "responses" (movements of the handle). In the second case, his movements repeated a remembered pattern. They appeared to fit the plan-driven neuro-cognitive model, where plans or commands from the mind-brain control handle movements, independent of events in the environment.

Mimicry: Felt Handle Position Equals *Seen Target Position*

People sometimes make their actions mimic those of other people. Some children who watch adults playing musical instruments use toy instruments and make exaggerated motions that they believe are the same as the adult's actions. Sometimes an inexperienced person attempts to perform without practice in a marching band or military unit by watching and duplicating the actions of others in the unit. In gatherings, sometimes individuals mimic what they see other people doing. In cases like these, people try to make their felt actions match actions or events they see in the environment.

During the present demonstration, the person made his felt, but unseen, handle movements match the movement of the target. By making his actions duplicate the movements of an environmental stimulus, he played the role of an S-R system; he functioned like a control system, making his presently perceived hand position match the presently seen position of the target. To help him play that role, the target function and the disturbance remained the same for 15 practice runs, and a piece of cardboard shielded his hand from view. He practiced nuking his felt hand movements match the movements of the target on the screen. When he decided that he was "ready," he did the run shown in Figure 2E and 11 similar runs, for a total of 12 runs.

The person's controlled handle movements generally resembled the pattern of target movements. The mean correlation for twelve sets of predicted and actual handle positions was .824 (S.D. = .089, range = .981 to .615, n = 1800 data pairs per set). By accepted standards in behavioral science, that mean correlation is extremely high, but agreement between predicted and actual handle positions is even higher when a person keeps a cursor aligned with a target. In a study with 104 sets of 1800 predicted and actual handle positions, Bourbon et al. (1990) reported a mean correlation of .996 (S.D. = .002). It is easier for people to make one seen environmental variable track another than to make their own felt actions "track" a seen variable.

There were obvious differences between movements of the target and handle. For example, the target always moved at one of three uniform velocities, but the person's handle velocities were not uniform. Also, before reversing direction, the target always moved the same distance above or below the center of the screen, but the person reversed handle movements at varying distances from the center of their range. He did not perfectly duplicate the performance of a pure S-R system. Even after 12 practice sessions, it was not easy for him to judge and control either the velocity of handle movements or the distances he moved the handle before reversing its direction.

During the present trials, the cursor was affected by a random disturbance and by the handle. It "wandered" around the position of the target. Cursor position was an accidental side-effect when the person controlled the position of the handle.

Repetition: Felt Handle Position Equals Remembered Handle Position

Sometimes people make their patterns of actions repeat a remembered or imagined pattern. Many self-improvement and rehabilitation programs urge clients to imagine themselves doing a desired action "perfectly;" then to do the action as imagined. When people attempt to move through a darkened familiar environment, they sometimes try to duplicate movements they remember from when they could see their surroundings. In a group that uses a device like a baton or banner to do synchronized routines, some people who drop the device try to continue making movements remembered from performances when they held it. In cases like these, people try to make the actions they feel match patterns they remember or imagine.

In the present demonstration, the person did the condition shown in

Figure 2, except that he made the pattern of his felt-but-unseen handle movements match the pattern he remembered from the condition shown in Figure 1C. By making his actions duplicate the earlier pattern, he imitated the performance of a neuro-cognitive plan-driven system. To help him act that role, he ran 22 replications of the undisturbed task shown in Figure 1C and kept the cursor aligned with the target. A piece of cardboard screened his hand from view, and he paid close attention to the tactile and kinesthetic sensations that accompanied successful tracking. He intended to repeat the practiced movements from memory when the screen was blanked during the next task When he was ready, the program started. The initial positions of all variables were displayed on the screen, then the screen went blank and he completed the run shown in Figure 2F and 15 additional runs, for a total of 16 runs.

Qualitatively, the pattern of the person's controlled handle movements resembled the one from Figure 1C. Quantitatively, the match between modeled and actual patterns of handle movement was atrocious. The mean correlation for sixteen sets of 1800 predicted and actual handle positions was -0.003 (S.D. = .118, range = .390 to -.223). It was *much* harder for the person to create a precise replica of a highly practiced regular pattern of handle movements than to make either a cursor or his hand movements match a seen target. This result has serious implications for all neuro-cognitive plan-driven models of behavior, but especially for those where people claim that the elimination of sensory "feedback" does not affect planned actions. In the present case, simply concealing the person's hand behind a piece of cardboard eliminated precise repetition of the desired pattern.

There were obvious differences between handle movements during the undisturbed run and this one. In the undisturbed run, where handle position was an accidental side-effect of control, the velocity of the person's handle movements necessarily approximated the uniform velocity of the target; in the plan-driven run, where he controlled the handle's positions, handle velocities were more erratic. Also, during the plan-driven run, he reversed the direction of the handle at varying distances from the center of its range; during the undisturbed run, when the position of the handle was an unintended side-effect, the reversals were more uniform. Even after 22 practice sessions, it was not easy for him to judge and control either the velocity of handle movements or the distances he moved the handle before he reversed its direction

The person labored under other serious burdens that confront every Plan-driven system. Such systems are extraordinarily sensitive to the slightest errors in the timing of actions and to the smallest deviations from the required values of any important variables. A deviation at any time during the running of such a system can quickly lead to actions and consequences that are the *reverse* of what they should be. We discussed this extreme sensitivity to small errors in "Worlds" (p. 59), but we did not show quantitative examples of its consequences. Plan-driven models cannot serve as general models of human behavior.

Comparing the Models and the Person

The person's handle positions (Figure 2E and 2F) were more variable than those of the corresponding lineal causal models (Figure 2B and 2C, respectively), in large part due to his not maintaining uniform velocities for the handle. Also, the person moved the handle through a pattern that was not centered in the range of movement, but the models centered their simulated handles. Finally, the plan-driven model perfectly "remembered" the pattern of target movement from the first run, but the person obviously did not; he reversed the direction of handle movement at the wrong times, compared to the ideal remembered pat tern. When it comes to controlling one's own actions, what happens is not always what the person remembers and intends.

The PCT Model Emulates the Person and the Causal Models

Pure causal systems, like the lineal models I explained earlier, cannot produce unvarying results in a variable environment. In 'Worlds," we described a rationale for making causal models succeed in a variable world:

> To modify cognitive or SR models so that, like living systems, they might thrive amidst change, we must... give each model an internal standard and a process for comparing present perceptions against that standard. But then the models would all be control systems, each controlling its input (Bourbon & Powers, 1993, p. 70).

We cannot modify either a pure S-R model or a pure plan-driven model so that it emulates the PCT model, yet simultaneously preserve its core structure. On the other hand, we can easily modify a PCT model so that it emulates either lineal causal model: All we have to do is change p*, the reference signal for the PCT model.

The PCT Model Emulates the S-R Model

To emulate the S-R model, where the position of the target determines the position of the model's handle, the PCT model makes its perceived handle position match the perceived position of the target. The reference signal, p*, becomes h - t = zero, where h and t are positions of the handle and target. Any perceived discrepancy (error signal) between h and t changes the position of h, according to the following program steps:

p: = h - t
error: = p* - p
h: = h + k • error • dt

With no other change, the PCT model will "impersonate" the S-R model (and the person, when he made his felt handle movements match seen target movements).

Figure 2G shows the results when the "modified" PCT model ran in simulation. It reproduced the results of the pure S-R model (Figure 2B): the disturbed and uncontrolled cursor no longer tracked the target, but handle movements, which were now controlled, accurately tracked target movements. This PCT model also reproduced general features of the person's attempt at impersonating a stimulus-driven system, shown in Figure 2E. However, the agreement between the PCT model and the person would be just as poor as that between the S-R model and the person.

The PCT Model Emulates the Plan-Driven Mode!

The PCT model can emulate the plan-driven model, where the computed or remembered pattern of previous target positions determines the position of the model's handle. In that role, the PCT model specifies that the perceived handle position at any moment matches the computed position. The reference signal, p*, for the PCT model becomes h - H = zero, where h is the present position of the handle, and H is the momentary computed or remembered ideal position. A perceived discrepancy between those positions produces movements of the handle, according to the following steps in the computer program:

p: = h - H
error: = p* - p
h: = h + k • error • dt

With no other change, the PCT model will emulate the plan-driven model (and the person, when he made his felt handle movements match a remembered pattern of handle movements).

Figure 2H shows the results when the "modified" PCT model ran in simulation. It accurately duplicated the results of the pure plan-driven model (Figure 2C). The PCT model also reproduced qualitative features of the person's attempt at impersonating a neuro-cognitive plan-driven system, shown in Figure 2F. However, the agreement between the PCT model and the person would be just as poor as that between the neuro-cognitive plan-driven model and the person.

Discussion

A person can act like a system where environmental stimuli control its actions, and like one where internal plans and commands control its actions; a PCT model can achieve the same results as the person, but neither a pure stimulus-controlled model nor a pure plan-driven model can duplicate all of the appearances of a person and of the other models. To make either cause-effect model to do that, we would need to radically change its core structure and convert it into a perceptual control system. However, for people or PCT models to act like lineal causal systems, their core structures do not change. All that changes for the person is the intended perception; for the PCT model, only the reference signal changes.

Generality of the PCT Model

In the present demonstration, a person used the experimental arrangement shown in Figure 1A to achieve several different controlled results. In the second stage of the demonstration, three different models of behavior each predicted one of the person's results: the PCT model kept its cursor aligned with a target, the S-R model made its handle movements match target movements, and the plan-driven model made its handle movements match a remembered plan. The success of all three models during that stage does not mean that we need a different model to explain the person's performance in each condition. To the contrary, in the final stage of the demonstration, a PCT model with a simple change in its reference signal duplicated all of the results of the person and the two lineal causal models. Perceptual control theory provides a *general* model of control behavior, while each of the lineal models applies only to a limiting case.

There is no defense for using either lineal causal model as a general model of behavior, but many behavioral scientists do. The settings where scientists believe the environment controls a person's behavior are diverse. They range from behavioral conditioning laboratories, where scientists say environmental stimuli control a person's actions, to social gatherings, where they say people "lose control" of their behavior, with control passing to presumed forces such as a "virus-like emotional contagion" or a "group mind." Instead of proving the legitimacy of a stimulus-response model, those are instances when, for whatever reasons, people intend to perceive their actions matching perceptions of a selected feature of the environment. Other consequences of a person's actions would, like the position of the cursor in the present demonstration, "go out of control." Events like these often catch the attention of observers, whether they are behavioral scientists or the local constabulary, but those observers are wrong if they assume that the person has "lost control" to "powerful" forces in the environment.

There are also many settings where scientists believe that a plan (command, trait, neural signal, gene, force) from the mind-brain controls a person's behavior. They range from concert halls, where many scientists say that some performers' actions occur too regularly and rapidly for the environment to affect them, to neurophysiological clinics, where they say that people with damaged spinal sensory nerves provide evidence that motor plans determine the course of behavior. Instead of proving the legitimacy of a plan-driven model, these are instances when, for whatever reasons, people intend to perceive their actions matching remembered or imagined patterns of movements. When they do, other consequences of their actions will, like the position of the cursor in the present demonstration, "go out of control."

How Could behavioral Scientists Overlook the Fact of Control?

I have shown that, depending on which perceptions a person controls, an observer can mistake the person for a stimulus-controlled system or a plan-driven system. That is one reason behavioral scientists might have overlooked the phenomenon of control. There are other reasons, and perceptual control theorists have described some of them.

For one thing, when scientific psychology began in the 1800s, psychologists followed a tradition several centuries old. They assumed that the lineal models of cause and effect explaining the actions of inanimate objects also explain human behavior. But as William Powers has written, the "orderly march of cause and effect from stimulus object to sensory receptor, and from muscle tension to the eventual behavioral result, does not exist" (Powers, 1973, p. 4). Powers described a fact that sometimes makes it difficult for informed observers to see the phenomenon of control and virtually guarantees that uninformed ones will not:

> In general an observer will *not*, therefore, be able to see what a control system is controlling. Rather, he will see an environment composed of various levels of perceptual objects reflecting his own perceptual organization and his own vantage point. He will see events taking place, including those he causes, and he will see the behaving organism acting to cause changes in the envi-

> ronment and the organism's relationship to the environment. The organism's activities will cause many changes the observer can notice, but what is controlled will only occasionally prove to be identical with any of those effects. Instead, it will normally be some function of the effects, and the observer's task is to discover the nature of that function (1973, p. 233, emphases in the original).

Powers has written much more about those ideas (see Powers, 1989, 1992). So have other perceptual control theorists. One of them, Wayne Hershberger (1987a, 1987b, 1988, 1989, 1990), has discussed the idea that when an organism controls its perceptions, observers often notice overt actions that seem either elicited by antecedent environmental stimuli, or emitted from within the organism Psychologists have treated elicited and emitted behaviors as distinct from one another and governed by different "laws"; they sometimes call elicited actions "involuntary" and emitted actions "voluntary." Hershberger emphasizes the fact that organisms voluntarily control many of their perceptions of environmental variables by using involuntary actions to eliminate effects of environmental disturbances acting on those variables. The illusory exclusivity of the two "classes" of behavior makes it difficult for many observers to notice that the organism is a controller.

In a series of ingenious experiments, Richard Marken (1982, 1989, 1992) has illuminated another point made by Powers: when an organism voluntarily controls its perceptions, its actions simultaneously produce many unintended consequences. It is not always obvious which of the many variables an organism affects are "under control." Marken has shown the procedures that an observer must follow to distinguish between intended and unintended consequences of behavior—between controlled and uncontrolled states of the environment.

Marken (1993) also has shown several circumstances where an observer can mistakenly think a perceptual control system is a reflexive stimulus-response system, or a reinforcement-controlled system, or a cognitive system. Mistakes like these are behind many lineal causal models in behavioral science, and they guarantee that scientists will "miss" the fact that organisms control many of their own perceptions. Marken suggests that theorists who advocate any of the three mutually exclusive lineal causal models are similar to the three legendary blind men who encountered an elephant: each observes part of the phenomenon of control, consequently, their various interpretations of *the phe*nomenon are incomplete and incorrect, but understandably so.

Conclusion

In the present demonstrations, a person and a PCT model emulated, or "acted like," lineal causal models used in nearly all behavioral theories. Similarly, in laboratories and clinics, people emulate nearly any kind of system a scientist thinks they should be. For more than a century, the clinical practices, research methods, and theoretical preferences of behavioral scientists have guaranteed they would not discover this obvious fact: a person is one kind of "thing" that an observer can mistake for any of the many kinds behavioral scientists have imagined. Every person controls perceptions; perceptual control theory explains and predicts the control of perception, even when a person impersonates a lineal causal system.

Appendix

The following behavioral models are from the paper "Models and Their Worlds" (Bourbon & Powers, 1993).

The S-R Model

From the person's data during the run in Figure 1, we calculated the slope (m) and offset (intercept, b) of the regression of the handle on the target. Target position is t and handle position is h. The S-R model for the person consists of

h: = mt + b

and

c: = h + d.

Target position, an independent variable, determines handle position, as a dependent variable. This model represents pure environmental control of behavioral actions.

The Plan-Driven Model

The plan-driven cognitive model "remembers" the average pattern of target movements during the run shown in Figure 1, then "computes" handle movements that perfectly match those target movements. The resulting model consists of

h: = H

and

$c: = h + d.$

In this model, a computed representation (H) of the pattern of previous target movements (t) causes the handle to move in a pattern identical to that of the computed representation.

The PCT Model

The computational steps for the PCT model are

$p: = c - t,$

$error: = p^* - p,$

$h: = h + k \cdot error \cdot dt,$

and $c: = h + d,$

where p is the perceptual signal, and p* is the reference signal or intended value of p. In "Worlds" (Bourbon & Powers, 1993, p. 61), we explained k, the integration factor that resents the velocity of handle movements when there is error, and dt, the sampling interval (here, 1 / 30 second). The reference signal specifies the perceptual signal; if they do not match, the resulting error causes handle movement.

Acknowledgements

I thank Richard Marken and Greg Williams for their critical reviews of versions of this paper, and William Powers for his encouragement when I designed the demonstrations.

References

Bourbon, W. T. (1993a). On the accuracy and reliability of predictions by control-system theory: Data gathered five years after the predictions. Experiment conducted at the annual meeting of the Control Systems Group, Durango, Colorado, 27 July-1 August 1993.

Bourbon, W. T. (1993b). A modest proposal to all behavioral scientists. Paper presented at the annual meeting of the Control Systems Group, Durango, Colorado, 27 July-1 August 1993.

Bourbon, W. T., Copeland, K. C., Dyer, V. R, Harman, W. K., & Mosley, B. L. (1990). On the accuracy and reliability of predictions by control-system theory. *Perceptual and Motor Skills,* 71,1331-1338.

Bourbon, W. T., & Powers, W. T. (1993). Models and their worlds. *Closed Loop, 3,* 47-72.

Hershberger, W. A. (1987x). Some overt behaviors are neither elicited nor emitted. *American Psychologist, 42,* 605-606.

Hershberger, W. A. (1987b). Of course there can be an empirical science of volitional action. *American Psychologist,* 42, 1032-1033.

Hershberger, W. A. (1988). Psychology as a conative science. *American Psychologist, 43,* 823-824.

Hershberger, W. A. (1989). The synergy of voluntary and involuntary action. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 3-20). Amsterdam: North-Holland.

Hershberger, W. A. (1990). Control theory and learning theory. *American Behavioral Scientist, 34,* 55-66.

Marken, R. (1982). Intentional and accidental behavior. *Psychological Reports, 50,* 647-&50.

Marken, R. S. (1989). Behavior in the first degree. In W. A. Hershberger (Ed.), *Volitional action: Conation and control (pp.* 299-314). Amsterdam: North-Holland.

Marken, R. S. (1992). *Mind readings: Experimental studies of purpose.* Gravel Switch, KY: Control Systems Group.

Marken, R. S. (1993). The blind men and the elephant: Three perspectives on the phenomenon of control. *Closed Loop, 3,* 37-46.

Powers, W. T. (1973). *Behavior: The control of perception.* Chicago: Aldine.

Powers, W. T. (1989). *Living control systems.* Gravel Switch, KY: Control Systems Group.

Powers, W. T. (1989). *Living control syste*ms II. Gravel Switch, KY: Control Systems Group.

Page 72 intentionally left blank

The Control Systems Group is a membership organization which supports the understanding of cybernetic control systems in organisms and their environments: *living control systems.* Academicians, clinicians, and other professionals in several disciplines, including biology, psychology, social work, economics, education, engineering, and philosophy, are members of the Group. Annual meetings have been held since 1985. The CSG Business Office is located at 73 Ridge Pl., CR 510, Durango, CO 81301; phone (303)247-7986.

The CSG logo shows the generic structure of cybernetic control systems. A Comparator (C) computes the difference between a reference signal (represented by the arrow coming from above) and the output signal from Sensory (S) computation. The resulting difference signal is the input to the Gain generator (G). Disturbances (represented by the black box) alter the Gain generator output on the way to Sensory computation, where the negative-feedback loop is closed.

Back cover