ODYSSEY

University of Idaho Library ILL

ILLiad TN: 432315

Borrower: ANTCH

Lending String:
*NTD,GND,NBU,RVA,KCC,UMR,THM,KEN,VLW,
AUM,GMU,SHH,VOD,IOH,TSC

Patron:

Journal Title: The Psychobiology of consciousness /

Volume: **Issue:**
Month/Year: **Pages:** 217-242

Article Author: Richard J Davidson; Julian M Davidson William T. Powers

Article Title: William T. Powers: A Systems Approach to Consciousness

Imprint: New York : Plenum Press, 1980

ILL Number: 166894702

Call #: BF311.P785

Location:

ODYSSEY ENABLED

Charge
Maxcost: 0.00IFM

Shipping Address:
Library - We Deliver
Antioch University
40 Avon St
Keene, New Hampshire 03431-3516
United States

Fax: 603-357-7345
Ariel:
Email: wedeliver@antioch.edu

10

A Systems Approach to Consciousness

WILLIAM T. POWERS

1. INTRODUCTION

The nature of consciousness will still be a mystery by the time this essay ends. We are concerned here with an approach, not an arrival. To some readers it may seem, for the first three quarters of the essay, that the approach begins unnecessarily far out in left field, but I hope that by the end, it will be seen that I have backed off so far in order to come at the subject from a new direction. Bear with me. The systems approach sometimes requires developing a lot of detailed groundwork before one can hazard any generalizations.

In applying the systems approach to anything as complex as a higher organism, the chief problem facing the theorist is that of finding a manageable level of generality—analyzing the whole system into units that are neither so detailed as to overwhelm one's limited ability to comprehend large assemblies of elements, nor so general that nothing surprising can possibly emerge from the final synthesis. Much of the controversy over consciousness has arisen, I believe, because of the gulf that lies between these two extremes. A neurologist deals with billions of tiny elements—neurons—and interactions that hold between only a few of them at a time. The humanistic psychologist deals with the brain as a single lump, and with interactions among global properties of that brain in a world of other brains. It is not surprising that there is a lack of common ground between these approaches.

2. SOME GENERALITIES ABOUT THE SYSTEMS APPROACH

The systems approach, as I use it, grew out of the hardware world, not out of biology, but it was developed for the same reasons it is needed in biology. When technologists began building complex electronic systems, they found it impossible to understand their own creations at the individual component level and, of course, unprofitable to describe

---

WILLIAM T. POWERS • 1138 Whitfield Road, Northbrook, Illinois 60062.

them only at the user's-manual level. An intermediate level of description, the block diagram, developed along with the complexity of electromechanical systems.

A block diagram describes a model of a physical system in terms of the major functions that, when assembled, constitute the whole system. Each block is intended to stand for a subsystem, not an abstraction, but the description of each block is given in the abstract. As a simple example, if two neural signals, f_1 and f_2, converged on a single neuron, and if the outgoing frequency of impulses that resulted was proportional to the weighted sum of the incoming frequencies, this arrangement would be drawn as a block with two inputs and one output, described by the equation

$$f_{\text{out}} = k_1 f_1 + k_2 f_2$$

The describing equation is abstract in the sense that it is a mathematical idealization and omits any description of the physical means by which this relationship is brought about. Yet, the meaning of the block diagram is not abstract; behind it is always supposed to be a physical device, located in space, which accounts for the described relationship. The arrows indicating inputs and outputs are not meant simply to lead the eye along a path through the diagram; they are intended to be schematic representations of physical pathways along which flow real neural signals. Block diagrams do not show sequences of transient events, as in the flowchart of a computer program, but fixed relationships among physical subsystems, as in a representation of the computer itself.

When a collection of such blocks has been defined, one has the pieces from which working models can be constructed. The output of one block can be connected to the inputs of other blocks. *Every* way of connecting such blocks will do *something;* that is, the behavior of any assembly will follow from the input–output rules governing each individual block and from the structure of interconnections. As was discovered early in this game, the behavior of a whole asembly is by no means self-evident in the properties of the individual blocks; it is not so much a matter of the whole's being greater than the sum of its parts as of the whole's being an entity *different from* the sum of its parts. When apples and walnuts are assembled into a system, one gets not the sum of apples and walnuts but Waldorf salad.

Every way of connecting building blocks results in *some* whole-system properties. Thus, the question of how to interconnect the blocks is of a different category from the question of what the blocks are. Even after one has properly identified a set of blocks and has verified that each block corresponds with an identifiable physical unit and that the functional description of the block is correct, one must still find an organization of the blocks that will reproduce the phenomena of interest. That organization cannot be found in the detailed description of how the blocks work. It can be found only in the structure that results from interconnecting the blocks in a particular way.

One of the most important discoveries of the systems approach, in retrospect, was the finding that the properties of a whole system are nearly independent of the properties of the building blocks that comprise it. Any simplistic notions of physical determinism had to be abandoned before the first electronic digital or analogue computer was constructed. The entities that proved important to understanding a large system were the functional blocks, but not the internal construction of those blocks. There is an infinity of different ways to build a physical device that is to perform according to a given functional description.

An elementary function could be designed to sum three signals, f_1, f_2, and f_3, and to represent the sum as a single output signal, f_4. This could be accomplished by a three-way convergence on a single summing element, or it could be accomplished in two stages; first, f_1 and f_2 are added together to produce a signal f_{2a}, and then f_{2a} and f_3 are added together to produce an output signal f_4. In this case, the time lag would distinguish the two internal constructions, but it would not distinguish a third alternative in which f_1 and f_3 are the two signals added together at the first stage especially if the remaining signal, whichever it is, passes through a time-delay element too.

The laws that govern the behavior of a complex system, therefore, are not the laws that govern the individual components of that system. The laws governing the individual components contain no statements about how those components shall be interconnected, nor do they impose any limits on possible interconnections (other than setting the number of *all possible* interconnections). Furthermore, there is no possible way to analyze a given system property into a *necessary* set of component properties, because any given system property could be brought about through assembling components in an immense variety of ways. Those ways are equivalent at the system level but not at the component level.

Not only are system properties independent of the laws governing individual components of the system, but they are independent of the *kinds* of physical elements involved. Once a function such as summation has been identified, it no longer matters to a system description whether the summation is done by neurons, by transistors, by vacuum tubes, or by fluid flow through channels cut in plastic blocks. The only essential aspect of the physical components is that they be able to create the relationship called summation between some set of input quantities and an output quantity.

Once the structure of a system is given, of course, and once the actual physical components have been identified, there is complete harmony between the system description and the component description. Given the structure, the behavior of the whole is seen as consistent with the behaviors of all the components. But the structure must be given; it is not enough to state the properties of the components. Structure is not a property of any one component; it is a property only of an assembly of components, just as temperature and pressure are properties only of an assembly of molecules or atoms.

The systems approach has therefore shown us very clearly that there is a hiearchy of laws of nature. The higher-level laws are not simply sums or averages of lower-level laws; they are laws that transcend lower-level laws and that cannot be described without introducing structural rules that have no meaning at the lower level of description (Brown, 1969; Pattee, 1973). We should not be surprised to find that this is true in biology; it is certainly true in physics. For example, conservation laws are laws of structure and cannot be found in the laws governing the movements of masses. When two elastic solids collide, their movements can be analyzed into accelerations due to forces generated by physical deformations and the spring constants of the materials involved. These laws say nothing about conservation of momentum, but conservation of momentum is found to hold true for *all* collisions of masses having *any* elastic properties. The nature of the physical interactions that actually took place during a collision cannot be deduced from an observation that momentum has been conserved, although when those details are known, they prove to be entirely consistent with conservation of momentum. The systems approach to organisms, therefore, has revealed only what has been a commonplace in physics for centuries.

The properties of an assembly of components structured in a particular way are consequences of laws of structure not effective at the component level. But there is no reason to stop there. Once a structure exists, and its properties are derived, that whole structure may become a component in a higher-level structure made of many such substructures. New laws of structure, one might guess, will come to light. But here the mind tends to boggle; Why would not such "structures of structures" be expressible simply in terms of a more complex statement of the same structural laws? It is very tempting to make a distinction once and for all between components and structures, creating not a hierarchy of laws but a simple dichotomy.

In principle, there is no reason that nature should not leave us with a dichotomy. But there is really no difficulty in imagining another level of structural law, when it is realized that the "component" level of laws is really a structural level, too, relative, say, to the laws of subatomic physics or quantum mechanics. Thus, it is equally true that in principle, there is no reason that nature should stop with a dichotomy, especially since it does not seem to have done so in the world described by physicists.

The systems approach itself is not limited to any particular number of hierarchical levels; it is simply a process of analysis of something into components that are individually describable, and then synthesis into structures that interconnect those components. Sometimes, this process results in discovery of a new level of laws, and sometimes, it simply expands the scope of levels already found. There is, however, no way to predict in advance which will happen; that is up to nature, not the analyst.

The systems approach, therefore, is not just an analysis of organisms that lies halfway between holistic concepts and molecular concepts. It is a method through which any number of intermediary stages can be constructed, as many as necessary to link both ends of the spectrum, if they are linkable. It is, thus our best hope for reconciling mechanistic approaches to (or avoidances of) the subject of consciousness with the intuitive approaches that seem for now to exist in a different universe of discourse.

3. A PARTICULAR SYSTEM MODEL

For what seems a long time (considering the progress I have made), I have been focusing on a particular set of laws of structure and their implications for behavioral science: the laws that govern the organization of functions known as a *negative feedback control system*. The applicability of these laws to an understanding of organisms was noticed at the outset of cybernetics; indeed, feedback control was the core concept from which Wiener (1948) developed most of his thinking on this subject. The fact that I have not gone on to the kinds of complexity and diversity that have characterized cybernetics since those early days probably marks me as backward, but I have been convinced that the concept of feedback control is exceptionally important and that its potential as an organizing principle has not been fully appreciated in the behavioral sciences. Others have noted and used the concept of negative feedback; many have. For the most part, however, the uses of these concepts have been limited and have been strongly conditioned by older concepts of cause and effect. Even in engineering psychology, where the methods of control system analysis have been applied to studies of tracking behavior for well over 20 years, the organism itself is still usually represented in block diagrams as a simple input–output or stimulus–response device, a "transfer function" that reduces, at low frequencies, to a single constant of proportionality. In a few instances (e.g., Kelley, 1968) the analysis has gone farther than that, several hierarchical levels of control being proposed for some peripheral neuromuscular systems, but generally there has been no concerted attempt to see what the concept of feedback control means when contrasted with the fundamental assumptions under which most of behavioral science still operates.

It is that level of analysis to which I have devoted my efforts (Powers, Clark, & McFarland, 1960; Powers, 1973). The interpretations that come out of a control-system analysis of behavior seem to point toward new approaches to many subjects, particularly two subjects of interest here, awareness and volition, with which behavioral science has dealt largely by avoiding the real issues. Like Nasrudin, behavioral science has been looking for its house key under a handy streetlight, not because it was lost there but because the light is better there.

There are many ways to order a discussion of control theory, but for present purposes, it will be profitable to start with a phenomenon that has puzzled behavioral scientists as long as there has been such a subject.

3.1. *Control Theory and Stabilized Consequences*

William James (1890) is credited with pointing American psychology toward the laboratory, where, as some say, psychology lost its mind. It certainly lost the concept of purpose. But in his introduction to *The Principles of Psychology*, James insisted in several different ways that the one distinguishing feature of living organisms was their ability to keep reaching a fixed aim by employing variable means. He was saying that living organisms demonstrate purpose, not blind reactions to external forces.

John Dewey (1896/1948), writing at the same time, saw that the reflex arc could not be characterized in terms of lineal cause and effect; it is a closed circle with no beginning and no end. Both Dewey and James, had they had the tools, could easily have gone on from there to found cybernetics, for the phenomena of purpose and the kind of organization that involves a closed circle of cause and effect are key elements of control theory. The rudiments of control-system analysis were available; James Clerk Maxwell had analyzed mechanical control systems (governors) in 1868. Unfortunately, no Norbert Wiener came forth to save scientific psychology its 11-decade pursuit of a different cause–effect model, by seeing how the organizational principles that apply to governors also permit the existence of purposive systems in a universe of physical laws. Simple determinism, involving external causation of behavioral acts, won the day.

But the phenomena that have always suggested inner purpose refused to disappear. All such phenomena can be characterized in one general way: they involve stabilized consequences of variable actions. Let us pause for a few definitions.

An *action* will be taken to mean some measure of behavior that appears to be unaffected by interfering environmental circumstances: at the lowest level of analysis, a train of neural signals entering a muscle could be called an action of the nervous system. It is usually possible to find more global measures: an applied force, a velocity of a limb, and so on. The main criterion is that an action should be attributable to the organism alone, either because there is nothing present that can interfere with it or because, for whatever reason, potential interferences are, in fact, ineffective.

We need to define a comparable quantity in the environment, a quantity we will call a *disturbance*. This is a measure of some physical variable that can vary or be fixed independently of what the behaving organism does. Gravity is a ubiquitous disturbance.

A *consequence* is technically defined for this discussion as a physical variable (or a set of them) that is a joint function of action and disturbance. The posture of an animal, for example, results from the combination of muscle forces with forces due to gravity, all acting on the masses of the body through the geometric linkages of the skeleton, according to the appropriate laws of mechanics. In general, a given action and a given disturbance will be found to have many identifiable consequences; they jointly affect physical variables of many kinds.

Consequences of action and disturbance that are not joint functions of both are taken as alternative measures of action or disturbance; the term *consequence* would not be used, even though it would be permitted in ordinary discourse. I shall say *joint* consequences to emphasize this special definition.

The great majority of the phenomena that are commonly called *behavior* can be analyzed as joint consequences of an action and a disturbance (or the resultant of several disturbances, which is what *the disturbance* will generally mean). They are not, however, identifiable with *all* joint consequences. A randomly selected joint consequence would, in general, be affected just as much by unpredictable variations in the environment as by variations in an organism's actions; there would be no *repeatable* pattern to the joint consequence unless the environment and the organism precisely repeated their separate patterns. In order for a joint consequence to show enough general repeatability to be recognized as behavior in a normally variable environment, there must be a special relationship between action and disturbance: whatever the disturbance, the action must adjust so that the joint consequence repeats.

This requirement has, of course, been noticed, and there have been qualitative explanations offered for the fact that it is so often met. One is the *compensatory response* explanation: the mechanical effects of a disturbance are accompanied by sensory effects that alter the behavior in just the way that cancels the mechanical effects. But like all such analyses that do not employ the principles of control theory, it fails when put to the quantitative test even in simple cases. Consider a person holding an arm out straight ahead, the finger being held within 5 mm of a target for 30 sec, with a net load of 10 kg referred to the wrist. The muscle forces acting upward will nearly cancel the force due to gravity, 98 n. How nearly? The average acceleration of the arm must be such that the arm moves by no more than 5 mm in 30 sec; by the relationship

$$s = \tfrac{1}{2}at^2,$$

we find the maximum acceleration to be 1.11×10^{-5} me/sec^2, or about one millionth of a gravity. The compensatory response explanation requires the nerves and muscles to retain an accuracy of calibration some four orders of magnitude better than the best that is ever measured; it is flatly untenable.

In this case, we have an action (upward force due to muscles) canceling the effects of a disturbance (downward force due to gravity) to create a stabilized joint consequence (arm position, the second integral of the acceleration of the moving mass). This is precisely the situation to which control theory applies. A model is easily constructed and is worth a brief look as preparation for what follows.

First, there must be a sensor that detects the current state of the consequence to be stabilized, arm position, the sensor representing that state as a signal, y. Next there must be a comparator that generates an error signal, e, that is proportional to the difference between the actual state as sensed and a reference state, y_0 (usually supplied in the form of another signal, inside the system as a whole), letting the proportionality constant be unity, we have $e = y_0 - y$. Finally, the error signal enters an output actuator, which produces a degree of action proportional to the amount and sign of error signal; in this case, we call the action f, the upward force, and assume that $f = ke$, k being a constant of proportionality having units of force per meter of error.

To finish the closed loop, we have to add the effects of the action together with the effects of the disturbance to create a state of the stabilized consequence (or, more conveniently, its representation, y). The acceleration, a, of the arm in the $+y$ direction is net force divided by mass, or $(f - mg)/m$, or $f/m - g$ (g is the acceleration due to gravity). Thus $a = f/m - g$. We have three system equations:

$$e = y_o - y \tag{1}$$
$$f = ke \tag{2}$$
$$\ddot{y} = f/m - g \tag{3}$$

The acceleration, *a*, corresponds to the second derivative of y-position, or $\ddot{y}$ (the double dot signifies the second derivation with respect to time).

This set of equations can be solved as a system; the solution is *not* correct. The solution predicts some kind of endless sinusoidal oscillation, which we observe does *not* occur in the real case. Thus, a system analysis is not better than any other kind; it does not automatically give right answers if the model is not basically correct. One must still attend to nature and ask, when the model doesn't work, *why* it doesn't work.

In this case, it doesn't work because we have left out friction and all other effects that create (or simulate) a drag proportional to the velocity of a movement. We have modeled, in effect, a mass suspended on a perfect spring. To make the model work, all that is needed is to introduce into one of the equations a term involving the first time derivative of y, which may amount only to changing the part of the system description that pertains to the physical properties of the environment through which the feedback occurs. When that change is made, the model is given another degree of freedom in the form of the coefficient of a *damping* term. That coefficient can then be adjusted to make the predicted behavior match actual behavior quite nicely. This has been done so often that I won't bother to repeat it here. (See, for example, Stark, 1968 or Maxwell, 1868/1965.)

The only aspect of the solution of interest here is the fact that, with sufficient damping present, there is a steady-state condition predicted. If there is a steady-state condition, all the derivatives will become zero, meaning that $\ddot{y}$ becomes zero in Eq. 3. Making the appropriate substitutions from Eqs. 1 and 2, this produces

$$0 = k(y_o - y)/m - g \tag{4}$$

or

$$y_o - y = mg/k$$

In this equation there are two parameters pertaining to the behaving system, y_o (the setting of the reference signal) and k (the output force per unit error). The parameter k indicates the amount of steady-state change in output force that results from a steady disturbance-caused deviation of arm position and so corresponds roughly to the sensitivity factor that must be accurately calibrated in the compensatory response model. In this control model, as can be seen, the only requirement on k

is that it be *large enough.* For $y_0 - y$ to be less than .005 meter, and with $mg = 98$ N, k must be at least 1.96×10^4 (i.e., 20 N/mm). If k were then *doubled,* the result would be that the error, $y_0 - y$, would be *halved.* "Compensation" requires the accurate balancing of large opposing effects; control requires only high sensitivity to error. The actual sensitivity can vary over a wide range without any significant effect on the observable results. Thus, the control model works both qualitatively and quantitatively, without requiring impossible precision of the neuromuscular system.

This extremely brief run through a control analysis has been meant only to give the general idea, and especially to show how the subject of *perception* enters the picture almost without being noticed. Whatever a control system controls, it must sense. Thus, it is often the case that one cannot tell what is being controlled without finding out, somehow, what is being sensed. It is this close relationship of perception to control that creates the appearances misinterpreted for so long as a direct stimulus–response, input–output, cause–effect relationship between external events and behavioral actions.

3.2. *Control Theory and Causation*

According to the control model, stabilized consequences are stabilized because they are sensed and compared with a reference. An indication of the error is used to cause the very output that opposes the error. Disturbances tending to alter the stabilized consequence are not directly sensed but are sensed only through the resultant small deviations of the stabilized consequence, that is, departures of the signal representing it from the reference signal representing the target state. A system with very high error sensitivity (k in Eq. 4) makes large changes in its output before a disturbance has affected the stabilized consequence by more than a small amount; those changes of output are opposed to the effects of the disturbance simply because they are opposed to the sign of the error, however the error is caused. As a result, the effects of output actions on the consequence are seen to be balanced, at all times, against the effects of any disturbance.

The rub is that this balancing of action against disturbance can be seen *only* in terms of the stabilized consequence, the joint effect of action and disturbance that is being sensed and controlled. If control is effective, that consequence will not vary appreciably, and since it does not, *it will fail to show significant correlations with either action or disturbance.* The traditional statistical approach cannot reveal controlled consequences.

The observer who is trying to find out *why* a given action follows upon, or accompanies, a given independent environmental occurrence must be completely aware of the possibility that control is involved; if he is not, he is very likely to fall into the trap that nature has baited for us. The appearance created by control behavior is that the organism simply *mediates* between disturbance and action. Since both disturbance and action may affect the stabilized (joint) consequence through delayed, indirect, and nonlinear paths, and since more than one disturbance may be acting, the fact that action is continuously canceling the net effects of disturbances on some stabilized consequence is not transparently self-evident. What *is* evident is that action bears some regular relationship to disturbance. In terms of effects on a properly defined stabilized consequence, that relationship is quantitative and precise *opposition,* but even in terms of direct measures, more than a chance correlation between action and consequence will be seen. Adding to the clear existence of such relationships the *absence* of any clear relationship to the stabilized consequence (because it *is* stabilized), we have a situation guaranteed to lead the uninformed observer astray. A sudden disturbance will lead, a fraction of a second later, to a sudden change of action; a continuing disturbance to a steady bias of action; a varying disturbance to covarying actions. The appearance is that of a cause–effect relationship.

The conclusion into which we are enticed by that painted hussy is that the organism must be sensing the disturbance; that the disturbance is a stimulus that acts on the nervous system to *make* it produce the variations of action that are observed. That picture places the organism between disturbance and action, as a mediating input–output device that converts the former into the latter. One could write transfer functions describing that relationship to the sixth decimal place—and still have it wrong.

Suppose that by a stroke of luck comparable to the birth of Norbert Wiener, someone had appeared on the scene in about 1869, someone who was aware of Fechner's work in about 1860 on the psychophysics of perception and also of Maxwell's analysis of mechanical control systems, and who was capable of putting the two schemes together. The model that would have resulted would, of course, have been not electromechanical but neuromechanical; that is of no consequence. The important thing is that the model would have shown a way in which a system could have an internal reference, with respect to which it *controlled what it sensed.* It would have been realized from the start that there could be *apparent* cause–effect relationships seen in behavior that were illusions as convincing and as incorrect as the optical illusions with which psychologists were to become so preoccupied.

Had that theoretical advance occurred, it would have been impossible to argue, as influential biologists, neurologists, and behaviorists were to do for the next century, that behavior is necessarily controlled by events impinging on the organism. When half a century later Watson stood contemplating the bird in his hand struggling to get back into the bush, he would not automatically have rejected the appearance of goal direction as illusory, and he would not have been struck by the unfortunate insight that all responses could in principle be predicted on the basis of stimuli. That would have been a fine insight if Watson had been studying systems that were *not* control systems, but that case had already been handled some centuries before when the physical sciences rejected animism. Although Watson had at least a chance to launch the scientific study of *animate* systems, he missed it.

Another outstanding chance to found cybernetics early came in 1938 (there were several others between). In that year, B. F. Skinner found that when organisms were given the means to control their own food inputs, they learned the required behaviors far more rapidly than had ever been seen before. To Skinner, however, the significant fact was not the strong tendency of changes in action to counter the effects of disturbances on food inputs but the fact that specifying the way in which action affected consequence (the schedule of reinforcement) had very reliable effects on the actions that came to be performed. Skinner's ambition was control of the animal's behavior, not understanding of the animal's control of the inputs that mattered to *it.* With that emphasis, coupled with the *a priori* assumption that a proper theory of behavior had to leave all causes in the environment, Skinner missed his chance. We will always be in his debt for his discovery of operant conditioning and the equally important phenomenon of shaping, but it is too bad that he does not also have the credit for introducing control theory to psychology. Operant conditioning demonstrated the first organismic control behavior that was well enough instrumented to permit systematic deduction of the laws of negative feedback; it was not interpreted that way.

It fell to Norbert Wiener to recognize the similarity between the structure of servomechanisms being designed by technologists and the structure of neuromuscular systems. But by this time, the mid-1940s, the concept of behavior as a consequence of external circumstances had developed a stranglehold on thinking in nearly every field that aspired to the adjective *scientific.* Even though Wiener and his immediate successors saw all of the major implications of control theory with regard to goal seeking and purpose (Buckley, 1968), they were unable to shake off that last concept of cause and effect. Wiener (1948) drew a diagram of a control system that shows an input coming in from the left and an output leaving to the right, with *internal* feedback serving to make the output a more reliable function of the input. The fact that the input he was talking about was really the reference signal and that neuroanatomically it came from centers higher in the brain and *not* from sensory receptors seems to have been missed, perhaps by Wiener himself and certainly by scores of others who have re-created the organization of that fateful diagram in the literature all during the past 30 years. The general relationship among action, disturbance, and consequence that we have been examining, although repeatedly approximated, has for all practical purposes remained unknown.

I don't wish to argue *ad hominem,* but there comes a point when recognition of ordinary human frailty can be important. I think there is a very good reason for the almost universal failure to recognize the basic implications of control theory. The result of such recognition by any person with two neurons to rub together would have been the realization that the basic cause–effect relationship assumed for an analysis of "irritable tissue" had been in error from the start. To visualize the consequences if that were true would be to imagine the skyscrapers of New York going down like dominoes. Scarcely a single "fact" about organisms would remain intact. So what do human beings do when they realize that a train of thought is leading toward disaster? They think about something else.

3.3. *Control Theory and Purpose*

For the first 40 years of its existence as a science, psychology was billed as the study of consciousness. Introspection was a legitimate tool. The idea that people and perhaps some animals were purposive was widely accepted; each person could examine his own experiences and see plenty of examples of actions directed toward preselected goals. Even after behaviorism came on the scene and preempted the term *scientific psychology,* there were many diehards who knew that purpose was a fundamental aspect of human organization but couldn't justify this belief scientifically. The result was a deep split between "hard" and "soft" studies of human nature.

The basic cause of this split was a mistaken concept of cause and effect, coupled with the naive epistemology that goes with it. The main argument against purpose or intentionality as a factor in shaping behavior was that physical determinism made no room for inner direction of behavior, regardless of the weight of circumstantial evidence. The basic organization of the nervous system seemed fairly well understood; stimuli affected sensory nerves, which sent signals to higher centers, which relayed them and elaborated on them and eventually sent them outward in patterns that excited the muscles, thus producing what we recognize as behavior. Without realizing it, those who offered this picture were assuming more than neurologists could know.

All that neurologists knew then about the structure of the nervous system (and most of what is known even now) consisted of the general concepts of synaptic transmission, a few highly localized relationships, and some of the main pathways of neural signal flow. In an approximate way, some crudely defined units of behavior such as "movements" of limbs were known to arise when specific areas of the brain were "excited." Superimposed on the facts known about the components of the nervous system was an *assumed structure,* which combined the components into a model that would support a preselected cause–effect organization. As I pointed out in the beginning, assembling components into a structure introduces laws peculiar to the structure. In this case, the law was exceedingly simple: output is a function of input.

Many exceptions to the general input–output flow were found as neurological investigations continued. Both excitation and inhibition occurred; neural pathways were found that provided shortcuts from input to output at many levels lower than the cortex, and pathways were found from outputs back to inputs both inside and outside the system. Of particular interest to us are the discoveries of the many stages in *output* processes where *inputs* enter with a sign opposite to "commands" from higher centers, resulting almost universally in negative feedback and cancellation of most of the so-called command signal. It is possible—and as time goes on, more and more strongly advisable—to abandon the old input–output structure and to adopt a control-system structure for models of the central nervous system.

That structure provides a place for inner purposes. They can be identified functionally as reference signals, signals that specify to a given control system what level of its inner representation to bring about and maintain. A variable reference signal specifies corresponding variations in the controlled consequence being sensed and represented. A fixed reference signal specifies a static condition of the inner representation, although not necessarily a static condition of external physical variables (rates, velocities, and sequences can easily be represented as steady signals: consider *amplitude* and *phase*).

This change in the structure assumed for the brain changes all interpretations of the meanings of signals in the brain, particularly the outbound signals that used to be thought of as stages in the relaying of patterns from the cortex to the muscles. Now, they are seen as error signals, determined in part by reference signals from higher centers but determined just as strongly by disturbances tending to alter controlled consequences. There are no patterns high in the cortex that are relayed intact to the muscles. The patterning of behavior takes place on the input side.

As a result of this change in assumed structure, we arrive at a new concept of the meaning of *purpose.* A purpose is not an intended *action* but an intended *consequence* of action. Furthermore, in the final analysis, it is an intended state of an inner representation of that consequence—an intended state of a perception. With all the work currently going on in neurophysiological laboratories, following the lead of Hubel and Wiesel (1965), and with the addition of just a little common sense, we should begin to suspect that the world of perception that is experienced is constructed by computers at various levels in the brain, that *all* we experience is an inner representation, a signal that is an unknown function of unknown variables that no one can sense directly. When Hubel and Wiesel found neurons that responded preferentially to certain visual objects such as edges or oriented lines, how did they know what the electrode recordings corresponded to? By looking with their own eyes at the test stimulus, of course. What they found is not a neural correlate of an external object but a correlation between an electronic data display and a *subjective perception.* They were comparing two ways of looking at neural signals. They were no more able to peek past their retinas and see what was really causing those electronic or neural "meter readings" than is anyone else.

If we admit that the world we perceive is, at best, a *function of* external stimuli and not those external stimuli themselves, the concept of purpose can be defended against all traditional criticisms. The philosophers of behaviorism at the turn of the century objected to the idea of purpose because they took purpose to mean (a) intended *outputs* and (b) predetermined future consequences of present acts. If they could show that an output was "caused by" some external event (i.e., depended on it in a regular way), there was no need to introduce purpose as a second cause of the same actions. And, if an organism intended for a certain consequence of its actions to occur, they argued, how could failure to achieve those results be explained? *Predetermination* to them meant predestination, in the sense that a spring-powered watch that is not wound is predestined to stop running. Since no given act has precisely the same future consequences twice in a row, went the objection, there is no way for an act to be intentional in the sense of having a predestined result in the future.

Control theory bypasses all those arguments by identifying purpose with the specification of *reference levels for inputs.* There is no prediction of the future involved; the organism simply acts at all times in the direction that will oppose the present-time error. When disturbances arise or external conditions change in any (reasonable) way, actions adjust quite automatically and by understandable mechanisms to continue opposing error. It is not as if a single spasmodic action had to produce a predestined future consequence. The control system is always right there, continually altering its actions to keep the sensed consequence what it is intended to be (one level of sensed consequence could be a steady approach toward some final relationship).

3.4. *Control Theory and Volition*

One of the main reasons for rejecting the idea that *any* inner phenomenon could shape behavior was very practical. According to the cause–effect model that was thought to be required by physical determinism, it should be possible to study the way behavior depended on external conditions and eventually to achieve the ability to predict and control behavior through manipulations of external circumstances. "Prediction and control" was the slogan of science. Since a great many people wished to study organisms in a scientific manner, it would obviously have been inconvenient to believe that organisms contained *internal* causative agencies that could not be manipulated by an external experimenter. It became fashionable to assume that if any internal causative agency existed, it would behave lawlessly or, as a favorite pejorative had it, "capriciously." The unthinkability of capriciousness was taken as its refutation.

This point of view was unwittingly supported by the proponents of purpose; somehow the argument came to be not a factual dispute about internal causation but a philosophical argument about free will versus mechanistic determinism. As a result, the intellectual debates concerned the last chapter of a book, the main bulk of which had not yet been written. As we shall see, the most likely analysis of internal causation does not support the concept of either mechanistic external determinism *or* "free will" (whatever that is).

The reference signal given to a control subsystem tells that subsystem how much of its perception to want to perceive; the error is the *want.* In a *hierarchical* control-system model, the reference signal reaching a system of level n is a function of the error signal in a system of level $n+1$; variation of the lower-level reference signal is the means by which the higher-level system controls its own representation of reality. At each higher level, says the model, the representations that are controlled are invariants abstracted (neurally, not verbally) from the next lower level of representations. Therefore, for all practical purposes, the environment of a system at a given level consists of the neural representations that are under control by systems of the next lower level. The higher-level system acts on that environment by emitting reference signals into it (for the lower-level systems); the organism learns through experience what reference signals must be emitted in order to control any representation with which it is concerned.

This model shares at least one major capacity with its modeler: a concern with control. I have always thought it odd that a science dedicated to prediction and control omitted both capacities from its models of human beings.

The term *volition* belongs to the world of common sense, primarily because science has recoiled from it. Common sense, however, occasionally turns out to have found the right answer while science pursues the complex consequences of erroneous assumptions. Now that we know of a way in which inner purposes and intentions can exist, it is not a long step to see how the general phenomena called *volition* could be allowed, once again, to exist.

To say that one has performed an act volitionally is to say that there is no immediate external cause for the act. The situation is not that simple, however. An act at one level of description is only a means at the next higher level. I may say that I volitionally extended my arm and be quite convinced that there was an internal predetermination to perform just that act; yet, in a larger context, I can see that the goal of catching a falling vase, together with the geometry of the prevailing external circumstances as perceived, made that act and no other mandatory. To catch the vase, I had to get my hand under it, and nowhere else.

This has always been the main point of confusion about volition, even for common sense. Any action is at the same time an *intended goal,* continuously achieved, and a *variable means* adjusted according to the requirements of higher-order goals and external disturbances. The degree of volition one senses depends on whether he is focusing on the intended action (as the goal state of a perception of action) or on the higher-order *reason* for the action, the higher-order goal served by the action. When attention is on the higher-order goal, the lower-level action is sensed as *output;* when attention is focused on the intentional nature of lower-level action, the same action is sensed as an *input,* a perceived and controlled consequence of an output of still lower level (say, "effort").

The control model shows that all these interpretations are correct; there is no contradiction. In fact, these commonsense descriptions have an uncanny congruence with the structure that arises out of a control-system model. The control model elucidates the commonsense descriptions, by showing how one and the same behavior can seem both volitional and dictated by circumstances. The behavior of any control system—that is, the output it generates—is always governed by precisely those two considerations, acting jointly: the inner goal, which determines the level to which the perception will be brought and maintained, and external circumstances, which determine the amount and direction of the action needed at any given moment to keep perception in that goal state. Behavior is not directed by inner volition alone or by external circumstances alone: it is a joint consequence of *both*. The specific states of stabilized consequences of actions cannot be explained in terms of external circumstances; the details of action cannot be explained in terms of the inner goal. We are dealing with a system in which both of these considerations operate *at the same time*. Either–or cause–effect thinking is totally inadequate to handle this phenomenon. For control theory, it's a piece of cake.

The structure of our lower levels of control organization does not seem to us to be our own structure but rather the structure of the world of direct experience. One has to extend a hand to catch a falling vase because that is the way the outside world, not the nervous system, works. For that reason, we tend not to recognize the same means–ends relationships at these lower levels; we externalize them. Often we do this just by omitting the statement of the goal that is involved. One says, "I ducked because he took a swing at me," as if taking a swing, through some law of nature, could cause someone else to duck. In fact, one would not duck unless doing so served to keep some perception at the intended level: the perception of being hit, perhaps, at the reference level of *zero*. Under unusual circumstances, which ought to be enlightening but usually aren't, one may select a different reference level for the same perception; a small amount of being hit (a movie stuntman) or even a large amount (a person intending to prove that he is impervious to pain or above being bullied).

This omission of the implicit goal from discussions of cause and effect makes relationships between disturbances and actions appear to be cause–effect laws that work just as behavioral scientists have been assuming that the nervous system works. I strongly suspect that it is this commonsense model of lower-level actions that is behind the initial scientific models. It takes a rather strange set of circumstances to cause a person to say, "I couldn't get out of the classroom because the teacher was looking right at me" and then add," . . . and I didn't want to leave while she was looking at me." When the goal *is* specified, it seems superfluous to mention; it goes without saying. Unfortunately, it also goes without *thinking*, which is the main problem here.

I should mention another reason for omitting implicit goals, one that is at least more defensible (and seemed final in 1913). *Any* cause–effect relationship can be described in terms of an implicit goal: the negatively charged pith ball moved because a negatively charged rod approached it, and the pith ball didn't want to be touched by the rod. Even the most exhaustive examination of pith balls and charged rods, however, will fail to reveal the structure of a control system. Therefore, attributing goals to this system is simply a blunder. All that happens can be explained in terms of direct interactions between the entities involved (and a little unavoidable magic—"fields"). Inanimate objects almost never possess the internal structure necessary to permit purposive behavior. One *should* omit hypotheses concerning inner goals when such goals are ruled out.

3.5. *Ultimate Purposes*

Anyone who proposes a hierarchical model has to realize, sooner or later, that he has a problem: the highest level. As long as one can remain comfortably in the middle or lower reaches of such a model there is no conceptual difficulty, but the model is not limited in the upward direction as it is in the downward direction. There is a temptation to achieve closure, to settle on some way of topping off the structure, even if one doesn't know how to.

I don't know how to; perhaps that admission is as important a part of this model as the parts that have been positively proposed. I am refusing to fish for some way to make the whole structure complete just to lend support to a philosophical prejudice (I tried it once, in my last book, and the discussion got so fuzzy as to be useless).

What can be said about the higher levels is mostly negative. We *do not* know the basis on which the highest-level goals are set. We are *incapable* of tracing them to any specific external circumstances, particularly not present-time circumstances. We can offer some reasonable conjectures about how biochemical and genetic factors enter, particularly in connection with learning, but we can *by no acceptable scientific means* show that those factors are "ultimate" determinants, not in any sense. It is time to stop trying to make everything fit 19th-century ideas of physical determinism, which are based on little more than an allergic reaction to religion. The upper regions of human organization are a mystery that we have barely begun to approach; we will never understand them on the basis of a jab-and-jerk model of behavior.

4. CONSCIOUSNESS

As long as behavioral science worked under the assumption that the brain is simply a link in the chain from sensory input to motor output, subjective perception, memory, and feelings of volition and intention had to be treated as epiphenomena. Like the lights on the front panel of a computer, such phenomena of consciousness informed the observer about what was going on but had no influence on the operation of the system. To "be scientific" about behavior meant, largely, to *redefine* subjective phenomena, finding words that would permit one to deny internal causes and assert external ones. Calling every action a *response* was just one such semantic ploy.

This denial of an active role to consciousness has, of course, created a gulf between scientific psychology and the common sense of laymen. Probably as a defensive measure against commonsense criticism, there has grown up an unspoken attitude that seems detectable whenever behavioral scientists get together to talk about behavior. If it were put into words, it might sound something like this:

> Sometimes it seems that organisms seek goals, want things, decide things, and act spontaneously. But you and I know that such appearances are illusions, because we know that natural laws of physics and chemistry are behind all such appearances. In ordinary affairs, we use ordinary language for convenience, but when we want to speak as scientists, we have to put metaphysical nonsense aside. If laymen criticize our scientific descriptions, that is only because of leftover superstitions, beliefs, and sloppy habits of thought from their primitive past. We scientists do not have any superstitions, beliefs, or sloppy habits of thought, especially since such things do not really exist in the first place. And if you won't call attention to my consciousness, I won't call attention to yours.

As I tried to show earlier, the systems approach brings out the clear existence of laws of structure that cannot be traced to, or deduced from, laws of physics and chemistry. For three-quarters of this essay, I have been trying to show that one particular systems approach, control theory, offers strong grounds for denying that the brain simply mediates between events affecting the senses and subsequent behavioral actions. Now, the argument can be carried to its logical conclusion: the phenomena of consciousness can no longer be dismissed by mutual agreement and must be studied as causative phenomena just as real as the physical events that are involved in behavior.

4.1. *Causative Factors in the Brain*

While it is not yet possible to trace reference signals to the highest level of the brain's control hierarchy, certain generalizations can be defended even without knowing the architecture of the higher levels. The main working hypothesis concerns relationships between levels.

A control hierarchy involves a hierarchy of perceptions—representations of higher-and-higher-order invariants constructed on the raw material of lower levels. It also involves a hierarchy of adjustable purposes: reference signals that specify to a given control system whether it is to keep its perception at a minimum, intermediate, or maximum level. As we have seen, the action performed by a given system (that is, the lower order reference signal or, for first-order systems only, the motor signal) reflects both the setting of a reference signal received from higher systems and the presence of disturbances that tend to alter the perceptual signal being controlled by the given system. For a fixed reference signal from above, the action is determined by disturbances; for a disturbance-free lower-order environment, the action is determined by the higher-order reference signal. We are concerned for now with this latter case, in which no significant disturbances are present and actions at the level of interest result primarily from the settings of reference signals.

Consider a control system that senses and controls one dimension of a repetitive pattern, say tapping out a steady rhythm on a tabletop with one finger. To implement a speed-of-tapping control system, there must be a perceptual device that reacts not to each tap but to the average rate of tapping; the output part of the system must respond to error by varying the speed with which the reference signal for finger position is switched back and forth between "up" and "down." If the rate is sensed at the "slow" end of the perceptual scale, and if the reference signal is set at a magnitude corresponding to "fast," the resulting error signal should cause the output function to increase its speed of alternation of position-reference signals. If the system is sufficiently sensitive to error, it will come to a steady-state condition with the sensed rate of tapping matching the steady reference signal, the sensed rate being in the form of a steady signal also.

This control system needs no external stimulus to keep the tapping going; it needs only a steady value of reference signal, supplied from higher systems. The behavior is "spontaneous," in that there is no series of environmental events that could be said to cause the individual tapping events, in any one-to-one way. If a disturbance occurred—say, the ambient medium were changed from air to cold molasses—the changes in reference signal for finger position might become more exaggerated and might advance in phase, and that change in the output pattern could be said to have been externally caused, even though the final result is thereby kept from being affected. But the central consequence, the controlled rate of finger-tapping as perceived, must be considered to have its immediate cause inside the person doing the tapping.

In any specific instance of a control phenomena, it makes no difference whether the controlled quantity is a static or a dynamic condition: the immediate cause of the phenomenon, in the absence of significant disturbances, must be attributed to the next higher level of organization, further removed from the external environment.

Even when disturbances are taken into consideration, the same kind of role is found for higher-order systems. A change of output in response to a disturbance always occurs so as to resist changes of the affected perception away from some given reference level, determined by a higher-order system. The response to a disturbance depends not on the disturbance alone, but on where the perceptual signal is relative to its reference level (above or below) and the effect of the disturbance relative to that reference level.

In a search for the causes of any given action, the control model thus always takes us farther from the periphery and into levels involving higher and higher orders of invariants constructed by the nervous system. The direction toward ultimate causes is not back toward the sensory inputs but in exactly the opposite direction, deeper into, or higher into, the hierarchy.

Somewhere in that hierarchy, we must eventually find all that is experiencable, and that includes not only representations of the physical and the physiological environments but relationships among elements of the environment, structures of logic concerning those relationships, abstract principles by which programs of logical thought are directed, and systems concepts that allow us to represent collections of principles in terms of models. In short, we must sooner or later come across *every object of experience.* Every time we ask why a given perception is controlled, the *why* will become *how* a higher-order perception is controlled. The ultimate model of the nervous system will be a model of our entire structure of purposes, and identically a model of the world on which we act to carry out those purposes. The whole will be a model of experience, everything to which a person can attend whether he refers to it as being inside or outside himself.

Clearly, the causes of behavior cannot be understood until we understand this entire structure, which I assume to be the structure imposed by the organization of the brain. It is not possible, in these terms, to speak of "the" cause of anything a person does; any attempt to find such causes leads inevitably to the necessity of understanding the entire hierarchy and its interactions with that world outside that we have barely begun to deduce. The laws of the external reality, as they affect us, are inextricably interwoven with the laws of structure that reside in our brains and mark us as human.

4.2. *Ultimate Causes*

We are not born with this inner structure completed. It grows, and takes on its adult form through decades of incessant interactions with the external world. When we fail to control perceptions of various types, either because we have not yet constructed those invariants or because we have not yet stumbled across a means, the failure has many consequences that affect us through routes other than the senses. To take a simple case, if we do not learn to control the efforts, tastes, sounds, movements, sequences, and relationships involved in eating, we suffer direct physiological effects that would occur even if all our senses were numbed. Somewhere inside us, there is a link between these unsensed effects and the way the hierarchy of control changes and grows. And inherent to this link, there must logically be some specification of the *proper* states of the physiological variables involved. That specification is certainly not acquired by an individual during his lifetime. It must have developed over countless thousands of years, through slow processes of genetically preserved change. It is not subject to present-time manipulation.

If there can be any ultimate determinant of the way we learn to think and act, it must be in this set of inherited specifications for the state of the physical organism that calls for no change in behavioral organization. When these specifications are not met, the result is change—change of the very structure that is the hierarchy we have been discussing. I have called these specifications *intrinsic reference levels,* and I have viewed the process of change as a kind of metacontrol system that alters structure as its means of controlling what it senses. Our highest purposes have to do with continuing to be human.

To some, this invocation of physiological variables may come as a relief—finally, we are back to physics and chemistry! But I don't think we can afford to be as simplistic as that. What we are talking about is *organization,* not components; we are talking about the same kind of organization of stored information that can make us grow kidneys and fingerprints. The physics and chemistry of DNA and RNA provide us with a model of only the tape and the tape recorder (more or less); they do not explain the origin or the evolution of the message. We may be dealing here with the highest-order invariants of all. The "state of the physical organism" can mean anything from the *p*H of the bloodstream to the elegance of a solution to a conflict. I am not willing to put any restrictions on the subject matter to which intrinsic reference levels might refer. Laws of structure are involved, and we know very little about the number of levels of structure that have laws of their own.

4.3. *The Nature of Consciousness: Point of View*

There is one fact about the nervous system that is very easy to forget, or ignore. The nervous system we know about through sensory experience bears practically no relationship to the nervous system we use in our theories. The nervous system of naive perception is a gray blob, which under a microscope becomes a mass of separate blobs and filaments. We see no organization there; we imagine it. We have pieced together a mental model of events in that gray blob, trimming, rearranging, and adding to the mental model until it behaves more or less as our instruments and logic combined report real brains to behave.

Somewhere in that gray blob, we must conclude, is the mental model of the gray blob. We don't know anywhere *else* it might exist. We deal, perforce, with what von Foerster (1974) has called a "recursive system." The system, as his students have put it, computes that it is computing. The experience experiences itself.

This bothers me. Something has been left unsaid, or unthought, or unnoticed. I think the problem lies not in our models but in a certain attitude toward models, one that encourages us to follow our own logic once around the loop and then forget to follow it around again. What is a "recursive system"? It is first of all *an idea about systems.* More: it is, for those who think about it, an *object of experience,* just as much as a toenail is, although it does not share the same space in which toenails exist. To think about recursive systems, or nervous systems, or control theory, or any conceptual scheme at all requires that someone adopt a particular structured point of view toward lower-order experiences.

That is what bothers me: the fact that there is always a point of view, and that it is structured. What is it that can adopt, and abandon, points of view? Where do we find it in this hierarchical model of perception and control of perception?

The answer is that we do not find it there. We do not yet have a model that can reproduce this phenomenon I call *point of view.* The structure that is involved in any given point of view is the acquired structure of the brain; that much is not hard to fit in. But the brain, once organized, contains *all* points of view possible to it, even those not currently operative. Moreover, points of view can shift up and down in the hierarchy suggested by this model; one may become a configuration-recognizer and may after a while become a relationship-recognizer. Yet there could be no relationships, such as "next to," if the lower levels were not still faithfully constructing the configurations that are related. Something moves. Something that is not yet represented in the model.

5. CONCLUSIONS

If these explorations have led me anywhere, it has been to a vivid sense of my ignorance. Control theory, I believe, turns us around and pushes us off in a new direction, and leads us to see how little progress toward understanding our own existence we have really made in the past 300 years. It does not allow us to see very far ahead in that new direction.

It does do something of value, however, by forcing us to look at ordinary experience as evidence of our own structure as much as evidence about an external reality. When thinking about hierarchical levels of perception, one is forced to move his point of view upward (or whatever the direction is) merely to try to comprehend the currently operative structure that gives form to the current, or just previous, point of view. It forces us to experience this phenomenon in action rather than just constructing words or block diagrams about it. It pins the label *model* on our models, which otherwise we would be inclined to accept without thinking as truths about a hypothetical external world.

Most important to me, it leads almost inevitably to the realization that the objects of experience, from the "concrete" to the "abstract," are indeed *objects of* experience, phenomena moving in and out of the field of experience so effortlessly and easily that we scarcely see any significance in their comings and goings. Yet, I think those comings and goings constitute a higher-level fact that is at least as important as the specific items that come and go. A relationship, such as the distance from your eyes to the page you are reading, does not come from anywhere when you notice it or go anywhere when you are tired of noticing it. It appears in experience when one adopts the point of view of a relationship-perceiver. Perhaps the signal was there all along; under some circumstances involving several levels of hierarchically related control tasks, the signal *has* to be there all along, even when not being experienced.

It would be gratifying to know how to test this phenomenon of point of view. If it exists, it must have effects; it is not there merely for the amusement of a passive occupant. I think I know of one way, and I am slowly getting organized to try it. Even though this is still just an idea, it may be worth mentioning so that others so inclined might poke around in this direction, too.

Point of view seems to me to involve both what we call *awareness* and what we call *volition*. Furthermore, for reasons that will probably not stand the light of day, I have a suspicion that awareness and volition are closely related to all the manifestations of *change of organization.* These thoughts lead to the idea of constructing some multileveled control experiments (experiments concerning control, that is) in which subjects are encouraged, asked, or underhandedly forced to concentrate on one of the hypothetical levels of perception involved, enough to drive the others out of immediate attention, much as a reader concentrating on the meanings of a string of words may fail to realize that all the *t*'s are taller than any of the *a*'s. If awareness has the properties of something that can move from one point of view to another, and if it is intimately or even necessarily involved in the process of reorganization, the parameters of control ought to become variable at the level where the current point of view is located. It is not excessively difficult to monitor a few basic parameters of control, such as sensitivity, phase shift, and RMS (root-mean-square) error on a continuous basis, and I have succeeded in doing this for some simple two-level (supposedly) control tasks. It will not be easy to prove that the control system in effect is multileveled and not just complex, but I think there are some ways involving simultaneous application of different kinds of random uncorrelated disturbances and others involving looking at reaction times. At any rate, this project seems worth pursuing, since I don't know of any other scientific approach that might reveal some *property* of consciousness. I will report on the results in due course, if I am not too old by the time there are any.

REFERENCES

BROWN G. S. *Laws of form.* Oxford: Blackwell's, 1969.

BUCKLEY, W. *Modern systems research for the behavioral scientist.* Chicago: Aldine, 1968.

CAMPBELL, D. T. "Downward causation" in hierarchically organized biological systems. In T. Dobzhansky & F. J. Ayala (Eds.), *The problem of reduction in biology.* London: Macmillan, 1973.

DEWEY, J. *The reflex arc concept in psychology,* In W. Dennis (Ed.), *Readings in the history of psychology.* New York: Appleton-Century-Crofts, 1948, pp. 355–365. (Originally published, 1896.)

HUBEL, D. N., & WIESEL, T. N. Receptive fields and functional architecture in two nonstriate visual areas of the cat. *Journal of Neurophysiology,* 1965, *28,* 228–289.

JAMES, W. *The principles of psychology.* New York: Holt, 1890.

KELLEY, J. R. *Manual and automatic control.* New York: Wiley, 1968.

MAXWELL, J. C. On governors. In *The scientific papers of James Clerk Maxwell, Vol.* 2. New York: Dover, 1965, pp. 105–120. (Originally published, 1868.)

PATTEE, H. H. *Hierarchy theory.* New York: Braziller, 1973.

POWERS, W. T. *Behavior: The Control of Perception.* Chicago: Aldine, 1973.

POWERS, W. T., CLARK, R. K., & MCFARLAND, R. L. A general feedback theory of human behavior, Parts I and II. L. V. Bertalanffy & A. Rapoport (Eds.), *General Systems: Yearbook of the Society for General Systems Research.* Ann Arbor, Mich.: Society for General Systems Research, 1960, pp. 63–83.

STARK, L. *Neurological Control Systems.* New York: Plenum, 1968.

VON FOERSTER, H., and students. *Cybernetics of Cybernetics.* Urbana: University of Illinois Biological Computer Laboratory, 1974.

WIENER, N. *Cybernetics.* New York: Wiley, 1948.