CONFERENCE WORKBOOK

for

"TEXTS IN CYBERNETIC THEORY"

An In-Depth Exploration of the Thought of

HUMBERTO R. MATURANA
WILLIAM T. POWERS
ERNST VON GLASERSFELD

A Conference of

THE AMERICAN SOCIETY FOR CYBERNETICS

October 18-23, 1988

Felton, California

TABLE OF CONTENTS

TEXTS IN CYBERNETIC THEORY

Designed to provide an opportunity for informal, serious study of three viewpoints in cybernetics, this special conference of the American Society for Cybernetics will devote each of the first three full days to reading, examining, elucidating, and discussing a specific text embodying the primary ideas of a particular cybernetic theoretician. The day's author will respond to questions of explication arising from small-group study of the text, as well as provide additional elaboration of his theoretical viewpoint and its implications in an evening lecture accompanied by further questions from the floor as well as general discussion. The aim of each "author day" will be to understand the author's viewpoint. The final day and a half will engage the authors in dialogue and discussion of issues that have emerged in the previous days.

In addition to promoting a deeper understanding of three major points of view in cybernetic theory, the conference will provide three of our theorists the rare opportunity of being heard very carefully--simultaneously offering each participant an opportunity to examine more deeply his or her own theoretical constructs. In short, the conference aims to foster a context in which all of us can learn and explore together, freeing each other from the stifling mode of "my ideas against your ideas" and instead working together against the ideas: to clarify as fully as possible some of the major current ideas in cybernetics, as well as their implications.

All conference participants will be expected to have read thoroughly each of the three papers contained in this conference workbook prior to arriving at the conference. Participants can facilitate discussion at the conference by making note of questions which arise as they familiarize themselves with the texts in the weeks preceding the conference. At the same time, since the whole point of the conference is learning together, no one should let relative lack of familiarity with all the minutiae of one or more of the theoretical viewpoints prevent their attendance. Also, out of courtesy to our speakers and to the total weave of the conversation, all participants are expected to attend the entire conference, and there will be no late arrivals.

Our three speakers have worked hard to provide us excellent self-contained presentations of all the major facets of their respective theoretical viewpoints--as you will see in the following pages. If in reading these texts we exercise half the care they devoted to writing them, we will have a fruitful conference indeed. Come prepared to work, and, as always in the ASC, to have fun.

---Rod Donaldson,
Conference Chair

An Outline of Control Theory

William T. Powers

Nearly 100 years ago, William James pointed out that organisms differ from every other kind of natural system in one crucial regard: they produce consistent ends by variable means. He made this observation just at the dawn of so-called scientific psychology: his words were quickly forgotten. In their eagerness to make the study of behavior into a science, the American psychologists who became the intellectual leaders of the movement called behaviorism decided to let pure reason govern their approach. In a physical universe, one seeks the LaGrangian: the summing-up of present causes in sufficient detail to allow prediction of future effects. Because the universe is lawful and regular, they reasoned, regularities in behavior must be caused by regular influences on the behaving organism. Thus to predict behavior, all we had to do was study the conditions under which it took place with sufficient precision and care. From such studies would come behavioral laws like the laws of physics. Using these laws the psychologist could then not only predict what behaviors would occur, but by manipulating the environment, control behavior.

From the very beginning, therefore, scientific psychology assumed a property of behavior that is precisely the opposite of the one William James noticed. The psychologists decided that if regularities of behavior occurred, they could be traced back to regular antecedents, and that by manipulating those antecedents they could cause the behaviors to occur again. In this way they created an imaginary kind of organism that behaves in a way that real organisms do not behave, and proceeded to spend the next nine decades -- so far -- trying to make real organisms act like the imaginary one.

This imaginary organism is in fact far older than behaviorism. It came into existence with Galileo and Descartes. The early successes of the physical sciences were based on the fact that in at least some regards, the non-living natural world behaves regularly when subjected to regular influences. The world is a mechanism, and mechanisms do only what they are made to do by outside forces. All of the sciences of life, as they firmed up, sought to apply the same successful methods to determining the mechanisms of life. Behaviorism was born of these earlier approaches; in fact it was directly shaped by the thinking of biologists.

To speak of the "mechanisms" of life is to make a number of subtle but powerful assertions. The subtlest is this: if organisms are mechanisms, they are operated by the world around them. To explain their behavior, therefore, we need look only at their surroundings, and of course at their physical makeup. The physical makeup, however, only establishes the physical thing on which the environment works: without some external force to act on it, the mechanism will do nothing. Whatever it does do, it is caused to do.

This conception of life meant, of course, that to explain behavior we needn't refer to anything inside the organism. No concept of consciousness, thought, or will was needed, because if all behavior could be explained by referring to visible causes, what more could we add to the explanation by assuming inner causes as well? What would be left for them to cause? This line of argument, of course, assumed something that was very far from accomplished: that we could, in fact, account for behavior in terms of external causes.

As the twentieth century got under way, and as more and more scientists pledged alliegance to the principle of external causation, a disinterested observer might have noticed a peculiar fact. Every single attempt to explain behavior in terms of external causation failed. Each one failed, that is, in any terms a physicist or an engineer might apply. Instead of regular responses to outside stimuli, experimentalists kept finding only irregular responses, so irregular that it often took hundreds of trials or hundreds of experimental subjects to reveal that some regularity might lurk beneath the otherwise random-looking data. By the 1930s it had become obvious that the regularities of behavior were all but hidden because of a new property that was named "variability."

So the sciences of behavior became mostly ways of applying statistics to ferret out suggestions of regularity. If there had not been such an enormous commitment to the causal picture of behavior, and so many earnest efforts to show that it was really correct, there would have come a time when these scientists would have stood back, assessed the situation, and given up the basic assumption as a failure. Any physical scientist would have done so long before.

By the 1930s the cause-effect assumption was, however, far too well established to be thrown out or even seriously questioned by mainstream scientists. Essentially all scientific work regarding behavior was based on looking for regular causes of regular behaviors -- or at least for correlation coefficients that might be taken as hinting at such a relationship. The scientific world had settled on a general picture of the mechanisms of behavior, and while there was continual wrangling about just how this or that cause affected behavior, there was no disagreement about causality itself.

To this point, the concept of mechanism had essentially only one meaning: a sequence of causal links that began with some primary effect and propagated, one link to the next, until it terminated in some observable event. One part of the mechanism affected the next, and so on to the final effect. But on the morning of August 2, 1929, a Bell Laboratories engineer named H. S. Black discovered a principle that brought a new kind of mechanism into view. On that morning, on the way to work, H.S. Black suddenly understood how to analyze negative feedback.

The artificial control system

Black didn't publish his discovery for four years, but it quickly became the foundation for a new approach to the design of physical systems. The basic problem Black had solved was this: given an electronic amplifier that had part of its output connected to subtract from its input, how could this feedback arrangement be stabilized, so it would not "run away?" Obviously, one answer is not to feed back very much of the output: if the feedback effect is very small, nothing untoward will happen. But what if the net amplification factor, tracing completely around the feedback loop, were very large -- say, 1000? This would seem to mean, under the old causal analysis, that any small disturbance would be fed back to the same place with 1000 times the amplitude -- and the next time around it would have become 1,000,000 times as large, and so on. Black showed how an amplifier with <u>any</u> magnitude of "loop gain" could be made stable, provided that the feedback effect opposed the initial disturbance -- that the feedback was negative, not positive. The trick Black discovered was how to make the feedback stay negative.

Systems with large amplification and stable negative feedback soon proved to have some fascinating properties. Their behavior seemed almost independent of their physical properties. Even though stabilizing them meant slowing their responses somewhere in the feedback loop, they were capable of far faster and more precise action than systems without feedback. The speed lost through the slowing factors was far more than made up by the fact that very high amplifications could be used.

Black was primarily a telephone systems engineer, looking for ways to build reliable long-lived amplifiers out of imperfect components. But there was another branch of electrical engineering that found a different use for his principles, the branch that eventually came to be known, early in World War II, as control-system engineering. During the 1930s some engineers were looking for ways of substituting automatic machinery for human beings in certain tasks, primarily tasks that took a whole human being's attention full-time just to keep some simple physical variable like steam pressure or airspeed under control. There was nothing in any existing theory of behavior that could explain how a human being managed to accomplish even the simplest of these tasks. Theories of behavior were long on metaphor and qualitative assertions, but very short on instructions for how to build a machine that would behave as organisms were assumed to behave.

An engineer, some engineer once said, is someone who learns what is necessary to get the job done. In this case, what the engineers had to learn was how organisms really work. They solved this problem from scratch, inventing in the process a new kind of machine. Being interested only in the machine, they didn't realize that they had revolutionized the sciences of life.

It is probably no coincidence that these engineers worked primarily with electronic systems. They were accustomed to systems in which there were no moving parts except at the output, systems in which everything interesting took place in the form of changing voltages and currents. An electronics engineer was perfectly happy to point to a circuit chassis and say, "That's the RF signal, and here's where it gets turned into the IF, and here is the detector that turns it into audio, and here is where the music comes out." In fact, all those currents and voltages were just currents and voltages, until they were named and given functional meaning by the engineer. So there is something appropriate about the fact that engineers working with networks of anonymous and essentially identical electronic signals managed to discover how to build machines that imitate, in a rudimentary way, the kinds of behavior that are accomplished by a brain: a brain in which there are no moving parts and everything that happens occurs in the form of networks of anonymous and essentially identical neural signals.

To shorten the story, the engineers eventually discovered that in order to control some physical variable, a control system had to have certain basic parts, connected in the right relationships. First, whatever was to be controlled had to be continuously represented by an electronic analogue signal. If a position of an object was to be controlled, some measuring device had to be attached to the object so that as the object moved from point A to point B, an electrical signal changed from magnitude A to magnitude B. This was the sensor.

Second, not surprisingly, the control system had to be able to affect whatever was to be controlled. An electronic signal inside the system had to be converted, through an effector, into some physical effect that acted on the variable to be controlled. If an object's position was to be controlled, then the effector would be a motor or a pneumatic piston or a solenoid. For the best control, the amount of action had to be essentially proportional to the amount of driving signal, although it was found that this proportionality could be very approximate.

Having thus dealt with the input and output processes, analogous to human senses and human muscles, the engineers then tackled the third problem, the heart of the matter. Exactly HOW did the sensory signal have to affect the output effector to get the result envisioned -- control of the external variable?

It's clear that if the sensor indicates that the position -- or whatever -- is in error, the sensory signal should operate the effector to make the position or whatever change back toward the right state. A positive deviation should lead to an effort having a negative effect on the deviation, and vice versa when negative and positive are interchanged. Negative feedback. The problem was that you can't simply connect the sensor's signal to the effector and get the right result. If you do that, the control system will energetically force the position/whatever toward the state that creates zero sensory signal. If all you want is to keep the position/whatever nailed to the low end of its range of variation that will do fine (although a nail would also work), but what if you want to control something around some state other than zero, or around a variable state?

Consider the poor stationary engineer whose job it is to stand with one hand on a valve wheel and keep a steam pressure gauge at a constant reading. He may not even know that the wheel changes the draft in a furnace and varies the boiling rate of water in the pressure vessel. His job is to keep that needle at the right reading, and all he has to know to do this job is that turning the wheel clockwise will raise the reading and turning it counterclockwise will lower it. Or is that all he has to know?

Actually, he has to know one more fact: the right reading. The dial tells him the present pressure, but not the right pressure. If the dial indicates 328 pounds per square inch, that is too much, and he has to turn the valve counterclockwise. If it indicates 326 pounds per square inch, that is too little and he has to turn the valve clockwise. Only if the reading is 327 pounds per square inch is it all right not to turn the wheel. As the factory is putting widely varying demands on the steam supply, the engineer hardly ever gets to leave the wheel alone and think about philosophy.

So how is the control-system engineer to get that "right reading" into the control system? It's just one position of the needle among all the positions the needle might have, and a phone call from the production manager might result in making some other reading the right one, so 327 pounds now calls for turning the wheel right or left. There is clearly a reference-reading against which the actual reading is being compared, and that reference reading, to have any effect, must be carried inside the human being's head. So the control-system engineers had to provide a reference signal inside the control system they were building. The reference signal represented the intended pressure.

The sensor represents the state of whatever is being controlled as a signal, a voltage with an analogous magnitude. It makes sense to compare one voltage to another, and that is what was done: the reference signal was also a voltage. In the nick of time, the 6SN7 vacuum tube came along and (in a circuit called a differential amplifier or "long-tailed pair") provided the basis for an electronic comparator that could generate an output voltage that was reliably proportional to the difference between two input voltages. One input voltage was the sensor signal, the other the reference signal. And now the output of the system could be zero when the input was NOT zero. A motor connected to the draft-adjusting valve could stop turning when the error signal coming out of the comparator was zero, which occurred when the sensor voltage was, say, 32.7 volts, just matching the reference voltage of 32.7 volts. The sensor and reference signals, of course, were calibrated so that one volt meant 10 pounds per square inch in this imaginary but generic design. The sensor didn't read the dial: it was the same pressure sensor that made the needle move.

Now if the pressure was too low the motor would turn one way, if it was too high the motor would turn the other way, and if it was "just right" -- meaning that the sensor signal matched the reference signal, whatever its setting -- the motor would not turn at all. The control-system engineer could then explain to the stationary engineer that his life of drudgery was over, and also that he had lost his job.

Verbal descriptions of the way control systems work are almost certain to be misleading unless critical details are spelled out with care. The sheer mechanics of speaking or writing stretches out the action so it seems that there is a sequence of well-separated events, one following the other. If you were trying to describe how a gun-pointing servomechanism works, you might start out by saying "Suppose I push down on the gun-barrel to create a position error. The error will cause the servo motors to exert a force against the push, the force getting larger as the push gets larger." That seems clear enough, but it's a lie. If you really did this demonstration, you would say "Suppose I push down on the gun-barrel to create an error ... wait a minute. It's stuck."

No, it isn't stuck. It's simply a good control system. As you begin to push down, the little deviation in sensed position of the gun-barrel causes the motor to twist the barrel up against your push. The amount of deviation needed to make the counteractive force equal to the push is so small that you can neither see nor feel it. As a result, the gun-barrel feels as rigid as if it were cast in concrete. It creates the appearance of one of those old-fashioned machines that is immovable simply because it weighs 200 tons, but if someone turned off the power the gun-barrel would fall immediately to the deck. Nothing but the effector, the motor's armature suspended on good bearings in a spinning magnetic field, is holding it in place. The motor does this because the control system is exceedingly sensitive to tiny deviations of sensed position away from the reference position.

The gun is so well-stabilized that it resists any amount of push you can exert, without a tremor.

The operator of this gun, on the other hand, can easily make it swivel from one position to another just by turning a knob between two fingers. The knob varies the reference signal. When the reference signal changes, the definition of "zero error" changes, and the control system acts instantly to make the sensed position stay in a match with the new definition. If the operator twiddles the knob idly back and forth, the motor and gears may scream and the lights may dim, but the gun-barrel will also twiddle idly back and forth under precise control.

World War II started only six years after Black published the secret of negative feedback, and sophisticated control systems were pointing gun-barrels before the war's end (I learned to troubleshoot and repair control systems during that war). Into the middle of this feverish development came Norbert Wiener, Arturo Rosenblueth, and Julian Bigelow. They were not the only people to see that control systems behaved in some mysterious fashion as if they were alive -- even teenaged Electronic Technician's Mates could see that -- but they were the only ones with an ingenious name for this phenomenon: cybernetics, from a Greek word for steersmanship.

Cybernetics

In 1948 Norbert Wiener published Cybernetics: Control and communication in the animal and the machine. In this book he showed that the organization of a negative feedback control system was in one-to-one correspondence with the organization of certain neuromuscular "reflex arcs;" he even suggested new ways of looking at purposive or directed behavior as a whole in terms of control theory. This topic interested many others, and soon gave rise to the Macy Conferences, at which gatherings of scientists explored not only control-system theory, but other topics such as information theory, communication, and self-organizing systems.

The next major publication was W. Ross Ashby's Design for a brain, in 1952. Here Ashby took the basic control-theoretic idea and expanded on it in detail. Among other important concepts, Ashby introduced the idea of "ultrastability", a special property that he gave to a multi-control-system model that enabled it to maintain itself as a control system under drastic changes in its surroundings, even in its own circuitry. This was the first clear statement of a model of organisms showing how they could be responsible for their own organization.

Unfortunately, engineers were under-represented in the early ranks of cyberneticists, one primary exception being Bigelow, who considered himself, however, a proponent of information theory.

Perhaps if engineering experts on control theory had been called in early in the game, their conventional and practical knowledge of control systems would have completely stifled the inventiveness that kept cybernetics going. But a price was paid for that intellectual freedom.

It was clear to all the early cyberneticists that control systems behaved in ways that were very different from any concept of behavior that had existed until then. Instead of action being the end of a causal chain, it was simply one part of a closed causal circle. The relationship between organism and environment, when organisms were seen as control systems, was no longer one of obedience to external forces. Instead, the organism itself became an active agent in the world, its inner organization being responsible for what it did. The early years of cybernetics were full of the excitement that comes from seeing a familiar phenomenon in a new light. The implications of circular causality were simply enormous. Studying behavior suddenly became far less important than studying the inner organization of the brain: its inner logic, its use of language, its capacity to do something with incoming stimuli beside respond to them in a blind mechanical way. Organisms began to appear autonomous.

All these new concepts followed, however, from a basic new conception of mechanism that few cyberneticists understood. Most of those who attached themselves to this movement were attracted by what seemed a series of exceptionally coherent insights into the nature of behavior, insights that came, apparently, from nowhere, or at least from a few outstandingly ingenious minds. Most of these cyberneticists understood that somewhere in the background was some technological stuff that had got the whole thing started, but they were not technologists and weren't very interested in machines. It was this new collection of concepts that caught their attention. So they began to guess about how such systems might be organized so as to behave in this new way.

There is where the price of ignorance started to be paid. In fact the basic principles of operation of closed-loop systems had been worked out in considerable detail before Wiener and his colleagues ever appeared on the scene. Machines that imitated the purposiveness of human behavior had been designed after a careful analysis of how human beings behaved in that same way (although without any intention of explaining human behavior). The mathematics needed to analyze circular causation, based largely on H. S. Black's work, had matured and was in regular use by engineers. The machines whose behavior inspired the birth of cybernetics were already understood. There was no need to guess about how these newly-appreciated phenomena came about.

What cybernetics had to add to this picture was not an explanation of closed-loop phenomena, but a creative exploration of the significance of these new principles as they applied to human behavior. In large part, and to the degree possible at the time, this was done. The way was paved for revising some of our most basic notions of what organisms are and what their actions mean. But at the same time, a body of spurious conjecture appeared, produced by people unaware of or uninterested in the existing knowledge about control systems (or else, aware of it in a peripheral way but convinced that its essence could be captured in a few cleverly-stated rules of thumb).

The most unfortunate aspect of the conjectures was that they were all grounded in the old cause-effect conception of behavior; the radical switch of viewpoint actually required was simply too fundamental to be accomplished without basic knowledge of the principles of control. Those principles, never firmly grasped, soon faded from view. The leaders of cybernetics began, without knowing they were doing so, misleading. One person, who later became a president of the American Society for Cybernetics, announced that he had always considered purposive behavior to be adequately modeled by a drop of water sliding down an inclined plane under the guiding influence of gravity. Another famous cyberneticist, summing up what had been learned during the Macy Conferences, announced that no closed-loop system could avoid runaway oscillations if the feedback factor were greater than unity. Still another proposed that the basic principle of regulation amounted to sensing the cause of a disturbance, and converting that information into a precisely-computed compensatory effect on the controlled variable. Many others proclaimed that control was based on sensing errors, as if error could be observed in the outside world. Others said that control amounted to calculating the precise program of action that would correct an error, and then executing it. Many others said that incoming sensory information "guides" behavior, and another very popular notion was that control consists of limit cycles or alternating sequences of error and corrective actions. Every mistake that could be made was made, authoritatively.

While these views missed the main point, some of them nevertheless contained a grain of truth, and served to keep alive the flavor, if not the substance, of control theory. The basic phenomenon of circular causality continued to be recognized, and its implications expanded. Furthermore, the idea that organisms are active agents was crucial in encouraging explorations of brain models, computer models for the most part, and in leading to the development of new philosophical stances, all pertinent to control theory. The weakness at the foundations was not fatal; at least the implications of control theory continued to be recognized, and continued to attract people who saw that this view made more sense than conventional ones, even if they could not defend it.

We now come to the real subject of this outline: the control-system model I am trying to introduce, or rather re-introduce, to cybernetics. It is not easy for cyberneticists to concede that there is something fundamental about their own discipline that they have missed, especially when the one who makes this claim seems to be an outsider. A certain amount of resistance, even hostility, is to be expected, and I assure the reader that I have already accepted it and discounted it. I have to do so, to remain consistent with the principles I believe to apply to human nature.

But something is demanded of cyberneticists, too; they must at least take under advisement the possibility of thinking the unthinkable. I ask no more than understanding of what I propose.

Cybernetic control theory

While I already knew a little about control theory at the time, my lifelong interest in applying it to human behavior began only after I read Wiener and Ashby in 1952. It seemed to me that they had uncovered a vastly important principle of behavior, new to the life sciences. Being unknown and feeling ignorant, I determined to learn more about control theory and its applications to behavior, so that some day I could enter those exalted halls of cybernetics with something to contribute.

This project began in 1953, in collaboration with a physicist, R. K. Clark. We were soon joined by a clinical psychologist, R. L. MacFarland, and began to learn control theory in depth, my role being that of an engineer/physicist who was designing and building control systems as part of the job of a medical physicist. Clark really made the whole project possible by finding us both a position at the V. A. Research Hospital in Chicago, where I worked as his assistant. MacFarland was the Chief Clinical Psychologist there, and made important contributions in translating our somewhat austere models into terms that conventional psychologists might conceivably understand.

Our first paper describing the control-system model was published in 1960, in the shadow of Miller, Galanter, and Pribram's book on the organization of behavior, where the TOTE unit acquired its unfortunate lease on life. I will not bore the reader with tales of the meager acceptance that greeted our publication: cyberneticists have had their own problems, for similar reasons, with the Establishment.

This brief review of my own history is by way of saying that my interest in control theory was originally inspired by cybernetics, and was always intended, at least as a background hope, for use by cyberneticists (as well as psychologists). I thought, for many years, that I was simply catching up.

Neither will I bore the reader by re-running the laborious process by which we arrived at the final model, after backing out of many blind alleys. I will pass over the ensuing years of intermittent discouragement, the regrouping that ended with my book in 1973, my subsequent tentative forays into the American Society for Cybernetics, and the rise of the Control Systems Group, that rumor of Visigoths poised on the borders of cybernetic civilization ready to plunder and rape and otherwise violate the comfortable ways of the ASC. None of these matters will be important if the basic concepts of this theory are clearly understood. We have all been through the wars. We are all on the same side. Let's get to it.

The nature of control

The first thing that must be understood is that control is something that a control system does, not something that is done to it. The second thing is that in a control system there is no "controller." Control is a phenomenon that arises when an active system, constructed in a specific way, interacts with its immediate environment. The third thing is that the relationship between control system and environment is not symmetrical. Even though each affects the other, only the control system controls. The word "environment" means here the passive physical environment that takes no action of its own, but behaves as it is made to behave by natural forces: the world of the physicist. The presence of other control systems is a complication we will take up later.

A control system senses its environment and acts on it. Sensing means representing, and representing, if it is to mean anything reasonable, means analogizing. A sensor responds to some specific aspect of its environment, some variable outside the sensor, by generating a signal that is a quantitative analogue of the state of the variable. Bear with me for now: this concept of representation will become more interesting.

Acting means generating some physical effect whose magnitude and direction depend smoothly on the magnitude and sign of a driving signal inside the control system. Again, bear with me: we are speaking of the foundations of more complex actions.

As explained earlier, the sensor signal representing the external variable is compared with an internal reference signal that is of the same physical nature as the sensor signal. The result is an error signal that is zero only when the sensor signal matches the reference signal.

The action of the system is driven by the error signal.

In order for control to appear, the parts of this system must act in specific ways. The sensor signal, for example, must vary over a range from minimum to maximum as the external variable goes through its whole possible range of change. This relationship establishes the range within which control is possible.

The action of the system must affect the external variable at least in the dimension that is sensed. If an action caused by a positive error signal changes the sensed variable in one direction, the action caused by a negative error signal must change the variable -- as sensed -- in the opposite direction.

The overall effect of these relationships must be that the action driven by either sign of error signal must tend to alter the external variable in the direction that makes the sensor signal come closer in magnitude to the reference signal, so that the error signal becomes smaller. This is the basic requirement for negative feedback.

These requirements give us the qualitative basis for control phenomena. But there is a critically important quantitative basis as well, which accounts for the asymmetry of control.

The error signal drives the output action. It makes a great deal of difference how much error is required to produce a given amount of action. The ratio of action to error is called the <u>error sensitivity</u> of the control system. The output function, the effector of the control system, not only converts from signal-units to physical-world units of effect, but it enormously increases the level of energy that is involved in all variations. The output function is a transducer, but it is also an <u>amplifier</u>.

The output action of the system is connected to the external variable through an environmental link. In this link the laws of thermodynamics prevail: no more comes out than went in. Between the action and the effect on the external variable there is usually some degree of <u>loss</u> of effect. There may be a change in energy level in passing from the external variable to its sensory representation, but if we normalize both variables to their total range of change, there is no amplification. Almost all of the amplification (that is not simply a change of units) that occurs in this control process occurs in the output function, in the conversion from error to action. Here thermodynamics means nothing: the system is supplied from outside with whatever amount of energy it expends. The books do not have to balance: this is a thermodynamically open system.

It is a peculiarity of control systems that causation often seems to reverse itself. If we compare two control systems with greatly different error sensitivities, our first guess might be that the system with the greater error sensitivity, all else being equal, would produce the greater amount of action. What actually happens is that the system with the greater error sensitivity contains the smaller error signal, and its action is essentially the same as what the other system produces. If you double the error sensitivity, the result is very nearly to halve the error signal, not to double the amount of action.

There is one last consideration that has nothing to do with the process of control itself, but which is one of the major reasons why control is necessary: disturbances. The external variable is affected not only by the system's action, but by the world in general. The temperature of a house is affected not only by the furnace's output, but by heat entering, leaving, or being generated by other sources in the building. The path of a car is affected not only by the driver's steering efforts, but by crosswinds, tilts and bumps in the road, soft tires, and misalignment of the wheels. A savings-account balance is affected not only by depositing and withdrawing money, but by service charges, computer errors, and crooked employees. Variables that organisms control are controlled because they will not spontaneously come to the states desired by the organisms, and even when brought to those states, will not stay there.

The physical environment is in a continuous state of variation, so much so that no specific action can have just one specific consequence. There can be no such thing as computing an action that will have a desired result, unless one has taken great pains to shield those results against all normal independent influences. That may be approximately possible in the laboratory, but it does not happen in normal environments.

Furthermore, as we are beginning to hear, the lawfulness of the physical world itself is largely illusory even discounting Heisenberg. Many natural phenomena are so sensitive to slight variations in initial conditions that even though we can prove, by backward reasoning, that they are lawful, we cannot establish initial conditions accurately enough to turn those deductions into reliable predictions. The behavior of higher organisms is clearly one of these phenomena. Behavior results from the application of muscle forces -- not very reproducible in themselves -- to the masses of the body. The result is not "movement" but acceleration. Even to turn an effort into a position requires a double time-integration, which vastly magnifies all force variations, and by greater and greater amounts as time progresses. And this does not begin to take into account the indirect effects of limb movements that, in order to produce the larger patterns of behavior, must be integrated again and again, all the while being subject to unpredictable disturbances. It is not necessary to invoke control theory to show that the old causal model of behavior is wrong: all we need do is look realistically at what is involved in making "the same behavior" occur twice in a row in a disturbance-prone and semi-chaotic universe.

If organisms simply behaved blindly, the consequences of their actions would be essentially unpredictable. The same action applied ten times in a row would have ten different consequences, in most cases radically different. The physical world, uncontrolled, drifts in a kind of gigantic Brownian movement, showing order on an intermediate time-scale but for the most part simply changing aimlessly. Control systems impose order on this aimless drift. The automobile, buffeted by winds, jolted by bumps, dragged by uneven friction, wearing out asymmetrically from one minute to the next, nevertheless clings to a path that deviates by no more than one or two feet from the right path in 100 miles. This regularity is wholly unnatural, and can be accounted for only by knowing that there is a control system at the steering wheel.

The fact that there is behavior at all shows us that there is control.

To grasp the behavior of a control system correctly, it is necessary to think of all parts of the system at once. Control is not a sequential process, but a process of continuously and energetically maintained equilibrium among all parts of the system and between the system and external influences. If a disturbance arises that tends to change the external variable being controlled, the system does not wait to act until the disturbance has finished its work. Instead, the action of the system begins to change the instant there is any deviation of the sensor signal from the reference signal. Because this action opposes the error, it also opposes the effect of the disturbance. As the disturbance increases and decreases, so does the action opposing it increase and decrease. The sensor signal, in this process, varies slightly away from the reference setting, but if the error sensitivity is reasonably high only a tiny amount of error is needed to keep the action balanced against the disturbance at all times. For all practical purposes the action prevents the disturbance from affecting the controlled variable.

You will notice that some familiar concepts customarily associated with control processes are missing here. The first missing factor is any ability of the control system to sense the <u>cause</u> of a disturbance of the external controlled variable. While a more complex system could sense the cause of the disturbance, doing so would not materially improve control. The control system responds only to deviations of its own sensor signal from the reference signal. Why there is a deviation, whether it is due to a single cause of disturbance or to the combined effects of a thousand independent causes all acting at once (the normal case), is irrelevant. All the control system needs to monitor is the controlled variable itself: if the controlled variable starts to depart from its correct state, the system acts directly on it to keep it where it belongs. There may be a few circumstances in which "feed-forward" would be advantageous, but it can never substitute for the basic process of control. I should add that pure compensation, in which only the state of the disturbance (not the controlled variable) is sensed and a compensating action is calculated and applied along with the effect of the disturbance, will not work at all in most circumstances. It may seem to work on paper, where we can represent variables by simple whole numbers and give the imaginary system knowledge of all disturbances acting (and of the links from each disturbance to the controlled variable), but in the real world it can't even come close to explaining what we observe.

Another missing factor is any provision inside the control system for computing the proper amount of output to correct a given error. The only thing approximating an output computation is the amplification of the error signal, the system's error sensitivity. In order to compute the right amount of output to produce a given effect on the controlled variable, the control system would need a great deal of information that its simple sensor signal does not carry. It would need to know the momentary properties of the physical link connecting its action to the controlled variable, and it would need to know what amount and direction of disturbance will be acting at the time when the output calculation is put into effect. To get the required information it would need a vast array of extra sensors and a very large computer programed with the laws of physics -- and the ability to predict future disturbances. Furthermore, it would need to know about its own properties, because the instant that the output computation began to have its effect, the input variable would change to a different state, making the computation obsolete. The concept of "computing the appropriate action" is not only superfluous, but amounts to a very poor design. In the real world, human beings often try to control complex events in this way, thinking that logically it has to work, but in fact such efforts usually prove fruitless, as witness the attempts of the Federal Reserve to regulate the economy by diddling interest rates.

Finally, also missing is the entire concept of a "controlled action." Control systems do not control their actions: they vary them. What they do control is the variable affected both by the action and by disturbances. And in the final analysis, what they really control is the sensor signal that represents the external variable. All the rest of the system functions to maintain the sensor signal in a match with the reference signal. The action of the system is determined at every moment by the nature of the feedback link to the controlled variable and by the amount and direction of net disturbance that is acting. If the action itself were controlled, the variable could not be stabilized against disturbance. If the driver of a car controlled the steering wheel instead of the position of the car, the car would go immediately into the ditch, because no one position of the steering wheel will keep the car on the road for very long.

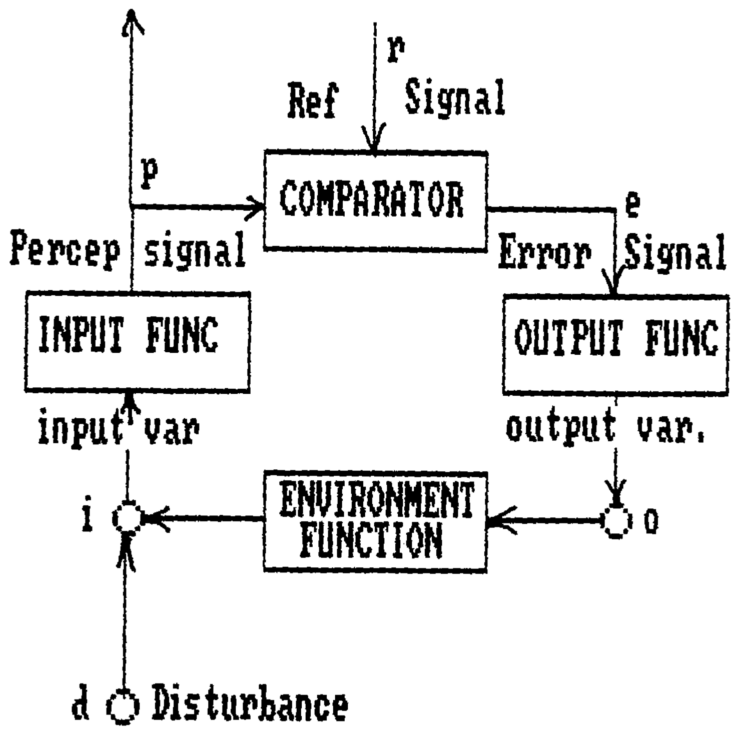

Fig. 1: Generic control-system diagram

Fig. 1 shows the basic relationships we have been talking about.

A hierarchy of control

What we have seen so far would probably be called a "homeostatic" system. We have a system that maintains a one-dimensional variable at a constant level matching a fixed reference signal. This system might behave very energetically as disturbances come and go, but the net result of its action would be a variable that is held constant.

By now, however, it should be clear that the control system's action focuses on maintaining its own sensor signal in a match with the reference signal. Nothing was said that specifies the setting of that reference signal, and nothing was said to limit the reference signal to a single fixed value.

If the reference signal varies in magnitude, the first effect will be to <u>create</u> error. Instead of the sensor signal departing from the reference signal, the reference signal departs from the sensor signal, but the result is precisely the same: an error signal that is highly amplified to produce action. The basic arrangement has not changed: the system will still be organized to alter the sensor signal in the direction that makes the error smaller. But now its action will have the effect of making the sensor signal change, rather than holding it constant.

In a well-designed control system, errors are never allowed to get very large. Consequently, when the reference signal changes, the output action will drive the controlled variable to change right along with the reference signal. This is the gun operator twiddling the control knob. Changing the reference signal is a way of changing the external controlled variable in a predetermined way -- namely, the way the reference signal changes. If the reference signal changes smoothly from a low value to a high value, so will the controlled variable change, quite without regard to any other physical influences acting on it. The control loop will automatically produce whatever fluctuations in action are required to make the controlled variable obey the reference signal rather than other influences.

So whatever is capable of manipulating the reference signal is also capable of manipulating a variable in the environment of the control system. The way that variable changes is determined by the cause of the reference-signal changes, and more important, ceases to be dependent on all the physical laws that would otherwise determine how it behaves. The control system has taken over that variable, cut it out of the normal flow of inanimate nature and made it behave as the control-system -- or as the manipulator of the reference signal -- wishes it to behave. The aimless drift that the variable would naturally exhibit is replaced by purposive change. Regularity has been imposed on Chaos.

Note that we still do not have purposive action. The actions of the system are still dictated by disturbances and by natural resistance of the variable to being changed. For any given state of the controlled variable, the action might be found anywhere within its possible range, depending on what else is doing something to the controlled variable, or trying to. Purpose can be seen only in the controlled variable itself -- in its variations that have been rendered immune to the normal forces affecting it. The purposiveness of a home thermostat is not to be seen in the furnace's turning on and off. It is to be seen in the steady temperature of the room where the sensor is located: 68 degrees in the daytime, and 62 degrees at night, when the little purposive computer lowers the reference signal for the temperature-control system. Rain or shine, summer or winter, the temperature stays at one or the other intended level. The furnace turns on and off as it must. Controlled variables, not actions, contain the evidence of purpose.

In the human body, at the lowest level of behavioral organization, there are something like 600 to 800 small control systems, each of which controls the sensed amount of strain in one tendon. The signal representing tendon strain is sent to the spinal cord, where it is compared (by subtraction) against a reference signal arriving from higher centers. The resulting error signal drives the muscle associated with the same tendon. These systems are small, but they are not weak: the range of strain that can be detected and controlled ranges from about a tenth of a gram up to something over 300 kilograms, in the system associated with a normal biceps muscle.

The reference signal that reaches the spinal comparator has been described regularly as a "command" signal, its function being to cause a specific amount of muscle contraction. But that is not how it works. The reference signal specifies how much signal is to be generated by the sensors that detect tendon strain. If disturbances alter that strain, the local control loop will automatically raise or lower the muscle tension to leave the net strain the same. It is the stretching of the tendon, not the contraction of the muscle, that is under control.

More specifically, it is the signal analogous to tendon strain that is controlled. In each case, this signal follows a branching path. One branch goes to the spinal comparator, as mentioned. The other branch continues upward, or inward, carrying a copy of the sensor signal in the direction from which the reference signal is coming. When everything is working properly, the upgoing copy of the sensor signal varies exactly as the descending reference signal varies. From the standpoint of the higher systems generating the reference signal, the effect is to control a sensation of effort simply by varying a signal standing for the amount of intended effort. The brain "wills" an effort by emitting a reference signal: immediately, that same amount of effort is experienced. The lag is imperceptible, amounting to no more than 20 milliseconds. It's no wonder that we have trouble separating the sense of willing an action from the experience of the action occurring. Paralysis, of course, makes the difference frighteningly obvious.

We have now created a class of control systems, the set of all effort-control systems. Everything that a human being does that could be called overt behavior is done by varying the reference signals reaching these systems. Everything. Whether a person is playing a piano concerto, painting the Mona Lisa, pressing the button that starts a war, making a lying speech to skeptical constituents, skating for an Olympic medal, or pounding on the keys of a word-processor, the acts involved are all accomplished by varying the reference signals reaching these 600 to 800 first-order control systems.

No system higher than the first order can act directly on the environment by generating physical forces. The actions of all higher systems consist entirely of generating outgoing neural signals. There are no moving parts in this system above the first level. There are only signals, and systems that receive, manipulate, and generate signals.

This is not the place to present 30 years of elaboration on this concept of levels of control. I will only try to sketch in the basic relationships that seem reasonable to propose. As far as I know, there is considerable neurological evidence in support of these suppositions, and nothing known to speak against them. But I am not pretending to be a brain researcher: I'm only trying to put together a feasible picture of an organization that has, within the bounds of what we know, a chance of actually existing. Perhaps these suggestions will raise some questions in the minds of real brain researchers. I'm far from the first to suspect control systems in the brain, but I don't believe that anyone else has approach the problem quite in this way (at least before I did). My little claim to fleeting fame.

Having isolated the first-order behavioral control systems, we now have a collection of incoming sensory signals, a subset of which is under control, and a collection of outgoing signals that become reference signals for the first-order systems. We can ignore the probability of cross-connections and other complications at this and other levels, in the interest of seeing the big picture first.

It's clearly possible now to think of a second level of control. At this second level, a control system would receive some set of first-order sensory signals (most of which come from receptors not involved in effort control), and would re-represent this set of signals by combining them in perceptual computing functions to create a new set of signals. Each second-order perceptual signal thus produced will represent some new type of invariant of the first-order world (every single-valued function of multiple variables generates some sort of invariant). I have reason to think, but will swallow the temptation to elaborate, that each new level actually represents a new type of variable in exactly the sense of Russell's Theory of Types.

Once an aspect of the first-order world has been represented as a one-dimensional second-order perceptual signal, we can quickly assemble a control system. We need a reference signal from still higher up, and a comparator to generate an error signal. And we need an output function that will amplify the error signal and send the result in the form of reference signals to <u>all the first-order control systems that can affect the second-order perceptual signal</u>. The effect may be direct, through pathways inside the body, or indirect, through pathways that include the external world. The effect may be achieved through altering the external world, or by altering the relationship of parts of the body to it, as when the eyes move.

How many second-order control systems might exist? A great many: a better question would be, how many can be active at the same time? Here there is a fundamental limit. The number of first-order control systems sets one limit on how many independent combinations of muscle tension can be produced at the same time. The number is large, but it is not infinite.

A second limit exists, set by the number of different functions, independent of each other, that are perceptually computed from the set of all first-order perceptual signals (at any one time). At most, 600 to 800 such signals might conceivably coexist, but in fact the likelihood of that many independent functions being discovered by the brain has to be very small. Let us just say that there is some number of independent dimensions of the first-order world that could be simultaneously computed, and that it must be considerably less than 600.

Why is this limit on numbers important? Because of a consideration left out of the discussion so far. Even just on anatomical evidence, we know that each spinal comparator neuron receives not just one reference signal, but in most cases hundreds of them. There can be only one net reference signal at a time for one first-order control system, but because the converging reference signals can have both positive and negative effects on the net setting, this net reference signal has to be considered as the weighted sum of the outputs of many higher-order systems. We can say "second-order" systems; there are arguments against reference signals skipping orders on the way

FIG. 2

A HIERARCHY OF CONTROL

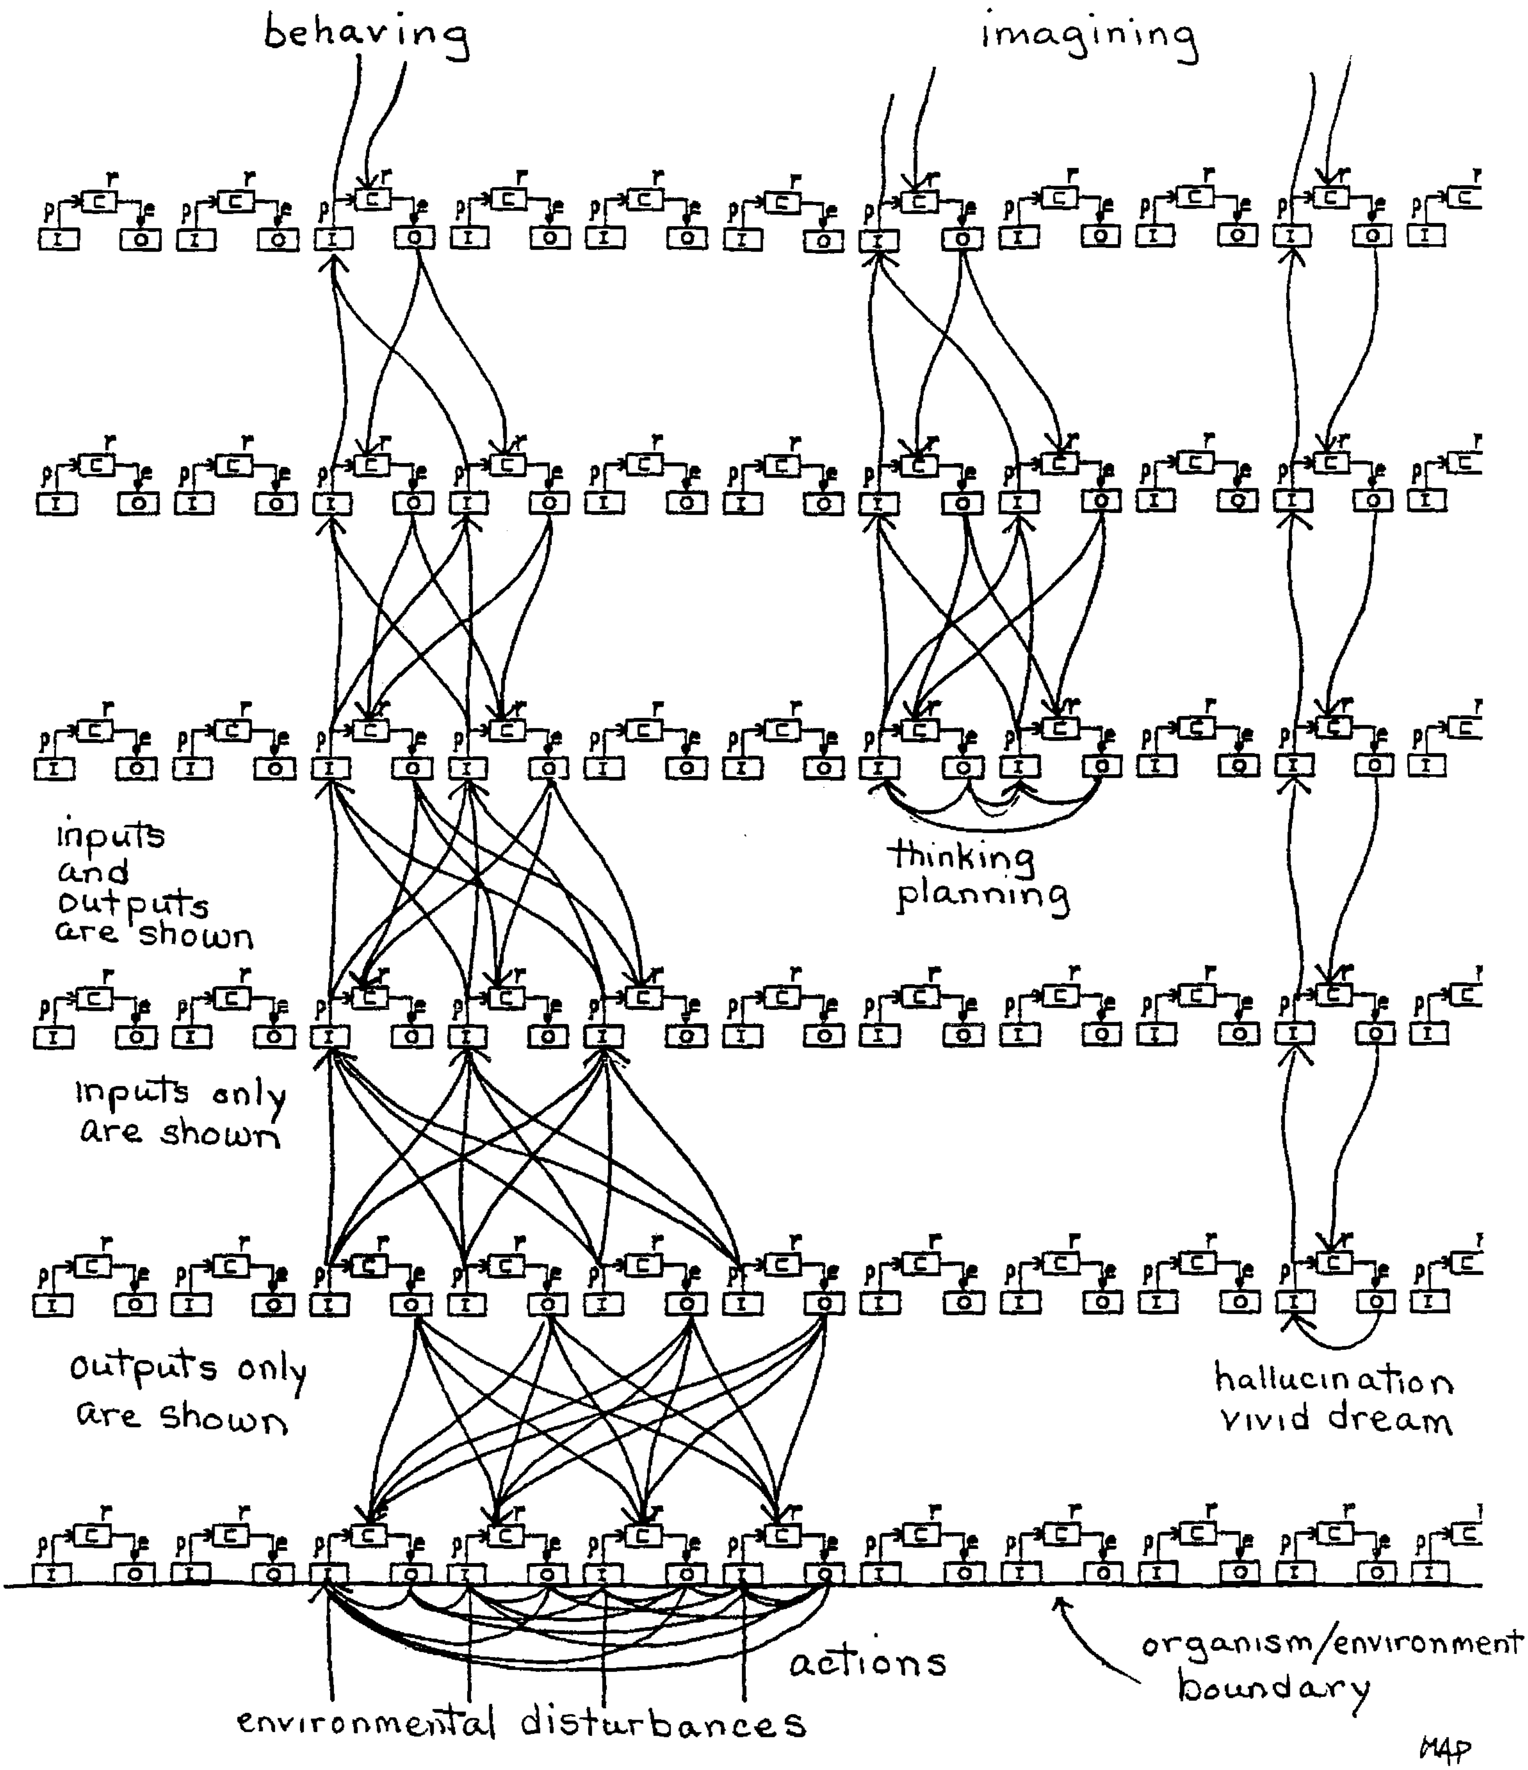

down in a control hierarchy (such signals would be treated as disturbances and canceled).

We thus have a picture in which some number of independently-acting second-order control systems act by sending multiple amplified copies of their error signals to many first-order control systems, specifically those whose actions can alter the second-order perceptual signal directly or indirectly. The second-order systems therefore share the use of overlapping sets of first-order systems. No one second-order system determines the net reference setting for any one first-order control system. The net reference setting for one first order system is always a compromise among the demands of all the second-order systems that affect it.

What's interesting about this arrangement is that it can actually work. The crucial part of this sharing of control is not the separation of output effects -- those are simply added together, with the appropriate sign to maintain negative feedback around each loop. What matters is that all the second-order input functions produce perceptual signals capable of independent variation: the _input_ functions must be linearly independent. Given these conditions, we have a well-known setup for the solution of large sets of simultaneous equations by analogue computation. Digital computers can be set up to do the same thing, far more slowly, using "methods of steep descent" and other arcanities. It is possible for many second-order control systems to maintain quite independent control of their own perceptual signals, despite having to act through a set of shared first-order control systems.

Fig. 2, thought up and drawn by Mary Powers and a handy program, shows a few of the arrangements possible in a large hierarchy of control systems. Of course only a few connections are shown, with some deliberately confined to input or output effects for clarity. In the middle and on the right are shown some short-circuit connections, in which the outgoing reference signals bend back to become inputs to the same systems, without involving lower systems or the environment. This is the "imagination connection" that enables us to think -- to envision the effect of doing things but without doing them. Above the level of the imagination connection, the perceptions are perfectly normal, except perhaps for the combinations in which they can occur. We have a sense "that" something is happening, without the lowest-level details to make it vivid or real.

This diagram has a vague resemblance to a real nervous system, which would become much stronger if at each level we stretched the connecting lines and clumped all the input functions together, and all the comparators and output functions together. Then we would have a realistic picture of the sensory nuclei, the motor nuclei, the upgoing and downgoing tracts, and the collaterals that run crosswise at every level in the brain.

At every level that may exist, we can expect the same sort of arrangement. Each new level of perception creates a new class of entities that can be controlled by varying reference signals at the next lower level. If you trace out any higher-order control system, you will see that the control loop always (except for imagination) involves effects in the external world. This permits us, as external observers organized in the same way, to discover the aspects of the shared world that are under control by another organism, even though those aspects be highly abstract. All that is required is that we learn to apply the same stages (or equivalent stages) of perceptual computation to the basic sensory input we are getting -- from what we presume to be a common environment. This is how we attack the problem of communication under the control system model.

There are obvious questions about the highest level of control, and obvious answers that I will not spend time on here. I hope it is suspected that far more could be said about this hierarchy than I have said. Most people take about two years to get the full picture of this model even when they're trying; we won't get that far in one paper.

There are two main subjects that still really need discussion -- I will abandon the notion of getting into the biochemical control systems and evolution, because this is already a very dense and long presentation. One subject is epistemology, which takes on a particularly important significance in this model, and the other is reorganization, the key to the development of an adult control hierarchy and also, although I won't go this far, the route to understanding physiological growth and the evolution of species. I want to show how the control model bears on two subjects that have become centrally important in cybernetics over the past ten or fifteen years. And I would like to say at least a word or two, at the end, about the picture of human existence and aspirations that control theory can give us.

The view from inside

To this point we have been looking at control systems and the world with which they interact from a viewpoint that is convenient but artificial. From where we stand, or float, we can see the physical environment surrounding the body, the brain and nervous system inside the body, and the signals spreading through millions of channels in the brain. Our X-ray eyes penetrate the skin to reveal muscles contracting and relaxing, putting stresses on tendons that give way slightly, exciting the little sensory nerves embedded in them. In the outside world we can see objects, but also the forces and influences that connected them together. When we put matters that way, it has to be clear that this entire picture is imaginary. It is, in short, a model: a model of a brain in a model of a world.

Here is a simple question: according to this model, where is the model? If you look at Fig. 2, you will see those imagination connections that allow higher systems in the brain to generate perceptual signals for themselves without causing them in the normal way by acting on the external world. The model says that this imaginary picture of the brain and the external world exists in the brain, and is created inside the brain. My brain. Perhaps, a little bit, your brain too.

In particular, the model implies that all these things we experience, whether in imagination or "really" (there is no remarkable difference), reside in the upgoing perceptual pathways. This leads me to make a proposal for which there can be, in the nature of things, no direct evidence, but that does make a lot of things fall into place rather neatly, It is this: the objects of experience of any kind exist in the form of perceptual signals continually rising through the brain.

This proposal in no way pins down who, what, or where the perceiver is, the noticer, the observer. It concerns only _that_ which is observed. The _objects of observation_, I am proposing, are neural perceptual signals in the brain.

If you were to spend a few decades systematically and skeptically examining the real solid three-dimensional physical world that you see, feel, hear, touch, and taste, I claim that you would find it to consist of a number of types of experience. From simplest to most complex, I claim that these types can be named roughly this way: intensity, sensation, configuration, transition, event, relationship, category, sequence, program, principle, and system. The words need some elaboration to make their intended meanings clear, but you get the flavor.

These types of experience have an interesting relationship to each other. The ones farther along in the list -- "higher" -- depend for their existence on the existence of types lower in the list (I do many things backward: my list goes from bottom to top, and I write it left to right). Furthermore, if you want to change a particular experience of a given type, you will find it necessary to change experiences of lower types. Those relationships, however, are not reciprocal: a lower type of experience does not depend on a higher one, and can be changed without changing a higher one. As we go up the list, the relationships between types are the relationships between successive stages of invariants, each stage abstracted from the previous one by a new rule, as in Russell's Theory of Types.

This, not by accident, is exactly the structure of the perceptual part of the control hierarchy in Fig. 2. It also seems to be the structure of the perceiving functions at various levels in the brain, give or take some topological transformations, and allowing for the fact that models are always neater than nature.

But this is not just a structure of perceptual functions: it is a structure of control systems. A control system at any level acts on a world consisting of lower-level control systems, the means of acting being to send varying reference signals to some of the lower systems. These control actions ultimately result in the lowest-level systems doing things to the outside world, and thus to the lowest level of perceptual input signals, the intensity signals. The first-order signals are abstracted to become second-order signals, and so on until we reach the system we began with, where the effect of that system's action is to maintain its own perceptual signal in a match with the reference signal it is being given from above.

But here we are floating in space again, while the point, if I haven't mentioned it, is to see how it is to _be_ a system like the one in Fig. 2.

When you _are_ a system like this, you find that by acting you can alter the world you perceive. When you learn its rules well enough, you can learn how to make many of those perceptions come to states you have experienced before and liked, or to stay away from states you have experienced before and didn't like. When you see a flower, you can move it to your nose or your nose to it, use your diaphragm to pull air in, and experience a scent that you judge as pleasant. If it's a pretty flower that ought to have a scent but doesn't, you can supply a scent in imagination. You can supply a scent at a low level, like a hallucination, or at a higher level, like an impression of niceness.

Whatever you do alters your perceptions: that's how you know you're doing something. You perceive your own efforts and their immediate effects such as skin pressure; you perceive effects of those efforts as objects change their (visible) positions, orientations, and velocities. You use your ability to control your limbs as ways of controlling other objects; you use control of other objects to create movements and events in relationship to other movements and events; you control movements and events to maintain certain categories of experience in the states you intend; you maintain these categories in sequences that constitute progressions of familiar kinds; you adjust these progressions according to rational decisions, choices, tests, and symbolic equivalences; you carry out rational processes in support of general principles that you defend, and you maintain those principles as a way of sustaining whole systems within which you live and experience and which you try to maintain, systems like a self, a science, a society, a culture, a world, a universe.

All of our actions, according to this control model, are part of a process of controlling perceptions. To understand this idea properly, you have to abandon all the meanings the word control has accumulated, meanings that represent, mostly, bad guesses as to what is going on. Controlling does not feel like trying: it is lack of control that feels like trying. Controlling is just doing. You don't have to "try" to look at something -- you just look. Your oculomotor control systems snap the object you want to look at to the center of your visual field, and there is no sense of trying at all. You don't "try" to write your name; you just write it. By far the majority of control processes that go on at all these levels are skillful, swift as thought, stable, and seemingly effortless. You form an idea of what is to happen and it happens at the same time. You just do it. There is no process of laboriously selecting some intended perception, figuring out a way to get closer to it, and then painfully working your way toward zero error. That only happens when you don't know what you're doing. Mostly our perceptions track our intentions with no perceptible lag. That's why, sometimes, they're hard to tell apart.

Of course at the higher levels of control, particularly the cognitive levels, things happen more slowly unless we're imagining. There has to be time for all the lower-level systems to bring their perceptions to whatever the momentary net reference signal specifies. The lowest level systems have a lag of perhaps 50 milliseconds, whereas the highest ones, operating at their fastest, may lag as much as half a second or a second. Some control processes may take much longer than that: I'm involved in one that has been going on for -- let's see, 1988 minus 1953 plus one -- thirty-six years. Of course a wise person doesn't tolerate protracted error; he or she redefines the controlled quantity so it can in fact be held at its reference level without large error. I'm making progress, that's more like it.

To say that behavior exists in order to control perceptions is not to say that all perceptions are under control. Much that we can see happening around us happens without benefit of our advice or effort. But we do come to "expect" the world to be a certain way; that is, even without specifically intending to do so, we set up reference signals against which we compare perceptions even when we have no direct way of affecting them: an inner model of how the world should be. As long as the world matches these expectations, we experience no error and go about our affairs normally. But just let the sun rise in the West one morning, and see how much error you would experience, and how frantically you would start to act to try to do something about this gross mistake.

You can see that this model implies an epistemology. If what we experience consists entirely of perceptual signals in the brain, it follows that we do not experience the causes of these signals. The causes lie outside, according to this model, beyond our sensory endings where we, the observers who experience only perceptual signals, have no contact. Our motor efforts disappear into that world, and we know nothing of what they do to the world until the effects return to cause changes in our intensity signals. What we can know of that external world consists only of what we can sense, and what we can imagine. Sensing and imagining occur inside, not outside, the brain.

How would a brain organized as this model is organized ever know that an external world, other than the apparent one, actually exists? There are at least two kinds of evidence available. One kind is the fact that in order to bring any perception under control , the brain must discover how to manipulate reference signals to have the required effects. This can be done only by trial and error, with perhaps a smidgin of genetic help. The relationship between what must be done and the result that it has constitutes a model of some "property" of the external world. The fact that stable properties can be found is evidence that there is something lawful and stable outside the boundaries of experience. In more formal surroundings, this is called "scientific method" (except in the behavioral sciences, where scientific method means assuming a cause-effect model and then throwing out all data that doesn't conform to it).

The second line of evidence is found in the very fact that control is necessary. The world will not usually meet our needs, desires, or expectations unless we do something do it, and even when we have learned how to maintain the world as we wish it to be (in the respects we can affect) we find that we still have to vary our actions in order to maintain it in any particular state. In other words, those perceptions are subject to influences other than our own actions. Disturbances. The driver of a car can deduce the direction and strength of a crosswind that he cannot sense in any way, simply by observing how he is holding the steering wheel. He quite automatically varies the position of the steering wheel in the way required to keep the scene in the windshield constant, showing that the car is in the right position on the road. He has no preference for wheel position. Thus he can "see" the crosswind, deduce it from his own control actions, without any other way of sensing it. He could, of course, be wrong: there could be something horribly wrong with the car.

That's really a third line of evidence: we can be wrong. We can go through half a lifetime or more thinking we have really got something nailed down, have full control and a competent model of what is happening, only to have some trifling incident turn our whole idea upside down, utterly destroying, for a while, our confidence in our ability to know anything. Such an experience, however, should give us more confidence, not less. What should make us lose confidence is finding that we can no longer detect the mistakes that tell us we can still, somehow, be in contact with reality.

This is certainly not a philosopher's approach to epistemology; it's a purely practical approach. I think that practicality, pragmatism of the right sort, is the key idea here. Knowing that it's all perception, we will think in new ways about most of our own experiences and actions. But will we then give up making models, just because we know they are "only" models? That would be foolish, because then we would be giving up the basis for giving up models, wouldn't we? I think the best course is to admit that what we call knowledge consists entirely of models, models of a body, of a brain, of a physical and chemical reality, of a society, of everything. Rather than giving up models, we should become conscious of the process of making models.

If we know we're making models, we won't go around telling people that they are wrong for trying out different models, or that they are right even if their models are sloppily constructed and unconnected with any other models. We should be looking to make all of our models consistent with each other, and worrying very seriously when they are not, and being fussy about what we will accept as a model. We ought to test the hell out of our models, because if they don't behave the way our experiences behave, they are worse than sloppy: they're delusory. They're useless. They're dangerous.

Of course when we know we're making models, we can be free to try out any ideas we please, as long as we realize that we're playing what is in the end a serious game. We are trapped in here, folks, and our very survival depends on making models that in some way reflect the regularities of the real universe that is right out there, but that we can know only approximately and only by way of models. Fun and games make life interesting, but somebody has to take out the garbage. But it's not that bad. Making models is really _fun_. Hello?

One last point before we leave this subject barely touched. I have made the claim that our experiences of the world fall into eleven types (more or less). Does this mean that the real universe is organized that way? I think my answer would be pretty obvious: of course not, although we can conjecture that there is some reason for these particular types to have evolved (the evolution-model). Basically, the types of perceptions are determined completely by the types of perceptual functions that are applied to each level of signal, and it is highly probable that each person organizes perceptions, within each type, quite differently. But there is a miracle going on that anyone interested in epistemology should acknowledge.

The miracle is that we can talk together at all about anything. Everything enters our nervous systems at the lowest level, becoming available to the brain first as a huge collection of identical intensity-signals. It takes many layers of information-processing before those intensities can be turned into the perception of a sentence, and more yet before a string of grammatically and syntactically ordered words can be used to evoke a non-verbal experience, the perceptual meaning of the sentence. I must turn my meaning into a sentence, and utter it, and you must turn the sound-intensities back into a sentence and the sentence into a meaning before anything resembling what is in my awareness springs up in your awareness. So how do we ever come to believe that the meaning you get is the one I intended?

Very often it's not the same. We only think it's the same, and sometimes fatally, assume it's the same. Finding out if it is the same is basically impossible, but even reaching some level of confidence in the sameness requires a long process of back-and-forthing, of cross-checking, of "If I understand you correctly, then when I do this, you'll do that."

Yet, look at this: over ten thousand words so far, and I still have some hope that you are with me. What you have made out of all these words, I will know only when you do or say something relevant to them -- as I intend their meanings. Epistemology is a very faint echo of the real problem, which is communication.

With that, let us pass on to the final topic.

Reorganization

I'm going to give short shrift to this subject partly because my endurance in sustaining this long narrative is beginning to wear down about as far as the reader's must be. This is a critically important subject; unfortunately, I don't know much about it, and can speak only in generalities. This is one place where I really wish I were a good mathematician.

The idea of reorganization is an essential part of this model, and has been since its beginnings. It was suggested -- laid out pretty completely -- by W. Ross Ashby in his notion of "ultrastability," and independently by Donald T. Campbell as "blind variation and selective retention.' The basic idea is simple, and older than either Ashby or Campbell.

There are many forms of learning, but the most fundamental is learning something for which there is and could be no basis in prior experience. This is the kind of learning that has to take place when you grab the knob of an unmarked door and try to open it. With no hints available, the door might require either a pull or a push: nothing in nature says it has to be either way. So what do you do to figure out how this door opens? You don't figure it out. You pull. Or you push. Whichever comes to mind first. If the door doesn't open, you have the information you need: do the opposite. If it opens, you also have the required information: don't change the way you did it. But before you could get that information, to select the right move out of the possible moves, you had to try something for no reason.

This is what I assume to be the basic principle of reorganization, which I could not put any better than Campbell did. Act at random, and select future actions on the basis of the consequences.

Another way of putting this is a little more systematic, and suggests at least some sort of organized system at work. Suppose we have a reorganizing system that is capable of acting on another system (of which it actually could be a part) to change the organization of that system. In this case I don't mean the organization of the behavior of the target system, but the very structure of that system, the physical connections in it. Changing the structure will, of course, change the behavior, but the reorganizing system doesn't act on the behavior directly. It acts on the behaving system. That's how Ashby's ultrastable homeostat works. It doesn't inject signals into the homeostat: it switches connections.

The reorganizing system must not only be able to alter physical connections in the target system, but it must know when to stop altering those connections. This is the "selective retention" part. Each change in the structure of the behaving system will alter the way that system interacts with its environment. The change in interaction will have many consequences, most of which, probably, are irrelevant to the system as a whole. Some of these changes, however, will have indirect effects on the welfare of the system itself, including the reorganizing system. These indirect effects are the basis for selection, and therefore the basis for starting and stopping the process of reorganization.

Selection necessarily implies a selection criterion. Some indirect effects of behavior are "good" and some are "bad", or at least "not good enough." But this reorganizing system has to be dumb. It has to work even when the system it is working on has only the barest suggestion of competence in it. It has to work without any theory, without any knowledge of an external world, without any memories of prior experience (from this lifetime, anyway).

So this system has to be told, somehow, what is good or not good enough, and perhaps even too good. It has to be given reference signals from somewhere. For lack of a better idea, I'll say DNA.

Furthermore, these reference signals have to have highly specific meanings. It won't do to posit a reference signal that says "survive!' How could a dumb reorganizing system with no linguistic capabilities know what "survive" means? It won't do to say "organize behavior." There isn't any behavior to organize at first. No, these basic reference signals have to be expressed in much more concrete terms that have direct meaning to the reorganizing system. They have to say things like "this much blood sugar" or "this body temperature" or "this carbon dioxide concentration in the blood." Of course they might also say more interesting things, like "no more than this amount of total error signal in the brain," or even "this pretty pattern of forms in your vision." We mustn't underestimate the power of a billion years of evolution. The selection criteria that make reorganization work as it does might prove to be extremely sophisticated.

But we know one thing they will not be: intelligent. Intelligence is something that gradually forms as the brain becomes organized into a hierarchy of perception and control. Intelligence is the <u>product</u>, not the cause, of reorganization. The intellectual skills found in the fully-formed adult control hierarchy are not available before it has been built. The reorganizing system has to work from the very beginning of life, so it can't take advantage of what it has not yet brought into being.

The reference signals -- let's call them "intrinsic" reference signals to distinguish them from the kind in the acquired hierarchy of control -- can have no effect by themselves; they are only specifications. The reorganizing system has to be able to <u>sense</u> the states of the variables that relate to the reference signals. And the sensed states have to be compared with the reference signals; the reorganizing system has to contain comparators, one for each intrinsic variable. Ashby called these intrinsic variables "critical variables." He saw the reference states as upper and lower limits, while I see them as single target values, but that's the sort of difference we might hope to settle through experiments, and isn't important here.

So we arrive at the idea that the reorganizing system is really a control system. It is, however, a very peculiar sort of control system, in that its output actions are random. It does not act "against" error: it continues to act until the error disappears. The error, of course, is simply the total absolute difference between the sensed intrinsic state and the set of all corresponding intrinsic reference signals. Ashby didn't spell out the perceptual functions or the reference signals in his ultrastable homeostat, but he did build them into it, perhaps without even realizing exactly what he had designed.

One helpful notion is the idea of <u>rate of reorganization</u>, which would be measured simply as so many changes per second, or hour. If there is a lot of intrinsic error, the reorganization rate will be high. As intrinsic error falls, assuming it does, the rate of reorganization will slow, until finally when the

Control theory is above all a theory of living autonomous systems. Living systems are all control systems, the only natural ones, and the essence of their lives is to control what happens to them, rather than leaving their fates to wind, tide, erosion, and entropy. But the human control systems that concern us most are also very new control systems, largely ignorant of their own nature and prone to treat other living systems, including human ones, as little more than objects to be moved or disturbances to be overcome. Indeed there have been times in human history when many people saw inanimate nature as full of purposive control systems, and human beings as only passive victims of nature's intentions.

It is not easy for control systems, human beings, to live together. Even when they attempt to cooperate, they end up pitted against each other over minor differences in perception or goal. They just can't help trying to keep their own errors corrected. To be with others one has to learn deliberately to loosen the control, to lay back, to tolerate error, to be a little less skillful. To expect less, perhaps, of group efforts than of individual ones, but to value them, perhaps, more. To let reorganization ease the strain. To realize how isolated we are; how miraculous it is that we have any contact at all, mind to mind. To appreciate the vast sea of mystery that fills the space between us, through which we would have great trouble steering without the touch of other human hands on the helm, the surprise of other human thoughts about the course.

Draft prepared for "TEXTS IN CYBERNETIC THEORY", *Special Fall Conference of ASC*, 1988

An Exposition of Radical Constructivism *

Ernst von Glasersfeld

SRRI, Hasbrouck Laboratory
University of Massachusetts
AMHERST, MA 01002

Michel de Montaigne is often listed among the sceptics. This is a little misleading because he actually used his outstanding wit and erudition to defend the realm of religious faith against the threat of the *Pyrrhoniens*. These "Pyrrhonists" were a subversive group of thinkers who had rediscovered Sextus Empiricus and his account of Pyrrho, the father of scepticism in the Hellenic world. Montaigne merely cut down to size the efforts of human reason. He put it as concisely as one can:

La peste de l'homme, c'est l'opinion de savoir.[1]

The translation that seems the most adequate to me would be:

Mankind's plague is the conceit of knowing.

Radical constructivism is an effort to eliminate that conceit. It does not deny the possibility of knowing, but it strives to show that knowledge is not the commodity the tradition of Western philosophy would have us believe. Indeed, constructivism is a theory of active knowing, not a conventional epistemology that treats knowledge as an embodiment of Truth that reflects the world "in itself", independent of the knower. The basic principles of radical constructivism are:

---

* This text is a composite and contains pieces from several papers, among which: "Cognition, Construction of Knowledge, and Teaching", to be published in *Synthese*, 1989; "The Reluctance to Change a Way of Thinking", to be published in *The Irish Journal of Psychology*, 1988; "Environment and Communication", prepared for the Action Group 1, *International Conference on Mathematics Education (ICME-6)*, Budapest, August 1988.

1 - a) Knowledge is not passively received either through the senses or by way of communication;

- b) knowledge is actively built up by the cognizing subject.

2 - a) The function of cognition is adaptive (in the biological sense of the term);

- b) cognition serves the subject's organization of the experiential world, not the discovery of an objective ontological reality.

To adopt these principles means to relinquish the mainstays of an inveterate conceptual network. It means getting out of habitual pathways and reconceptualizing a different rational view of the world. In short, it involves a good deal of thinking and, as Betrand Russell once said, people would rather die than think, and they do.

Knowledge and "Reality"

One of the differences between the vocational practice of theology and religion, on the one hand, and the lay practice of science and philosophy, on the other, we were told in school, is that the first is founded on a dogma that is held to be absolute and immutable because it stems from divine revelation; all lay practice, instead, is tentative because it develops theories that are always open to refutation by new ways fo thinking, new findings, or novel experiments with things or thoughts. Scientists and philosophers, therefore, are supposed to be open-minded and to welcome the solution of stubborn problems, even if the new solutions entail a change of ideas and the demise of concepts that seemed well established in the past. Montaigne and some of his contemporaries had a very clear view of this dichotomy. They were sceptics with regard to the rational knowledge of science and believers with regard to the traditional tenets of religion.

A look at the subsequent history of ideas, however, quickly shows that scientists and philosophers do not always live up to this ideal open-mindedness. The concepts and methods they have grown up with frequently seem to be as unshakable as any matter of religious faith and the perpetrators of innovation tend to be treated as heretics. This happened to Darwin and his theory of evolution, to Einstein when he first published on relativity, and it happened to Alfred Wegener when he suggested the idea of continental drift. In these instances the break with tradition advocated by the new theory was unmistakable and, consequently, triggered violent indignation on the part of those who were anxious to maintain the established dogma. The new theories won eventually, because they enabled scientists to do things they were not able to do before and to cover a larger area of experience with fewer assumptions.

New ideas in philosophy do not often gain decisive victories. I would not be so presumptuous as to offer an explanation why it is apparently so easy to live with profound epistemological contradictions. Being a little cynical in that regard, I feel that the quip I earlier quoted from Bertrand Russell says something to the point. Besides, there is a German saying (from Wilhem Busch, I believe) which, to me, seems relevant: *"Der liebe Gott muss immer ziehn, dem Teufel fällt's von selber zu."* (The good Lord must forever pull, but to the devil things fall quite by themselves.) If one replaces the ethical connotation with one of straight thinking, it fits our situation well: It seems quite natural that philosophical problems will be shelved when there are cherries to be picked or enemies to be fought.

Be this as it may, it is an historical fact that the pre-Socratics already saw very clearly what was to remain the major problem of Western philosophy for the next 2500 years: If one assumes, (*a*) that a fully structured world exists "out there", independent of any experiencing subject, and (*b*) that the cognizing subject has the task of finding out what that world is "really" like, one hangs the millstone of an irreducible paradox around one's neck. Whatever the subject perceives or conceives will necessarily be the result of the subject's ways and means of perceiving and conceiving--and there is no way ever to compare these results with what there was in the first place. The sceptics have not ceased to reiterate this, but it has not deterred philosphers from trying to find a way around the impasse. The urge to persevere in the quest for that unobtainable "objective knowledge" seems almost ineradicable. (And therefore, some philosophers concluded, objective knowledge must be obtainable after all.)

The trouble was (and is) that the sceptics' arguments have always focused on the negative. By reiterating that true knowledge of the objective world is impossible, they have helped to perpetuate the idea that knowledge, in order to be any good, would *have* to be about the objective world. This idea is at the very core of the Western epistemological tradition and the occasional dissidents, who tried to get away from it, have had virtually no effect.

The last three decades, however, have manifested symptoms that may indicate a change. It is certainly not the first time that scientific developments are having an influence on the professional thinking of philosophers, but I believe it *is* the first time that scientists are raising serious questions about the kind of epistemology philosophers have been defending. The disruption shows itself in the discipline that has become known as the Philosophy of Science and awareness of trouble was spread to a much wider public by Kuhn's *The Structure of Scientific Revolutions*. There, undisguised and for everyone to read, was the explicit statement that

> ... research in parts of philosophy, psychology, linguistics, and even art history, all converge to suggest that the traditional epistemological paradigm is somehow askew. That failure to fit is also made increasingly apparent by the historical study of science ... . None of these crisis-promoting subjects has yet produced a viable alternate to the traditional epistemological paradigm, but they do begin to suggest

> what some of that paradigm's characteristics will be. (Kuhn, 1970, p.121)

While the troubles of the "traditional epistemological paradigm" have shown no sign of subsiding in the years since Kuhn's publication, one could not honestly say that any substitute has been generally accepted. In most highschools and Universities teaching continues as though nothing had happened and the quest for immutable objective Truths were as promising as ever. For some of us, however, a different view of knowledge ***has*** emerged, not as a new invention but rather as the result of pursuing suggestions made by much earlier dissidents. This view differs from the old one in that it deliberately discards the notion that knowledge could or should be a representation of an observer-independent world-in-itself. It replaces it with the demand that the conceptual constructs we call knowledge be ***viable*** in the world as the knowing subject experiences it. (This is, in fact, quite similar to what the pragmatists have been saying.)

Ludwig Fleck, whose monograph of 1935 Kuhn acknowledged as a forerunner, wrote an earlier article in 1929 that went virtually unnoticed although it already contained much that presages what the Young Turks have been proposing in recent years:

> The content of our knowledge must be considered the free creation of our culture. It resembles a traditional myth (Fleck 1929, p. 425).
>
> Every thinking individual, insofar as it is a member of some society, has its own reality according to which and in which it lives (p.426).
>
> Not only the ways and means of problem solutions are subject to the scientific style, but also, and to an even greater extent, the choice of problems (p. 427).

Vico - The First Constructivist

The notion of cognitive "creation" or, as I prefer to say, ***construction***, was adopted in our century by Mark Baldwin and then extensively elaborated by Jean Piaget. Piaget's constuctivist theory of cognitive development and cognition, to which I shall return later, had, unbeknownst to him, a striking forerunner in the Neapolitan philosopher Giambattista Vico. Vico's epistemological treatise (1710) was written in Latin and remained almost unknown. Yet no present-day constructivist can afford to ignore it, because the way Vico formulated certain key ideas and the way they were briefly discussed at the time is, if anything, more relevant today then it was then.

The anonymous critic who, in 1711, reviewed Vico's first exposition of a thoroughly constructivist epistemology expressed a minor and a major complaint. The first--with which any modern reader might agree--was that Vico's treatise is so full of novel ideas that a summary would turn out to be almost as long as the work itself (e.g., the introduction of developmental stages and the incommensurability of ideas at different historical or individual stages, the origin of conceptual certainty as a result of abstraction and formalization, the role of language in the shaping of concepts). The reviewer's second objection, however, is more relevant to my purpose here, because it clearly brings out the problem constructivists run into, from Vico's days right down to our own.

Vico's treatise *De antiquissima Italorum sapientia* (1710), the Venetian reviewer says, is likely to give the reader "an idea and a sample of the author's metaphysics rather than to prove it." By ***proof***, the 18th-century reviewer intended very much the same as so many writers seem to intend today, namely a solid demonstration that what is asserted is ***true*** of the "real" world. This conventional demand cannot be satisfied by Vico or any proponent of a radically constructivist theory of knowing: one cannot do the very thing one claims to be impossible. To request a demonstration of ***Truth*** from a radical constructivist shows a fundamental misunderstanding of the author's explicit intention to operate with a different conception of knowledge and of its relation to the "real" world.

One of Vico's basic ideas was that epistemic agents can ***know*** nothing but the cognitive structures they themselves have put together. He expressed this in many ways, and the most striking is perhaps: *"God is the artificer of Nature, man the god of artifacts."* Again and again he stresses that "to know" means ***to know how to make.*** He substantiates this by saying that one knows a thing only when one can tell what components it consists of. Consequently, God alone can know the ***real*** world, because He knows how and of what He has created it. In contrast, the human knower can know only what the human knower has constructed.

For constructivists, therefore, the word ***knowledge*** refers to a commodity that is radically different from the objective representation of an observer-independent world which the mainstream of the Western philosophical tradition has been looking for.[1] Instead, ***knowledge*** refers to conceptual structures that epistemic agents, given the range of present experience within their tradition of thought and language, consider ***viable.***

The most frequent objection to radical constructivism, at the beginning of the 18th century as well as now, takes the form of discarding it as a kind of solipsism. It is the main objection that George Berkeley had to contend with when he published his major epistemological work, *A treatise concerning the principles of human knowledge*, in 1710 (by a strange coincidence, it was the same year that Vico published his treatise at the other end of Europe). If one keeps Berkeley's title in mind, it will be clear that when he declares *"esse est percipi"* (to be is to be perceived), the *being* he is talking about is the only one the human knower can conceive of, and that is *being* in the world of experience, *being* constituted by the kind of permanence that results from invariants created by an experiencer's successful assimilation (I shall explain this term later in the context of Piaget's theory). But Berkeley's opponents, just as today's critics of constructivism, reacted as though he had been talking about the "world-in-itself" rather than about *the principles of human knowledge.*

Both Vico and Berkeley were concerned with the human activity of knowing. Both had strong ties with the religious dogma that claims an absolute, eternal order of the universe. Their way of reconciling their blatantly subjectivist theories of knowledge with the requirement of an immutable objective world were parallel and equally ingenious. For Berkeley the unity and permanence of ontological *existence* was assured by God's perception which, because God is considered omniscient, was ubiquitous and all-encompassing. Vico, instead, maintained that, while the human mind could know only what the human mind itself had constructed, God knew the world as it is, because He had created it.

Some Recent Elements

Radical constructivism is less imaginative and more pragmatic. It does not deny an ontological "reality"--it merely denies the human experiencer the possibility of acquiring a "true" representation of it. The human subject may meet that world only where a way of acting or a way of thinking fails to attain the desired goal--but in any such failure there is no way of deciding whether the lack of success is due to an insufficiency of the chosen approach or to an independent ontological obstacle. Warren McCulloch expressed it very simply: "To have proved a hypothesis false is indeed the peak of knowledge" (1965, p.154). What we call "knowledge", then, is the map of paths of action and thought which, at that moment in the course of our experience, have turned out to be *viable* for us. Such a limitation of the scope of human understanding is, of course, considered dangerous heresy by all who, in spite of the sceptics age-old warnings, still cling to the hope that human reason will sooner or later unravel the mystery of the universe.

Richard Rorty, in his Introduction to *Consequences of Pragmatism*, announces this shift of focus in terms that fit the constructivist's position just as well as the pragmatist's:

> He (the pragmatist) drops the notion of truth as correspondence with reality altogether, and says that modern science does not enable us to cope because it corresponds, it just enables us to cope. (Rorty 1982, p.XVII)

Constructivism *is* a form of pragmatism and shares with it the attitude towards knowledge and truth; and no less than pragmatism does it go against "the common urge to escape the vocabulary and practices of one's own time and find something ahistorical and necessary to cling to" (Rorty 1982, p. 165).

When the anonymous reviewer complained that Vico did not ***prove*** his thesis, he in fact reproached Vico for not having claimed for his "metaphysics" (which was actually a theory of knowing) the correspondence with an ahistorical ontic world ***as God might know it.*** But this notion of correspondence was precisely what Vico--like the pragmatists--intended to drop.

Present-day constructivists, however, if pressed for corroboration rather than proof in the traditional sense, have an advantage over Vico. They can claim ***compatibility*** with scientific models that enable us to "cope" remarkably well in specific areas of experience. For instance, one might cite the neurophysiology of the brain and quote Hebb's:

> At a certain level of physiological analysis there is no reality but the firing of single neurons (Hebb 1958, p. 461).

This is complemented by von Foerster's (1970) observation that all sensory receptors (i.e. visual, auditory, tactual, etc.) send physically indistinguishable "responses" to the cortex and that, therefore, the "sensory modalitities" can be distinguished only by keeping track of the part of the body from which the responses come, and ***not*** on the basis of "environmental features". Such statements make clear that contemporary neurophysiological models may be compatible with a constructivist theory of knowing but can in no way be integrated with the notion of transduction of "information" from the environment that any realist epistemology demands.

Cognition as an Adaptive Function

Constructivism differs from pragmatism in its predominant interest in *how* the knowledge that "enables us to cope" is arrived at. The work of Jean Piaget, the most prolific constructivist in our century, can be interpreted as one long struggle to design a model of the generation of viable knowledge. In spite of the fact that Piaget has reiterated innumerable times (cf. 1967a, pp.210ff) that, from his perspective, cognition must be considered an ***adaptive function***, most of his critics argue against him as though he were concerned with the traditional quest for "true" knowledge, i.e., knowledge that could be said to correspond to an ontological reality.

This misinterpretation is to some extent due to a misconception about adaptation. The technical sense of the term that Piaget intended comes from the theory of evolution. In that context, ***adaptation*** refers to a state of organisms or species that is characterized by their ability to survive in a given environment. Because the word is often used as a verb (e.g. this or that species *has* adapted to such and such an environment), the impression has been given that adaptation is an activity of organisms. This is quite misleading. In phylogeny no organism can actively modify its genome and generate characteristics to suit a changed environment. According to the theory of evolution, the modification of genes is always an accident. Indeed, it is these accidental modifications that generate the variations on which natural selection can operate. And nature does ***not***--as even Darwin occasionally slipped into saying (Pittendrigh 1958, p.397)--select "the fittest", it merely lets live those that have the characteristics necessary to cope with their environment and lets die all that have not.

This interpretation of the theory of evolution and its vocabulary is crucial for an adequate understanding of Piaget's theory of cognition. Knowledge for Piaget, as for

Vico, is never (and can never be) a "representation" of the real world. Instead it is the collection of conceptual structures that turn out to be adapted or, as I would say, ***viable*** within the knowing subject's range of experience. (Note: Piaget nevertheless uses the word "representation"-- and so do I; but it is intended to refer to a ***re-presentation*** of a prior experience, not a picture of the "external" world; hence, I spell it with a hyphen, which Piaget did only occasionally.)

Both in the theory of evolution and the constructivist theory of knowing, "viability" is tied to the concept of equilibrium. Equilibrium in evolution indicates the state of an organism or species in which the potential for survival in a given environment is thought to be genetically assured. In the sphere of cognition, though indirectly linked to survival, equilibrium refers to a state in which an epistemic agent's cognitive structures have in the past yielded expected results, and continue to do so, without bringing to the surface conceptual conflicts or contradictions. In neither case is equilibrium necessarily a static affair, like the equilibrium of a balance beam, but it can be and often is dynamic, as the equilibrium maintained by a cyclist.

To make the Piagetian definition of knowledge plausible, one must immediately take into account (which so many interpreters of Piaget seem to omit) that a human subject's experience always includes the social interaction with other cognizing subjects. But introducing the notion of social interaction, raises a problem for constructivists. If what a cognizing subject knows cannot be anything but what that subject has constructed, it is clear that, from the constructivist perspective, the ***others*** with whom the subject may interact socially cannot be posited as an ontological given. I shall return to this problem, but first I want to exlicate the basis of a Piagetian theory of learning.

The Context of Scheme Theory

Two of the basic concepts of Piaget's theory of cognition are ***assimilation*** and ***accommodation***. Piaget's use of these terms is not quite the same as their common use in ordinary language. Both terms must be understood in the context of his constructivist theory of knowing. Unfortunately, contemporary textbooks in developmental psychology, most of which devote at least a few pages to Piaget, often fail to explain this. Thus one reads, for instance:

> Assimilation is the process whereby changing elements in the environment become incorporated into the structure of the organism. At the same time, the organism must accommodate its functioning to the nature of what is being assimilated. (Nash 1970, p. 360)

This is not at all what Piaget meant. One reason why assimilation is so often misunderstood is that its use as an explanatory postulate ranges from the unconscious to the deliberate. Another stems from disregarding the fact that Piaget uses that term, as well as "accommodation", within the framework of his theory of schemes. An example may help to illustrate the two extreme forms of assimilation.

When the nail that holds up the wire to my computer falls out of the wall in my study and I use my shoe to hammer it in again, I am deliberately ***assimilating*** the shoe to the function of a hammer. It may work, or it may not, but even if it does work I am not led to believe that the shoe ***is*** a hammer. The other form of assimilation--the one so many developmental psychologsts have misappropriated from Piaget--lacks that awareness. It is this second form that is epistemologically more interesting.

An infant quickly learns that a rattle it was given makes a rewarding noise when it is shaken. This provides the infant with the ability to generate the noise. Piaget sees this as the "construction of a ***scheme***" which, like all schemes, consists of three parts:

(1) Recognition of a certain situation (e.g. the presence of a graspable item with a rounded shape at one end);

(2) association of a specific activity with that kind of item (e.g. picking it up and shaking it);

(3) expectation of a certain result (e.g. the rewarding noise).

It is very likely that this infant, when placed in its high-chair at the dining table, will pick up and shake a graspable item that has a rounded shape at one end. We call that item a spoon and may say that the infant is ***assimilating*** it to its rattling scheme; but from the infant's perspective at that point, the item ***is*** a rattle. It is a rattle because what the infant perceives of it is just those aspects that fit the rattling scheme--and not what an adult would perceive as the characteristics of a spoon.[2] Then, however, when the infant shakes the item, it does not produce the result the infant expects: it does not rattle. This generates a perturbation ("disappointment"), and perturbation is one of the conditions that set the stage for cognitive change. In our example it may simply focus the infant's attention on the item in its hand, and this may lead to the perception of some aspect that will enable the infant in the future to recognize spoons as non-rattles. That development would be an ***accommodation***, but obviously a rather modest one. Alternatively, given the situation at the dining table, it is not unlikely that the spoon, being vigorously shaken, will hit the table and produce a different but also very rewarding noise. This, too, will generate a perturbation (we might call it "enchantment") which may lead to a different ***accommodation***, a major one this time, that initiates the "spoon banging scheme" which most parents know only too well.

This simple illustration of scheme theory also shows that the theory involves, on the part of the observer, certain presuppositions about cognizing organisms. The organism is supposed to possess at least the following capabilities[3]:

-- The ability and, beyond that, the tendency to establish recurrences in the flow of experience; this, in turn, entails at least two capabilities,
-- remembering and retrieving (re-presenting) experiences,
-- and the abilty to make comparisons and judgements of similarity and difference;
-- apart from these, there is the presupposition that the organism likes certain experiences better than others, which is to say, it has some elementary values.

The first three of these are indispensable in any theory of learning. Even the parsimonious models of classical and operant conditioning could not do without them. As to the fourth, the assumption of elementary values, it was explicitly embodied in Thorndikes ***Law of Effect***: "Other things being equal, connections grow stronger if they issue in satisfying states of affairs" (Thorndike 1931/1966, p.101). It remained implicit in psychological learning theories since Thorndike, but the subjectivity of what is "satisfying" was more or less deliberately obscured by behaviorists through the use of the more objective sounding term "reinforcement".

The learning theory that emerges from Piaget's work can be summarized by saying that cognitive change and ***learning*** take place when a scheme, instead of producing the expected result, leads to perturbation, and perturbation, in turn, leads to accommodation that establishes a new equilibrium. Learning and the knowledge it creates, thus, are explicitly instrumental. But here, again, it is crucial not to be rash and too simplistic in interpreting Piaget. His theory of cognition involves a two-fold instrumentalism. On the ***sensory-motor*** level, action schemes are instrumental in helping organisms to achieve goals in their interaction with their experiential world. On the level of ***reflective abstraction***, however, operative schemes are instrumental in helping organisms achieve a coherent conceptual network that reflects the paths of acting as well as thinking which, at the organisms' present point of experience, have turned out to be viable. The first instrumentality might be called "**utilitarian**" (the kind philosophers have traditionally scorned). The second, however, is strictly "**epistemic**". As such, may be of some philosophical interest--above all because it entails a radical shift in the conception of "knowledge", a shift that eliminates the paradoxical conception of Truth that requires a forever unattainable ontological test. The shift that substitutes viability in the experiential world for correspondence with ontological reality applies to knowledge that results from inductive inferences and generalizations. It does not affect deductive inferences in logic and mathematics. In Piaget's view, the certainty of conclusions in these areas pertains to mental operations and not to sensory-motor material (cf. Beth & Piaget 1961; Glasersfeld, 1985*b*).

The Social Component

In connection with the concept of viability, be it "utilitarian" or "epistemic", social interaction plays an important role. Except for animal psychologists, social interaction refers to what goes on among humans and involves language. As a rule it is also treated as essentially different from the interactions human organisms have with other items in their experiential field, because it is more or less tacitly assumed that humans are from the very outset privileged experiential entities. Constructivists have no intention of denying this intuitive human prerogative. But insofar as their theory of knowing attempts to model the cognitive development that provides the individual organism with ***all*** the furniture of his or her experiential field, they want to avoid assuming any cognitive structures or categories as innate. Hence, there is the need to hypothesize a model for the conceptual genesis of "others".

On the sensory-motor level, the schemes a developing child builds up and manages to keep viable will come to involve a large variety of "objects". There will be cups and spoons, building blocks and pencils, rag dolls and teddy bears--all seen, manipulated, and familiar as components of diverse action schemes. But there may also be kittens and perhaps a dog. Though the child may at first approach these items with action schemes that assimilate them to dolls or teddy bears, their unexpected reactions will quickly cause novel kinds of perturbation and inevitable accomodations. The most momentous of these accommodations can be roughly characterized by saying that the child will come to ascribe to these somewhat unruly entities certain properties that radically differentiate them from the other familiar objects. Among these properties will be the ability to move on their own, the ability to see and to hear, and eventually also the ability to feel pain. The ascription of these properties arises simply because, without them, the child's interactions with kittens and dogs cannot be turned into even moderately reliable schemes.

A very similar development may lead to the child's construction of schemes that involve still more complex items in her experiential environment, namely the human individuals who, to a much greater extent than other recurrent items of experience, make interaction unavoidable. (As we all remember, in many of these inescapable interactions, the schemes that are developed aim at avoiding unpleasant consequences rather than creating rewarding results.) Here again, in order to develop relatively reliable schemes, the child must impute certain capabilities to the objects of interaction. But now these ascriptions comprise not only perceptual but also cognitive capabilities, and soon these formidable "others" will be seen as intending, making plans, and being both very and not at all predictable in some respects. Indeed, out of the manifold of these frequent but nevertheless special interactions, there eventually emerges the way the developing human individual will think both of "others" and of him- or herself.

This reciprocity is, I believe, precisely what Kant had in mind when he wrote:

> It is manifest that, if one wants to imagine a thinking being, one would have to put oneself in its place and to impute one's own subject to the object one intended to consider ... (Kant 1781, p.223)

My brief account of the conceptual construction of "others" is no doubt a crude and preliminary analysis but it at least opens a possibility of approachng the problem without the vacuous assumption of innateness. Besides, the Kantian notion that we impute the cognitive capabilities that we isolate in ourselves to our conspecifics, leads to an explanation of why it means so much to us to have our experiential reality confirmed by others. The use of a scheme always involves the expectation of a more or less specific result. On the level of reflective abstraction, the expectation can be turned into a prediction. If we impute planning and foresight to others, this means that we also impute to them some of the schemes that have worked well for ourselves. Then, if a particular prediction we have made con-cerning an action or reaction of an other turns out to be corroborated by what the other does, this adds a second level of viability to our scheme; and this second level of viablity strengthens the experiential reality we have constructed. (cf. Glasersfeld 1985a, 1986)

A Perspective on Communication

The technical model of communication (Shannon 1948) established one feature of the process that remains important no matter from what orientation one approaches it: The physical signals that travel from one communicator to another--for instance the sounds of speech and the visual patterns of print or writing in linguistic communication--do not actually ***carry*** or ***contain*** what we think of as "meaning". Instead, they should be considered instructions to select particular meanings from a list which, together with the list of agreed signals, constitutes the "code" of the particular communication system. From this it follows that, if the two lists and the conventional associations that link the items in them are not available to a receiver before the linguistic interaction takes place, the signals will be meaningless for that receiver.

From the constructivist point of view, this feature of communication is of particular interest because it clearly brings out the fact that language users must individually ***construct*** the meaning of words, phrases, sentences, and texts. Needless to say, this semantic construction does not always have to start from scratch. Once a certain amount of vocabulary and combinatorial rules ("syntax") have been built up in interaction with speakers of the particular language, these patterns can be used to lead a learner to form novel combinations and, thus, novel conceptual compounds. But the basic elements out of which an individual's conceptual structures are composed and the relations by means of which they are held together cannot be transferred from one language user to another, let alone from a proficient speaker to an infant. These building blocks must be abstracted from individual experience; and their interpersonal fit, which makes possible what we call communication, can arise only in the course of protracted interaction, through mutual orientation and adaptation (cf. Maturana, 1980).

Though it is often said that normal children acquire their language without noticeable effort, a closer examination shows that the process involved is not as simple as it seems. If, for instance, you want your infant to learn the word "cup", you will go through a routine that parents have used through the ages. You will point to, and then probably pick up and move, an object that satisfies your definition of "cup", and at the same time you will repeatedly utter the word. It is likely that mothers and fathers do this "intuitively", i.e., without a well-formulated theoretical basis. They do it because it usually works. But the fact that it works does not mean that it has to be a simple matter. There are at least three essential steps the child has to make.

The first consists in focusing attention on some specific sensory signals in the manifold of signals which, at every moment, are available within the child's sensory system; the parent's pointing provides a merely approximate and usually quite ambiguous direction for this act.

The second step consists in isolating and coordinating a group of these sensory signals to form a more or less discrete visual item or "thing". The parent's moving the cup greatly aids this process because it accentuates the relevant figure as opposed to the parts of the visual field that are to form the irrelevant ground.[4]

The third step, then, is to associate the isolated visual pattern with the auditory experience produced by the parent's utterances of the word "cup". Again, the child must first isolate the sensory signals that constitute this auditory experience from the background (the manifold auditory signals that are available at the moment); and the parent's repetition of the word obviously enhances the process of isolating the auditory pattern as well as its association with the moving visual pattern.

If this sequence of steps provides an adequate analysis of the initial acquisition of the meaning of the word "cup", it is clear that the child's meaning of that word is made up exclusively of elements which the child abstracts from her own experience. Indeed, anyone who has more or less methodically watched children acquire the use of new words, will have noticed that what they isolate as meanings from their experiences in conjunction with words is often only partially compatible with the meanings the adult speakers of the language take for granted. Thus the child's initial concept of cup often includes the activity of drinking, and sometimes even what is being drunk, e.g., milk. Indeed, it may take quite some time before the continual linguistic and social interaction with other speakers of the language provides occasions for the accommodations that are necessary for the concept the child associates with the word "cup" to become adapted to the adults' extended use of the word, for instance, in the context of golf greens or championships in sport.

The process of accommodating and tuning the meaning of words and linguistic expressions actually continues for each of us throughout our lives. No matter how long we have spoken a language, there will still be occasions when we realize that, up to that point, we have been using a word in a way that now turns out to be idiosyncratic in some particular respect.

Once we come to see this essential and inescapable subjectivity of linguistic meaning, we can no longer maintain the preconceived notion that words ***convey*** ideas or knowledge; nor can we believe that a listener who apparently "understands" what we say must necessarily have conceptual structures that are identical with ours. Instead, we come to realize that "understanding" is a matter of fit rather than match. Put in the simplest way, to understand what someone has said or written means no less but also no more than to have built up a conceptual structure that, in the given context, appears to be ***compatible*** with the structure the speaker had in mind--and this compatibility, as a rule, manifests itself in no other way than that the receiver says and does nothing that contravenes the speaker's expectations.

Among proficient speakers of a language, the individual's conceptual idiosyncracies rarely surface when the topics of conversation are everyday objects and events. To be considered proficient in a given language requires two things among others: to have available a large enough vocabulary, and to have constructed and sufficiently accommodated and adapted the meanings associated with the words of that vocabulary so that no conceptual discrepancies become apparent in ordinary linguistic interactions. When conversation turns to predominantly abstract matters, it usually does not take long before conceptual discrepancies become noticeable--even among proficient speakers. The discrepancies generate perturbations in the interactors, and at that point the difficulties become insurmountable if the participants believe that the meanings they attribute to the words they use are ***true representations*** of fixed entities in an objective world apart from any speaker. If, instead, the participants take a constructivist view and assume that a language user's meanings cannot be anything but subjective constructs derived from the speaker's individual experiences, some accommodation and adaptation is usually possible.

From this perspective, the use of language--for instance in teaching--is far more complicated than it is mostly presumed to be. It cannot be a means of ***transferring*** "information" or knowledge to the student. As Rorty says: "The activity of uttering sentences is one of the things people do in order to cope with their environment" (1982, p.XVII).

This inherent and inescapable indeterminacy of linguistic communication is something the best teachers have always known. Intuitively, they have also been aware of the fact that students whose cognitive structures are not "perturbed" by an experience they themselves register as a failure, will not "accommodate" to form new concepts and new understanding, but continue to "assimilate" new experiences to the structures they already have. Thus, independently of any epistemological orientation, they have always known that "telling" is not enough, because understanding is not a matter of passively receiving but of actively building up.

Some Further Considerations

The pattern of assimilation is a pattern of maintaining categorizations, concepts, and, indeed, whole theories until some experience makes their adequacy questionable. It is a universal pattern from the constructivist point of view. Whenever a thinking subject has theories and concepts that have proved useful in the past, the subject has, as it were, a considerable vested interest in maintaining the status quo. That is to say, the holders of a theory will assimilate new experiences as long as they possibly can, even in the face of considerable perturbations.

Silvio Ceccato, the Italian pioneer in the analysis of mental operations and construction, once after a public discussion of his work, overheard an aged philosopher say: "If

Ceccato were right, the rest of us would be fools!"[5]

Most readers of the works of Piaget and the contemporary constructivists are not as direct and outspoken. Instead, they desperately try to assimilate what they read and hear, disregarding all sorts of clues and bending the interpretation of words to their own notions; and when this proves impossible, they conclude that the author is contradicting himself, because what he says is no longer compatible with their own conceptual construction. In this vein, they often brand constructivism as just another form of solipsism. But in doing so, they of course disregard the fact that constructivism does not ***deny*** an ontological reality--it merely holds that no such reality can be ***known***.

Radical constructivism is unashamedly ***instrumentalist*** (in the philosophical sense of that term) and this cannot but offend advocates of the maxim "Truth for Truth's sake". Consequently, if they don't call it solipsism, they dismiss it as cheap materialism. But this, again, is inappropriate. The instrumentalism embodied in constructivism is not to be equated with materialism. The second principle listed above states that the function of the cognitive activity is ***adaptive***. The concept of adaptation intended here, as I said before, is the basic biological concept in the Darwinian theory of evolution. It refers to the ***fit*** with the environment, which is to say, every species or organism found alive and capable of reproducing must, by that very fact, be considered adapted at that moment in the history of living organisms. To be adapted, therefore, means no more and no less than to be ***viable***.

For the observing biologist, of course, this viability refers to the fit with an environment that is external to the organism. But precisely because the biologist is an observer, this environment cannot be the "world-in-itself". From the constructivist point of view, to ***observe*** means to focus attention on a specific part of one experiential field. Usually, the focusing of attention involves categorizing what one focuses on as an item of a particular kind, a property, a relation, a thing, a process, etc. The moment such a categorization is made, the rest of one's experiential field becomes the item's ***environment.*** This is analogous to the simultaneous creation of a "figure" and its "ground". There, too, the moment the line drawn by the pencil becomes categorizable as an image of a particular item, the rest of the sheet is inevitably seen as "ground". The observer's discrimination of figure and ground or organism and environment is, of course, quite legitimate, and so is the biologists subsequent observation of specific relationships between the observed organisms and their environment. But once it is understood that all this discriminating, categorizing, and establishing of relationships takes place ***within the observer's experiential field***, it becomes clear that no result of these operations can pertain to the world as such, that is, the world as it might "exist" objectively ***without*** the observer's activties.

FOOTNOTES

1. Montaigne wrote this in his *Apologie de Raymond Sebond* (1575-76) and it can be found on p.139 of volume 2 of the complete edition of his ***Essais***, edited by Pierre Michel (Paris, 1972).
2. This notion of assimilation seems quite compatibe with the view of philosophers of science who maintain that all observation is necessarily "theory-laden".
3. Piaget nowhere lists these presuppositions, but they are implicit in his analysis of conceptual development (cf., for instance, Piaget, 1937 and 1967*b*).
4. Note that, even if the child has coordinated sensory signals to form such a "thing" in the past, each new ***recognition*** involves isolating it in the current experiential field.
5. I owe this anecdote to a personal communication: Silvio Ceccato told it to me shortly after the event, sometime about 1960.

REFERENCES

Anonymous: 1711, 'Osservazioni', *Giornale de' Letterati d'Italia* (Venice), 5, article VI.

Berkeley, G.: 1710. *A Treatise Concerning the Principles of uman Understanding*. Open Court, La Salle, Illinois (1963).

Beth, E.W., & Piaget, J.: 1961. *Epistemologie Mathematique et Psychologie*. Presses Universitaires de France, Paris.

Fleck, L.: 1929. 'Zur Krise der "Wirklichkeit"', *Die Naturwissenschaften, 17*, 23, pp.425-430. (The excerpts from this paper were translated by E.v.G.)

Fleck, L.: 1935. *Entstehung und Entwiklung einer wissenschaftlichen Tatsache*. Benno Schwabe, Basel, Switzerland. (Reprinted: Suhrkamp, 1980, Nördlingen, Germany.)

Foerster, H.von: 1970. 'Thoughts and Notes on Cognition', in P.Garvin (Ed.), *Cognition: A Multiple View*, pp.25-48. Spartan Books, New York.

Glasersfeld, E.von: 1985. 'Representation and Deduction', in L.Streefland (ed.), *Proceedings of the 9th Conference for the Psychology of Mathematics Education, Vol.1*, State University of Utrecht, The Netherlands, pp.484-489.

Glasersfeld, E.von: 1986. 'Steps in the Construction of "Others" and "Reality", in R.Trappl (ed.), *Power, Autonomy, Utopia*, Plenum, London, pp. 107-116.

Hebb, D.O.: 1958. 'Alice in Wonderland or Psychology among the Biological Sciences', in Harlow & Woolsey (eds.), *Biological and Biochemical Bases of Behavior*, University of Wisconsin Press, Madison, pp.451-467.

Kant, I.: 1781. *Kritik der reinen Vernunft, 1. Auflage* (Gesammelte Schriften, Bd.IV). Königl. Preussische Akademie, 1910ff., Berlin.

Kuhn, T.S.: 1970. *The Structure of Scientific Revolutions*. University of Chicago Press (2nd edition), Chicago. (First published 1962.)

McCulloch, W.S.: 1965. *Embodiments of Mind*. M.I.T. Press, Cambridge, MA.

Maturana, H.R.: 1980. 'Biology of Cognition', in H.R.Maturana & F.J.Varela (eds.), *Autopoiesis and Cognition*, Reidel, Dordrecht/Boston, pp. 5-58.

Nash, J.: 1970. *Developmental Psychology*. Prentice-Hall, Englewood Cliffs, New Jersey.

Piaget, J.: 1937. *La construction du réel chez l'enfant*. Delachaux et Niestlé, Neuchâtel, Switzerland.

Piaget, J.: 1967 a. *Biologie et connaissance*. Gallimard, Paris.

Piaget, J,: 1967 b. *Six Psychological Studies*. Vintage Books, New York.

Pittendrigh, C.S.: 1958, 'Adaptation, Natural Selection, and Behavior", in A.Roe & G.G.Simpson (eds.), *Behavior and Evolution*, Yale University Press, New Haven, pp.390-416.

Rorty, R.: 1982. *Consequences of Pragmatism*. University of Minnesota Press, Minneapolis.

Shannon, C.: 1948, 'The Mathematical Theory of Communication', *Bell Systems Technical Journal*, *27*, 379-423 and 623-656.

Thorndike, E.: 1966. *Human Learning*. M.I.T.Press, Cambridge, Massachusetts (first published, 1931).

Vico, G.: 1710, *De antiquissima Italorum sapientia*. (With Italian translation by F.S.Pomodoro, Stamperia de' Classici Latini, Naples, 1858).

TABLE OF CONTENTS

CONTENT PAGE

ONTOLOGY OF OBSERVING:

THE BIOLOGICAL FOUNDATIONS OF SELF CONSCIOUSNESS AND THE PHYSICAL DOMAIN OF EXISTENCE

Humberto R. Maturana

1. PURPOSE

My purpose in this essay is to explain cognition as a biological phenomenon, and to show, in the process, how language arises and gives origin to self consciousness, revealing the ontological foundations of the physical domain of existence as a limiting cognitive domain. In order to do this, I shall start from two unavoidable experiential conditions that are at the same time my problem and my explanatory instruments, namely: a) that cognition, as is apparent in the fact that any alteration of the biology of our nervous system alters our cognitive capacities, is a biological phenomenon that must be explained as such; and b) that we, as is apparent in this very same essay, exist as human beings in language using language for our explanations. These two experiential conditions are my starting point because I must be in them in any explanatory attempt; they are my problem because I choose to explain them; and they are my unavoidable instruments because I must use cognition and language in order to explain cognition and language.

In other words, I propose not to take cognition and language as given unexplainable properties, but to take them as phenomena of our human domain of experiences that arise in the praxis of our living, and that as such deserve explanation as biological phenomena. At the same time, it is my purpose to use our condition of existing in language to show how the physical domain of existence arises in language as a cognitive domain. That is, I intend to show that the observer and observing, as biological phenomena, are ontologically primary with respect to the object and the physical domain of existence.

2. THE PROBLEM

I shall take cognition as the fundamental problem, and I shall explain language in the process of explaining cognition.

We human beings assess cognition in any domain by specifying the domain with a question and demanding adequate behavior or adequate action in that domain. If what we observe as an answer satisfies us as adequate behavior or as adequate action in the domain specified by the question, we accept it as an expression of cognition in that domain, and claim that he or she who answers our query knows. Thus, if someone claims to know algebra--that is, to be an algebraist--we demand of him or her to perform in the domain of what we consider algebra to be, and if according to us she or he performs adequately in that domain, we accept the claim. If the question asked is not answered with what we consider to be adequate behavior or adequate action in the domain that it specifies, the being asked to perform (the algebraist) disintegrates or disappears, it loses its class identity as an entity existing in the operational domain specified by the question, and the questioner proceeds henceforth according to its nonexistence. In these circumstances, since adequate behavior (or adequate action) is the only criterion that we have and can use to assess cognition, I shall take adequate behavior or adequate action in any domain specified by a question, as the phenomenon to be explained when explaining cognition.

3. NATURE OF THE ANSWER

I am a biologist, and it is from my experience as a biologist that in this essay I am treating the phenomenon of cognition as a biological phenomenon. Furthermore, since as a biologist I am a scientist, it is as a scientist that I shall provide a biological explanation of the phenomenon of cognition. In order to do this: a) I shall make explicit what I shall consider as an adequate behavior in the context of what I consider is a scientific explanation (section 4), so that all the implications of my explanation may be apparent to the reader and she or he may know when it is attained; b) I shall make explicit my epistemological standing with respect to the notion of objectivity (section 5), so that the ontological status of my explanation may be apparent; c) I shall make explicit the notions that I shall use in my explanation by showing how they belong to our daily life (section 6), so that it may be apparent how we are involved as human beings in the explanation that I shall provide; and d) I shall make explicit the nature of the biological phenomena involved in my explanation (section 7), so that it may be apparent how we are involved as living systems in the explanation as well as in the phenomenon of cognition itself. Finally, in the process of explaining the phenomenon of cognition as a biological phenomenon I shall show how it is that scientific theories arise as free creations of the human mind, how it is that they explain human experience and not an independent objective world, and how the physical domain of existence arises in the explanation of the praxis of living of the observer as a feature of the ontology of observing (sections 8 to 11).

4. THE SCIENTIFIC DOMAIN

We find ourselves as human beings here and now in the praxis of living, in the happening of being human, in language languaging, in an *a priori* experiential situation in which everything that is, everything that happens, is and happens in us as part of our praxis of living. In these circumstances, whatever we say about how anything happens takes place in the praxis of our living as a comment, as a reflection, as a reformulation, in short, as an explanation of the praxis of our living, and as such it does not replace or constitute the praxis of living that it purports to explain. Thus, to say that we are made of matter, or to say that we are ideas in the mind of God, are both explanations of that which we live as our experience of being, yet neither matter nor ideas in the mind of God constitute the experience of being that which they are supposed to explain. Explanations take place operationally in a metadomain with respect to that which they explain. Furthermore, in daily life, in the actual dynamics of human interactions, an explanation is always an answer to a question about the origin of a given phenomenon, and is accepted or rejected by a listener who accepts or rejects it according to whether or not it satisfies a particular implicit or explicit criterion of acceptability that he or she specifies. Therefore, there are as many different kinds of explanations as there are different criteria of acceptability of reformulations of the happening of living of the observers that the observers specify. Accordingly, every domain of explanations as it is defined by a particular criterion of acceptability, constitutes a closed cognitive domain as a domain of acceptable statements or actions for the observers that accept that criterion of acceptability. Science, modern science, as a cognitive domain is not an exception to this. Indeed, modern science is that particular cognitive domain that takes what is called the scientific explanation as the criterion of validation (acceptability) of the statements that pertain to it. Let me make this explicit.

i) <u>Scientific explanations</u>. Scientists usually do not reflect upon the constitutive conditions of science. Yet, it is possible to abstract, from what modern scientists do, an operational (and, hence, experiential) specification of what constitutes a scientific explanation as the criterion of validation of what they claim are their scientific statements. Furthermore, it is possible to describe this criterion of validation of scientific statements as a reformulation of what is usually called the scientific method.

<u>A</u>. Different domains of human activities entail different intentions. Thus, as the intention of doing art is to generate an aesthetic experience, and the intention of doing technology is to produce, the intention of doing science is to explain. It is, therefore, in the context of explaining that the criterion of validation of a scientific explanation is the conjoined satisfaction, in the praxis of living of an observer, of four operational conditions, one of which, the proposition of an <u>ad hoc</u> mechanism that generates the phenomenon explained as a phenomenon to be witnessed by the observer in his or her praxis of living, is the scientific explanation. And, it is in the context of explaining that it must be understood that the scientific explanation is the criterion of validation of scientific statements. Finally, it is also in the context of explaining that it must be recognized that a modern scientific community is a community of observers (henceforth called standard observers) that use the scientific explanation as the criterion of validation of their statements. Now, the criterion of validation of

scientific explanations entails four operational conditions:

a) The specification of the phenomenon to be explained through the stipulation of the operations that a standard observer must perform in his or her praxis of living in order to also be a witness of it in his or her praxis of living.

b) The proposition, in the domain of operational coherences of the praxis of living of a standard observer, of a mechanism, a generative mechanism, which when allowed to operate gives rise as a consequence of its operation to the phenomenon to be explained, to be witnessed by the observer also in his or her praxis of living. This generative mechanism, that is usually called the explanatory hypothesis, takes place in the praxis of living of the observer in a different phenomenal domain than the phenomenal domain in which the phenomenon to be explained is witnessed, and the latter as a consequence of the former stands in an operational metadomain with respect to it. Indeed, the phenomenon to be explained and its generative mechanism take place in different nonintersecting phenomenal domains in the praxis of living of the observer.

c) The deduction, that is, the computation, in the domain of operational coherences of the praxis of living of the standard observer entailed by the generative mechanism proposed in (b), of other phenomena that the standard observer should be able to witness in his or her domain of experiences as a result of the operation of such operational coherences, and the stipulation of the operations that he or she should perform in order to do so.

d) The actual witnessing, in his or her domain of experiences, of the phenomena deduced in (c) by the standard observer who actually performs in his or her praxis of living the operations stipulated also in (c).

If these four operational conditions are conjointly satisfied in the praxis of living of the standard observer, the generative mechanism proposed in (b) becomes a scientific explanation of the phenomenon brought forth in (a). These four operational conditions in the praxis of living of the observer constitute the criterion of validation of scientific explanations, and science (modern science) is a domain of statements directly or indirectly validated by scientific explanations. Accordingly, it follows that there are no such things as scientific observations, scientific hypotheses, or scientific predictions; there are only scientific explanations and scientific statements. It also follows that the standard observer can make scientific statements in any domain of his or her praxis of living in which he or she can make scientific explanations.

<u>B</u>. According to <u>A</u>, a scientific statement is valid as a scientific statement only within the community of standard observers that is defined as such because they can realize and accept the scientific explanation as the criterion of validation of their statements. This makes scientific statements consensual state-

ments, and the community of standard observers a scientific community. That in principle any human being can belong to the scientific community is due to two facts of experience: one is that it is as a living human being that an observer can realize and accept the scientific explanation as the criterion of validation of his or her statements and become a standard observer, the other is that the criterion of validation of scientific statements is the operational criterion of validation of actions and statements in daily life, even if it is not used with the same care in order to avoid confusion of phenomenal domains. Indeed, these two experiential facts constitute the fundament for the claim of universality that scientists make for their statements, but what is peculiar to scientists is that they are careful to avoid confusion of phenomenal domains when applying the criterion of validation of scientific statements in the praxis of living.

C. Scientists and philosophers of science usually believe that the operational effectiveness of science and technology reveals an independent objective reality, and that scientific statements reveal the features of an independent universe, of an objective world. Or, in other words, many scientists and philosophers of science believe that without the independent existence of an objective reality, science could not take place. Yet, if one does, as I have done above, a constitutive, an ontological, analysis of the criterion of validation of scientific statements, one can see that scientific explanations do not require the assumption of objectivity because scientific explanations do not explain an independent objective reality. Scientific explanations explain the praxis of living of the observer, and they do so with the operational coherences brought forth by the observer in his or her praxis of living. It is this fact that gives science its biological foundations and that makes science a cognitive domain bound to the biology of the observer with characteristics that are determined by the ontology of observing.

ii) Science. In conclusion, the operational description of what constitutes a scientific explanation as the criterion of validation of scientific statements, reveals the following characteristics of scientific statements in general, and of science as a domain of scientific statements in particular:

A. Scientific statements are consensual statements valid only within the community of standard observers that generates them; and science as the domain of scientific statements does not need an independent objective reality, nor does it reveal one. Therefore, the operational effectiveness of science as a cognitive domain rests only on the operational coherence that takes place in the praxis of living of the standard observers that generate it as a particular domain of consensual coordinations of actions in the praxis of their living together as a scientific community. Science is not a manner of revealing an independent reality; it is a manner of bringing forth a particular one bound to the conditions that constitute the observer as a human being.

B. Since the members of a community of standard observers can generate scientific statements in any phenomenal domain of the praxis of living in which they can apply the criterion of validation of scientific statements, the universality of a particular body of scientific statements within the human domain will depend on the universality in the human domain of the standard observers that can generate such a body of scientific statements. Finally, scientific statements are valid only as long as the scientific explanations that support them are valid, and these are valid only as long as the four operational conditions that must be conjointly satisfied in their constitution are satisfied for all the phenomena that are deduced in the praxis of living of the standard observers in the domain of operational coherences specified by the proposed generative mechanism.

C. It is frequently said that scientific explanations are reductionist propositions, adducing that they consist in expressing the phenomena to be explained in more basic terms. This view is inadequate. Scientific explanations are constitutively non-reductionist explanations because they consist in generative propositions and not in expressing the phenomena of one domain in phenomena of another. This is so because in a scientific explanation the phenomenon explained must arise as a result of the operation of the generative mechanism, and cannot be part of it. In fact, if the latter were the case, the explanatory proposition would be constitutively inadequate and would have to be rejected. The phenomenon explained and the phenomena proper to the generative mechanism constitutively pertain to nonintersecting phenomenal domains.

D. The generative mechanism in a scientific explanation is brought forth by a standard observer from his or her domain of experiences in his or her praxis of living as an *ad hoc* proposition that in principle requires no justification. Therefore, the components of the generative mechanism, as well as the phenomena proper to their operation, have a foundational character with respect to the phenomenon to be explained, and as such their validity is in principle accepted *a priori*. Accordingly, every scientific domain as a domain of scientific statements is founded on basic experiential premises not justified in it, and constitutes in the praxis of living of the standard observer a domain of operational coherences brought forth in the operational coherences entailed in the generative mechanisms of the scientific explanations that validate it.

5. OBJECTIVITY IN PARENTHESES

If one looks at the two shadows of an object that simultaneously partially intercepts the paths of two different lights, one white and one red, and if one has trichromatic vision, then one sees that the area of the shadow from the white light that receives red light looks red, and that the area of the shadow from the red light that receives white light looks blue-green. This experience is compelling and unavoidable, even if one knows that the area of the shadow from the red light should look white or gray because it receives only white light. If one asks how it is that one sees blue-green where there is white light only, one is told by a reliable authority that the experience of the blue-green shadow is a chromatic illusion because there is no blue-green shadow to justify it as a perception. We live numerous experiences in our daily life that we class like this as illusions or hallucinations and not as perceptions, claiming that they do not constitute the capture of an independent reality because we can disqualify them by resorting to the opinion of a friend whose authority we accept, or by relying upon a different sensory experience that we consider as a more acceptable perceptual criterion. In the experience itself, however, we cannot distinguish between what we call an illusion, a hallucination, or a perception: illusion, hallucination, and perception are experientially indistinguishable. It is only through the use of a different experience as a metaexperiential authoritative criterion of distinction, either of the same observer or of somebody else subject to similar restrictions, that such a distinction is socially made. Our incapacity to experientially distinguish between what we socially call illusion, hallucination, or perception, is constitutive in us as living systems, and is not a limitation of our present state of knowledge. The recognition of this circumstance should lead us to put a question mark on any perceptual certainty.

i) An invitation. The word "perception" comes from the Latin expression per capire, which means "through capture" and carries with it the implicit understanding that to perceive is to capture the features of a world independent of the observer. This view assumes objectivity, and, hence, the possibility of knowing a world independent of the observer, as the ontological condition on which the distinction between illusion, hallucination, and perception that it entails is based. Therefore, to question the operational validity in the biological domain of the distinction between illusion, hallucination, and perception, is to question the ontological validity of the notion of objectivity in the explanation of the phenomenon of cognition. But, how then to proceed? Any reflection or comment about how the praxis of living comes about is an explanation, a reformulation of what takes place. If this reformulation does not question the properties of the observer, if it takes for granted cognition and language, then it must assume the independent existence of what is known. If this reformulation questions the properties of the observer, if it asks about how cognition and language arise, then it must accept the experiential indistinguishability between illusion, hallucination, and perception, and take as constitutive that existence is dependent upon the biology of the observer. Most philosophical traditions pertain to the first case, assuming the independent existence of something, such as matter, energy, ideas, God, mind, spirit,...or reality. I invite the reader to follow the second path, and to take seriously the constitutive condition of the biological condition of the observer, following all the consequences that this constitutive condition entails.

ii) Objectivity in parentheses. The assumption of objectivity is not needed for the generation of a scientific explanation.

Therefore, in the process of being a scientist explaining cognition as a biological phenomenon I shall proceed without using the notion of objectivity to validate what I say; that is, <u>I shall put objectivity in parentheses</u>. In other words, I shall go on using an object language because this is the only language that we have (and can have), but although I shall use the experience of being in language as my starting point while I use language to explain cognition and language, <u>I shall not claim</u> that what I say is valid because there is an independent objective reality that validates it. I shall speak as a biologist, and as such I shall use the criterion of validation of scientific statements to validate what I say, accepting that everything that takes place is brought forth by the observer in his or her praxis of living as a primary experiential condition, and that any explanation is secondary.

iii) <u>The universum versus the multiversa</u>. The assumption of objectivity, objectivity without parentheses, entails the assumption that existence is independent of the observer, that there is an independent domain of existence, the <u>universum</u>, that is the ultimate reference for the validation of any explanation. With objectivity without parentheses, things, entities, exist with independency of the observer that distinguishes them, and it is this independent existence of things (entities, ideas) that specifies the truth. Objectivity without parentheses entails unity, and, in the long run, reductionism, because it entails reality as a single ultimate domain defined by independent existence. He or she who has access to reality is necessarily right in any dispute, and those who do not have such access are necessarily wrong. In the universum, coexistence demands obedience to knowledge.

Contrary to all this, objectivity with parentheses entails accepting that existence is brought forth by the distinctions of the observer, that there are as many domains of existence as kinds of distinctions the observer performs: objectivity in parentheses entails the <u>multiversa</u>, entails that existence is constitutively dependent upon the observer, and that there are as many domains of truths as domains of existence she or he brings forth in her or his distinctions. At the same time, objectivity in parentheses entails that different domains of existence constitutively do not intersect because they are brought forth by different kinds of operations of distinction, and, therefore, it constitutively negates phenomenal reductionism. Finally, under objectivity in parentheses, each versum of the multiversa is equally valid if not equally pleasant to be part of, and disagreements between observers, when they arise not from trivial logical mistakes within the same versum but from the observers standing in different versa, will have to be solved not by claiming a privileged access to an independent reality but through the generation of a common versum through coexistence in mutual acceptance. In the multiversa, coexistence demands consensus, that is, common knowledge.

6. BASIC NOTIONS

Everything said is said by an observer to another observer that could be him or herself. Since this condition is my experiential starting point in the praxis of living as well as my problem, I shall make explicit some of the notions that I shall use as my tools for explaining the phenomena of cognition and language, and I shall do so by revealing the actions in the praxis of living that they entail in our daily life when we do science. Indeed, by revealing what we do as observers _I am making explicit the ontology of the observer as a constitutive human condition_.

i) _The observer_. An observer is, in general, any being operating in language, or, in particular, any human being, in the understanding that language defines humanity. In our individual experience as human beings we find ourselves in language, we do not see ourselves growing into it: we are already observers by being in language when we begin as observers to reflect upon language and the condition of being observers. In other words, whatever takes place in the praxis of living of the observer takes place as distinctions in language through languaging, and this is all that he or she can do as such. One of my tasks is to show how the observer arises.

ii) _Unities_. The basic operation that an observer performs in the praxis of living is the operation of distinction. In the operation of distinction an observer brings forth a unity (an entity, a whole) as well as the medium in which it is distinguished, and entails in this latter all the operational coherences that make the distinction of the unity possible in his or her praxis of living.

iii) _Simple and composite unities_. An observer may distinguish in the praxis of living two kinds of unities, simple and composite unities. A simple unity is a unity brought forth in an operation of distinction that constitutes it as a whole by specifying its properties as a collection of dimensions of interactions in the medium in which it is distinguished. Therefore, a simple unity is exclusively and completely characterized by the properties through which it is brought forth in the praxis of living of the observer that distinguishes it, and no further explanation is needed for the origin of these properties. A simple unity arises defined and characterized by a collection of properties as a matter of distinction in the praxis of living of the observer.

A composite unity is a unity distinguished as a simple unity that through further operations of distinction is decomposed by the observer into components that through their composition would constitute the original simple unity in the domain in which it is distinguished. A composite unity, therefore, is operationally distinguished as a simple unity in a metadomain with respect to the domain in which its components are distinguished because it results as such from an operation of composition. As a result, the components of a composite unity and its correlated simple unity are in a constitutive relation of mutual specification. Thus, the properties of a composite unity distinguished as a simple one entail the properties of the components that constitute it as such, and conversely, the properties of the components of a composite unity and their manner of composition determine the properties that characterize it as a simple unity when distinguished as such. Accordingly, there is no such thing as the distinction of a component independently of the unity that it integrates, nor can a simple unity distinguished as a composite one be decomposed into an arbitrary set of components disposed in an arbitrary manner of composition. Indeed, there is no such thing as a free component floating around independently of the composite unity that it integrates. Therefore, whenever we say that we treat a simple unity as a composite one, and we claim that we do so by distinguishing in it elements that when put together do not regenerate the original unity, we in fact are not decomposing the unity that we believe that we are decomposing but another one, and the elements that we distinguish are not components of the composite unity that we say that they compose.

iv) Organization and structure. A particular composite unity is characterized by the components and relations between components that constitute it as a composite unity that can be distinguished, in a metadomain with respect to its components, as a particular simple unity of a certain kind. As such, a particular composite unity has both organization and structure. These can be characterized as follows:

a) The relations between components in a composite unity that make it a composite unity of a particular kind, specifying its class identity as a simple unity in a metadomain with respect to its components, constitutes its organization. In other words, the organization of a composite unity is the configuration of static or dynamic relations between its components that specifies its class identity as a composite unity that can be distinguished as a simple unity of a particular kind. Therefore, if the organization of a composite unity changes, the composite unity loses its class identity; that is, it disintegrates. The organization of a composite unity is necessarily an invariant while it conserves its class identity, and *vice versa*, the class identity of a composite unity is necessarily an invariant while the composite unity conserves its organization.

b) In a composite unity, be this static or dynamic, the actual components plus the actual relations that take place between them while realizing it as a particular composite unity characterized by a particular organization, constitute its structure. In other words, the structure of a particular composite unity is the manner in which it is actually made by actual static or dynamic components and relations in a particular space, and a particular composite unity conserves its class identity only as long as its structure realizes in it the organization that defines its class identity. Therefore, in any particular composite unity the configuration of relations between components that constitutes its organization must be realized in its structure as a subset of all the actual relations that hold between its components as actual entities interacting in the composition.

It follows from all this that the characterization of the organization of a composite unity as a configuration of relations between components says nothing about the characteristics or properties of these components other than that they must satisfy the relations of the organization of the composite unity through their interactions in its composition. It also follows that the structure of a composite unity can change without it losing its class identity if the configuration of relations that constitutes its organization is conserved through such structural changes. At the same time, it also follows that if the organization of a composite unity is not conserved through its structural changes, the composite unity loses its class identity, it disintegrates, and something else appears in its stead. Therefore, a dynamic composite unity is a composite unity in continuous structural change with conservation of organization.

v) Structure determined systems. Since the structure of a composite unity consists in its components and their relations, any change in a composite unity consists in a structural change, and arises in it at every instant necessarily determined by its structure at that instant through the operation of the properties of its components. Furthermore, the structural changes that a composite unity undergoes as a result of an interaction are also determined by the structure of the composite unity, and this is so because such structural changes take place in the interplay of the properties of the components of the composite unity as they are involved in its composition. Therefore, an external agent that interacts with a composite unity only triggers in it a structural change that it does not determine. Since this is a constitutive condition for composite unities, nothing external to them can specify what happens in them: there are no instructive interactions for composite unities. Finally, and as a result of this latter condition, the structure of a composite unity also determines with which structural configurations of the medium it may interact. In general, then, everything that happens in a composite unity is a structural change, and every structural change occurs in a composite unity determined at every instant by its structure at that instant. This is so both for static and for dynamic composite unities, and the only difference between these is that dynamic composite unities are in a continuous structural change generated as part of their structural constitution in the context of their interactions, while static ones are not. It follows from all this that composite unities are structure determined systems in the sense that everything that happens in them is determined by their structure. This can be systematically expressed by saying that the structure of a composite unity determines in it at every instant:

a) the domain of all the structural changes that it may undergo with conservation of organization (class identity) and adaptation at that instant; I call this domain the instantaneous domain of the possible changes of state of the composite unity.

b) the domain of all the structural changes that it may undergo with loss of organization and adaptation at that instant; I call this domain the instantaneous domain of the possible disintegrations of the composite unity.

c) the domain of all the different structural configurations of the medium that it admits at that instant in interactions that trigger in it changes of state; I call this domain the instantaneous domain of the possible perturbations of the composite unity.

d) the domain of all the different structural configurations of the medium that it admits at that instant in interactions that trigger in it its disintegration; I call this domain the instantaneous domain of the possible destructive interactions of the composite unity.

These four domains of structural determinism that characterize every structure determined system at every instant are obviously not fixed, and they change as the structure of the structure determined system changes in the flow of its own internal structural dynamics or as a result of its interactions. These general characteristics of structure determined systems have several additional consequences of which I shall mention six. The first is that during the ontogeny of a structure determined system, its four domains of structural determinism change following a course contingent to its interactions and its own internal structural dynamics. The second is that some structure determined systems have recurrent domains of structural determinism because they have recurrent structural configurations, while others do not because their structure changes in a nonrecurrent manner. The third is that although the structure of a structure determined system determines the structural configurations of the medium with which it may interact, all its interactions with independent systems arise as coincidences, and these coincidental interactions cannot be predicted from the structure of the structure determined system alone. The fourth is that a composite unity exists only while it moves through the medium in interactions that are perturbations, and that it disintegrates at the first destructive interaction. The fifth is that since the medium cannot specify what happens in a structure determined system because it only triggers the structural changes that occur in the system as a result of the system's interactions, all that can happen to a composite unity in relation to its interactions in the medium is that the course followed by its structural changes is contingent upon the sequence of these interactions. Finally, the sixth is that since mechanistic systems are structure determined systems, and since scientific explanations entail the proposition of mechanistic systems as the systems that generate the phenomena to be explained, in scientific explanations we deal, and we can only deal, with structure determined systems.

vi) Existence. By putting objectivity in parentheses, we accept that constitutively we cannot claim the independent existence of things (entities, unities, ideas, etc.), and we recognize that a unity exists only in its distinction, in the praxis of living of the observer, that brings it forth. But we also recognize that the distinction takes place in the praxis of living of the observer in an operation that specifies simultaneously the class identity of the unity distinguished, either as a simple unity or as a composite one, and its domain of existence as the domain of the operational coherences in which its distinction makes sense as a feature of his or her praxis of living. Since the class identity of a composite unity is defined by its organization, and since this can be realized in a composite unity only while it interacts in a domain of perturbations, existence in a composite unity entails the conservation of its organization as well as the conservation of its operational structural correspondence in the domain of operational coherences in which it is distinguished. Similarly, since the class identity of a simple unity is defined by its properties, and since these are defined in relation to the operational domain in which the simple unity is distinguished, existence in a simple unity entails the conservation of the properties that define it and the operational structural correspondence in which these properties are realized.

vii) <u>Structural coupling or adaptation</u>. I call structural coupling or adaptation the relation of dynamic structural correspondence with the medium in which a unity conserves its class identity (organization in the case of a composite unity, and operation of its properties in the case of a simple one), and which is entailed in its distinction as it is brought forth by the observer in his or her praxis of living. Therefore, conservation of class identity and conservation of adaptation are constitutive conditions of existence for any unity (entity, system, whole, etc.) in the domain of existence in which it is brought forth by the observer in his or her praxis of living. As constitutive conditions of existence for any unity, conservation of class identity and conservation of adaptation are paired conditions of existence that entail each other so that if one is lost the other is lost, and the unity exists no more. When this happens, a composite unity disintegrates and a simple unity disappears.

viii) <u>Domain of existence</u>. The operation of distinction that brings forth and specifies a unity, also brings forth and specifies its domain of existence as the domain of the operational coherences entailed by the operation of the properties through which the unity is characterized in its distinction. In other words, the domain of existence of a simple unity is the domain of operational validity of the properties that define it as such, and the domain of existence of a composite unity is the domain of operational validity of the properties of the components that constitute it. Furthermore, the constitutive operational coherence of a domain of existence as the domain of operational validity of the properties of the entities that define it, entails all that such validity requires. Accordingly, a simple unity exists in a single domain of existence specified by its properties, and a composite unity exists in two--in the domain of existence specified by its properties as it is distinguished as a simple unity, and in the domain of existence specified by the properties of its components as it is distinguished as a composite unity. The entailment in the distinction of a unity of its domain of existence as the domain of all the operational coherences in the praxis of living of the observer in which it conserves class identity and adaptation, is a constitutive condition of existence of every unity. A unity cannot exist outside its domain of existence, and if we imagine a unity outside its domain of existence, the unity that we imagine exists in a different domain than the unity that we claim that we imagine.

ix) Determinism. To say that a system is deterministic is to say that it operates according to the operational coherences of its domain of existence. And this is so because due to our constitutive inability to experientially distinguish between what we socially call perception and illusion, we cannot make any claim about an objective reality. This we acknowledge by putting objectivity in parentheses. In other words, to say that a system is deterministic is to say that all its changes are structural changes that arise in it through the operation of the properties of its components in the interactions that these realize in its composition, and not through instructive processes in which an external agent specifies what happens in it. Accordingly, an operation of distinction that brings forth a simple unity brings forth its domain of existence as the domain of operational applicability of its properties, and constitutes the simple unity and its domain of existence as a deterministic system. At the same time, the operation of distinction that brings forth a composite unity brings forth its domain of existence as a domain of determinism in terms of the operational applicability of the properties that characterize its components, in the praxis of living of the observer. Accordingly, the operation of distinction that brings forth a composite unity brings forth the composite unity as well as its domain of existence, as deterministic systems in the corresponding domains of operational coherences of the praxis of living of the observer.

x) Space. The distinction of a unity brings forth its domain of existence as a space of distinctions whose dimensions are specified by the properties of the unities whose distinctions entail it as a domain of operational coherences in the praxis of living of the observer. Thus, a simple unity exists and operates in a space specified by its properties, and a composite unity exists and operates in a space specified by its properties as a simple unity if distinguished as such, and in a space specified by the properties of its components if distinguished as a composite unity. Accordingly, as a simple unity exists and operates in a single space, a composite unity exists and operates in two. Finally, it follows that without the distinction of a unity there is no space, and that the notion of a unity out of space, as well as the notion of an empty space, are nonsensical. A space is a domain of distinctions.

xi) Interactions. Two simple unities interact when they, as a result of the interplay of their properties, and in a manner determined by such interplay, change their relative position in a common space or domain of distinctions. A composite unity interacts when some of its components as a result of their interactions as simple unities with other simple unities that are not its components, change their manner of composing it, such that it undergoes a structural change. It follows that a simple unity interacts in a single space, in the space that its properties define, and that a composite unity interacts in two, in the space defined by its properties as a simple unity, and in the space that its components define through their properties, also as simple unities, as they constitute its structure.

xii) Phenomenal domains. A space is constituted in the praxis of living of the observer when he or she performs a distinction. The constitution of a space brings forth a phenomenal domain as the domain of distinctions of the relations and interactions of the unities that the observer distinguishes as populating that space. A simple unity operates in a single phenomenal domain, the phenomenal domain constituted through the operation of its properties as a simple unity. A composite unity operates in two phenomenal domains, the phenomenal domain constituted through the operation of its properties as a simple unity, and the phenomenal domain constituted through the operation of the properties of its components, which is where its composition takes place. Furthermore, the two phenomenal domains in which a composite unity operates do not intersect and cannot be reduced one to the other because there is a generative relation between them. The phenomenal domain in which a composite unity operates as a simple unity is secondary to the composition of the composite unity, and constitutes a metaphenomenal domain with respect to the phenomenal domain in which the composition takes place. Due to this circumstance, a composite unity cannot participate as a simple unity in its own composition.

xiii) Medium, niche, and environment. I call the medium of a unity the containing background of distinctions, including all that is not involved in its structure if it is a composite one, with respect to which an observer distinguishes it in his or her praxis of living, and in which it realizes its domain of existence. The medium includes both that part of the background that is distinguished by the observer as surrounding the unity, and that part of the background the observer conceives as interacting with it, and which it obscures in its operation in structural coupling (in its domain of existence). I call this latter part of the medium operationally defined moment by moment in its encounter with the medium in structural coupling, the niche of the unity. Accordingly, a unity continuously realizes and specifies its niche by actually operating in its domain of perturbations while conserving adaptation in the medium. As a consequence, the niche of a unity is not a fixed part of the medium in which a unity is distinguished, nor does it exist with independency of the unity that specifies it; it changes as the domain of interactions of the unity changes (if it is a composite one) in its dynamics of structural change (section v c). In these circumstances, an observer can distinguish the niche of a unity, regardless of whether it is simple or composite, only by using the unity as an indicator of it. Finally, I call the environment of a unity all that an observer distinguishes as surrounding it. In other words, while the niche is that part of the medium that a unity encounters (interacts with) in its operation in structural coupling, and obscures with its presence from the view of the observer, the environment is that part of the medium that an observer sees around a unity. Thus, a dynamic composite unity (like a living system), as it is distinguished in the praxis of living of the observer, is seen in an environment as an entity with a changing niche that it specifies while it slides through the medium in continuous structural change with conservation of class identity and adaptation. A composite unity in its medium is like a tightrope walker that moves on a rope in a gravitational field, and conserves its balance (adaptation) while its shape (structure) changes in a manner congruent with the visual and gravitational interactions that it undergoes as it walks (realizing its niche), and falls when this stops being the case.

7. BASIS FOR THE ANSWER: THE LIVING SYSTEM

The answer to the question of cognition requires now that we reflect upon the constitution and operation of living systems, and that we make some additional epistemological and ontological considerations about the conditions that our understanding of living systems must satisfy.

i) Science deals only with structure determined systems. To the extent that a scientific explanation entails the proposition of a structure determined system as the mechanism that generates the phenomenon to be explained, we as scientists can deal only with structure determined systems, and we cannot handle systems that change in a manner specified by the external agents that impinge upon them. Accordingly, whatever I say about living systems will be said in the understanding that all the phenomena to which they give rise arise through their operation as structure determined systems in a domain of existence also brought forth as a structure determined system by the observer's distinction.

ii) Regulation and control. As was indicated in section 6 xii, the distinction of a composite unity entails the distinction in the praxis of living of the observer of two phenomenal domains that do not intersect because the operation of a composite unity as a simple one is secondary to its composition. As a result, the whole cannot operate as its own component, and a component cannot operate in place of the whole that it integrates. In these circumstances, notions of control or regulation do not connote actual operations in the composition of a composite unity, because such operations take place only in the realization in the present of the properties of the composite unity's components in their actual interactions. Notions of regulation and control only connote relations taking place in a descriptive domain as the observer relates mappings in language of his or her distinctions of a whole and its components in his or her praxis of living.

iii) Living systems are structure determined systems. In order to explain the phenomenon of cognition as a biological phenomenon, I must treat living systems-as structure determined systems. I consider that to do so is legitimate for several reasons. I shall mention three. The first is an operational one: we know as a feature of our praxis of living that any structural change in a living system results in a change in its characteristics and properties, and that similar structural changes in different members of the same species result in similar changes in their characteristics and properties. The second is an epistemological one: if we do not treat living systems as structure determined systems we cannot provide scientific explanations for the phenomena proper to them. The third is an ontological one: the only systems that we can explain scientifically are structure determined systems; therefore, if I provide a scientific explanation of the phenomenon of cognition in living systems, I provide a proof that living systems are structure determined systems in our praxis of living as standard observers, which is where we distinguish them.

iv) Determinism and prediction. The fact that a structure determined system is deterministic does not mean that an observer should be able to predict the course of its structural changes. Determinism and predictability pertain to different operational domains in the praxis of living of the observer. Determinism is a feature that characterizes a system in terms of the operational coherences that constitute it, and its domain of existence, as it is brought forth in the operations of distinction of the observer. Accordingly, there are as many different domains of determinism as domains of different operational coherences the observer brings forth in her or his domain of experiences. At difference with this, a prediction is a computation that an observer makes of the structural changes of a structure determined system as she or he follows the consequences of the operation of the properties of the components of the system in the realization of the domain of determinism that these properties constitute. As such, a prediction can only take place after the observer has completely described the system as a structure determined system in terms of the operational coherences that constitute it in his or her domain of experiences. Therefore, the success or failure of a prediction only reflects the ability or inability of an observer to not confuse phenomenal domains in his or her praxis of living, and to indeed make the computation that constitutes the prediction in the phenomenal domain where he or she claims to make it. In these circumstances, there are two occasions in which an observer who does not confuse phenomenal domains in dealing with a structure determined system will not be able to predict its structural changes. One occasion is when an observer knows that she or he is dealing with a structure determined system by virtue of experience, in the praxis of living, with its components, but cannot encompass it in his or her descriptions, and, thus, cannot effectively treat it as such in its domain of existence and compute its changes of state. The other occasion is when an observer in his or her praxis of living aims at characterizing the present unknown state of a system assumed to be structure determined, by interacting with some of its components. By doing this the observer triggers in the system an unpredictable change of state that he or she then uses to characterize its initial state and predict in it a later one within the domain of determinism specified by the properties of its components. Therefore, since the domain of determinism of a structure determined system as the domain of operational coherences of its components is brought forth in its distinction in the praxis of living of the observer, and since in order to compute a change of state in a system the observer must determine its present state through an interaction with its components, any attempt to compute a change of state in a structure determined system entails a necessary uncertainty due to the manner of determination of its initial state within the constraints of the operational coherences of its domain of existence. This predictive uncertainty may vary in magnitude in different domains of distinctions, but it is always present because it is constitutive of the phenomenon of cognition as a feature of the ontology of observing and not of an objective independent reality. With this I am also saying that the uncertainty principle of physics pertains to the ontology of observing, and that it does not characterize an independent universe because, as I shall show further on, the physical domain of existence is a cognitive domain brought forth in the praxis of living of the observer by the observer as an explanation of his or her praxis of living.

v) Ontogenic structural drift. It is said that a boat is drifting when it slides floating on the sea without rudder and oars, following a course that is generated moment after moment in its encounter with the waves and wind that impinge upon it, and which lasts as long as it remains floating (conserves adaptation) and keeps the shape of a boat (conserves organization). As such, a drifting boat follows a course without alternatives that is deterministically generated moment after moment in its encounters with the waves and the wind. As a consequence of this, a drifting boat is also always, and at any moment, in the only place where it can be, in a present that is continuously emerging from the sequence of its interactions in the drift. The deterministic process that generates the course followed by a drifting boat takes place as a feature of the structural dynamics of the structure determined system constituted by the boat, the wind, and the waves, as these are brought forth by the observer in his or her praxis of living. Therefore, if an observer cannot predict the course of a drifting boat, it is not because his or her distinction of the boat, the wind, and the waves, in his or her domain of experiences, does not entail a structure determined system in which the course followed by the boat arises in a deterministic manner, but because he or she cannot encompass in his or her description of the interactions between the boat, the wind, and the waves, the whole structure of the structure determined system in which the course followed by the boat is a feature of its changes of structure.

What happens with the generation of the course followed by a drifting boat, is the general case for the generation of the course followed by the structural changes of any structure determined system that the observer distinguishes in his or her praxis of living, as it interacts in the medium as if with an independent entity with conservation of class identity (organization) and adaptation (structural coupling). Since living systems are dynamic structure determined systems, this applies to them; and the ontogeny of a living system, as its history of structural changes with conservation of organization and adaptation, is its ontogenic structural drift. All that applies to the course followed by a drifting boat applies to the course followed by the structural changes that take place in the ontogeny of a living system and to the course followed by the displacement of a living system in the medium during its ontogeny. Let me make this clear. In general terms, a drift is the course followed by the structural changes of a structure determined system that arise moment after moment generated in the interactions of the system with another independent system, while its relation of correspondence (adaptation) with this other system (medium) and its organization (class identity) remain invariant. Accordingly, the individual life history of a living system as a history of continuous structural changes that follows a course generated moment after moment in the braiding of its internally generated structural dynamics with the structural changes triggered in it by its recurrent interactions with the medium as an independent entity, and which lasts as long as its organization and adaptation are conserved, takes place as a structural drift. Similarly, since the course of the displacement of a living system in the medium is generated moment after moment as a result of its interactions with the medium as an independent entity while its organization and adaptation are conserved, the displacement of a living system in the medium while it realizes its niche takes place as a drift. Living systems exist in continuous structural and positional drift (ontogenic drift) while they are alive, as a matter of constitution.

As in the case of a drifting boat, at any moment a living system is where it is in the medium, and has the structure that it has, as the present of its ontogenic drift in a deterministic manner, and could not be anywhere other than where it is, nor could it have a structure different from the one that it has. The many different paths that an observer may consider possible for a drifting boat to follow at any instant, or the many different ontogenic courses that an observer may consider possible for a living system at any moment, are possible only as imagined alternatives in the description of what would happen in each case if the conditions were different, and not actual alternatives in the course of the boat or in the ontogeny of the living system. A drift is a process of change, and as is the case with all processes of change in structure determined systems, it follows a course without alternatives in the domain of determinism in which it is brought forth by the distinctions of the observer. Indeed, such imagined alternatives are imaginable only from the perspective of the inability of the observer to treat the boat, the wind, and the waves, or the living system and the medium, that he or she brings forth in his or her praxis of living, as a known structure determined system whose changes of structure he or she can compute. If we are serious about our explanations as scientists, then we must accept as an ontological feature of what we do as observers that every entity that we bring forth in our distinctions is where it is, and has the structure that it has, in the only manner that it can be, given the domain of operational coherences (domain of determinism) that we also bring forth as its domain of existence in its distinction.

Finally, let me mention several implications of all this for the entities that we bring forth as living systems in our praxis of living: a) Since for a living system a history of interactions without disintegration can only be a history of perturbations, that is, a history of interactions in the niche, a living system while living necessarily slides in ontogenic drift through the medium in the realization of its niche. This means that aim, goal, purpose, or intention, do not enter into the realization of a living system as a structure determined system. b) Since the structure of a living system is continuously changing, both through its internal dynamics and through the structural changes triggered in it in its interactions with operationally independent entities, the niche of a living system (the features of the medium that it actually encounters in its interactions) is necessarily in continuous change congruent with the continuous structural drift of the living system while it remains alive. Furthermore, this is so regardless of whether the observer considers that the environment of the living system changes or remains constant. This means that as an observer brings forth a living system in her or his praxis of living, it may appear to her or him as continuously changing in its use of a constant environment, or, conversely, as unchanging in a continuously changing environment, because the observer cannot see the encounter of a living system and its niche, which is where conservation of adaptation takes place. c) Conservation of adaptation does not mean that the manner of living of a living system remains invariant. It means that a living system has an ontogeny only while it conserves its class identity and its dynamic structural correspondence with the medium as it undergoes its interactions, and that there is no constitutive restriction about the magnitude of its moment after moment structural changes other than that they should take place within the constraints of its structural determinism and its conservation of organization and adaptation. Indeed, I could speak of the laws of conservation of organization and adaptation as ontological conditions for the existence of any structure determined system in the same manner as physicists speak of the laws of conservation in physics as ontological conditions for the occurence of physical phenomena.

Every living system, including us observers, is at any moment where it is, has the structure that it has, and does what it does at that moment, always in a structural and relational situation that is the present of an ontogenic drift that starts at its inception as such in a particular place with a particular structure, and follows the only course that it can follow. Different kinds of living systems differ in the spectrum of ontogenies that an observer can consider possible for each of them in his or her discourse as a result of their different initial structures and different starting places, but each ontogeny that takes place takes place as a unique ontogenic drift in a process without alternatives.

vi) Structural intersection. When an observer brings forth a composite unity in his or her praxis of living, he or she brings forth an entity in which the configuration of relations between components that constitutes its organization, is a subset of all the actual relations that take place between its components as these realize its structure and constitute it as a whole in the domain of existence in which they are brought forth (see section 6, iv). As such, the organization of a composite unity does not exhaust the relations and interactions in which the components that realize it may participate in their domain of existence. The result of this circumstance is that in the structural realization of a composite unity, its components may participate, through other properties than those that involve them in the realization of its organization, in the realization of the organization of many other composite unities which, thus, intersect structurally with it. Furthermore, when the components of a composite unity are themselves composite unities, the composite unity may participate in structural intersections that take place through the components of its components. In any case, when an observer distinguishes two or more structurally intersecting systems, he or she distinguishes two or more different composite unities realized through the same body.

Structurally intersecting systems exist and operate as simple unities in different phenomenal domains specified by their different organizations. Yet, depending on how their structural intersection takes place, structurally intersecting composite unities may exist as such in the same or different domains of existence. Thus, when two composite unities structurally intersect through their components, they share components and have as composite unities the same domain of existence. But, when two composite unities structurally intersect through the components of the components of one or both, they do not share components and as composite unities have different domains of existence. Nevertheless, since in a structural intersection there are components or components of components, or both, that simultaneously participate in the structure of several systems, structural changes that take place in one of several structurally intersecting systems as part of its ontogenic drift may give rise to structural changes in the other intersecting systems and thus participate in their otherwise independent ontogenic drifts. In other words, structurally intersecting systems are structurally interdependent because, either through the intersection of their domains of structural determinism, or though the intersection of the domains of structural determinism of their components, or through both, they affect each other's structures in the course of their independently generated structural changes, and although they may exist as composite unities in different domains, their ontogenic drifts intersect forming a network of coontogenic drifts. Thus, an observer may distinguish in the structural realization of a human being as a living system the simultaneous or successive intersection of a mammal, a person, a woman, a doctor, and a mother, all of which are different composite unities defined by different organizations that are simultaneously or successively conserved while they are realized in their different domains of existence, with particular characteristics that result from the continuous braiding of their different ontogenic drifts through the continuous interplay of their structural changes. Furthermore, these structural intersections result in dependent domains of disintegrations as well as dependent domains of conservations which need not be reciprocal, when the conservation of one class identity entails the conservation of structural features that are involved in the conservation of another. For example, in the structural intersection of a student and a human being in a living system, the conservation of the class identity "student" entails the conservation of the class identity "human being," but not the reverse: the disintegration of the student does not entail the disintegration of the human being, but the disintegration of the human being carries with it the disintegration of the student. Also, a particular composite unity may disintegrate through different kinds of structural changes, like disintegrating as a student through failing an examination or through attaining the final degree, with different consequences in the network of structural intersections to which it belongs.

The structural intersection of systems does not mean that the same system is viewed in different manners from different perspectives, because due to their different organizations structurally intersecting systems exist in different phenomenal domains and are realized through different structural dynamics. It only means that the elements that realize a particular composite unity as its components through some of their properties as simple unities, participate through other of their properties as simple unites as components of other unities that exist as legitimately different ones because they have different domains of disintegrations. The interactions and relations in which the components of a system participate through dimensions other than those through which they constitute it, I call orthogonal interactions and relations, and it is through these that structurally intersecting systems may exist in nonintersecting phenomenal domains and yet have unidirectional or reciprocal relations of structural dependency. Finally, it is also through the orthogonal interactions of their components that structurally independent systems that exist in nonintersecting phenomenal domains may also have coontogenic drifts.

vii) The living system. In 1970 I proposed that living systems are dynamic systems constituted as autonomous unities through being closed circular concatenations (closed networks) of molecular productions in which the different kinds of molecules that compose them participate in the production of each other, and in which everything can change except the closed circularity of the concatenation of molecular productions that constitutes them as unities (see Maturana 1970, in Maturana and Varela 1980). In 1973 Francisco Varela and I expanded this characterization of living systems by saying: first, that a composite unity whose organization can be described as a closed network of productions of components that through their interactions constitute the network of productions that produce them and specify its extension by constituting its boundaries in their domain of existence, is an autopoietic system; and second, that a living system is an autopoietic system whose components are molecules. Or, in other words, we proposed that living systems are molecular autopoietic systems and that as such they exist in the molecular space as closed networks of molecular productions that specify their own limits (see Maturana and Varela 1973, in Maturana and Varela 1980; and Maturana 1975). Nothing is said in this description of the molecular constitution of living systems as autopoietic systems about thermodynamic constraints, because the realization of living systems as molecular systems entails the satisfaction of such constraints. In fact, the statement that a composite unity exists as such in the domain of existence of its components, implies the satisfaction of the conditions of existence of these components.

The recognition that living systems are molecular autopoietic systems carries with it several implications and consequences of which I shall mention a few:

A. Implications: a) Living systems as autopoietic systems are structure determined systems, and everything that applies to structure determined systems applies to them. In particular this means that everything that occurs in a living system takes place in it in the actual operation of the properties of its components through relations of neighborhood (relations of contiguity) constituted in these very same operations. Accordingly, notions of regulation and control do not and cannot reflect actual operations in the structural realization of a living system because they do not connote actual relations of neighborhood in it. These notions only reveal relations that the observer establishes when he or she compares different moments in the course of transformations in the network of processes that take place in the structural realization of a particular living system. Therefore, the only peculiar thing about living systems as structure determined systems is that they are molecular autopoietic systems. b) Autopoiesis is a dynamic process that takes place in the ongoing flow of its occurrence and cannot be grasped in a static instantaneous view of distribution of components. Therefore, a living system exists only through the continuous structural transformation entailed in its autopoiesis, and only while this is conserved in the constitution of its ontogeny. This circumstance has two basic results: one is that living systems can be realized through many different changing dynamic structures, the other is that in the generation of lineages through reproduction, living systems are constitutively open to continuous phylogenic structural change. c) A living system either exists as a dynamic structure determined system in structural coupling in the medium in which it is brought forth by the observer, that is, in a relation of conservation of adaptation through its continuous structural change in the realization of its niche, or it does not exist. Or, in other words, a living system while living is necessarily in a dynamic relation of correspondence with the medium through its operation in its domain of existence, and to live is to glide through a domain of perturbations in an ontogenic drift that takes place through the realization of an ever changing niche. d) A living system as a structure determined system operates only in the present--that is, it is determined by the structure that it has at any instant in the structural realization of its autopoiesis in the molecular space--and therefore it is necessarily open to the flow of molecules through it. At the same time, a living system as an autopoietic system gives rise only to states in autopoiesis; otherwise it disintegrates. Therefore, living systems are closed systems with respect to their dynamics of states.

B. Consequences: a) To the extent that a living system is a structure determined system, and everything in it takes place through neighborhood relations between its components in the present, notions of purpose and goal that imply that at every instant a later state of a system as a whole operates as part of its structure in the present do not apply to living systems and cannot be used to characterize their operation. A living system may appear to operate as a purposeful or goal-directed system only to an observer who, having seen the ontogeny of other living systems of the same kind in the same circumstances in his or her praxis of living, confuses phenomenal domains by putting the consequences of its operation as a whole among the processes that constitute it. b) Because they are structure determined systems, for living systems there is no inside or outside in their operation as autopoietic unities; they are in autopoiesis as closed wholes in their dynamics of states, or they disintegrate. At the same time, and for the same reason, living systems do not use or misuse an environment in their operation as autopoietic unities, nor do they commit mistakes in their ontogenic drifts. In fact, a living system in its operation in a medium with conservation of organization and adaptation as befit it as a structure determined system, brings forth its ever changing niche as it realizes itself in its domain of existence, the background of operational coherences which it does not distinguish and with which it does not interact. c) Living systems necessarily form, through their recurrent interactions with each other as well as with the nonbiotic medium, coontogenic and cophylogenic systems of braided structural drifts that last as long as they conserve their autopoiesis through the conservation of their reciprocal structural couplings. Such is biological evolution. As a result, every living system, including us human beings as observers, is always found in its spontaneous realization in its domain of existence in congruence with a biotic and a nonbiotic medium. Or, in other words, every living system is at every instant as it is and where it is a node of a network of coontogenic drifts that necessarily involves all the entities with which it interacts in the domain in which it is brought forth by the observer in his or her praxis of living. As a consequence, an observer as a living system can only distinguish an entity as a node of the network of coontogenic drifts to which it belongs, and where it exists in structural coupling. d) The only thing peculiar to living systems is that they are autopoietic systems in the molecular space. In these circumstances, a given phenomenon is a biological phenomenon only to the extent that its realization entails the realization of the autopoiesis of at least one autopoietic system in the molecular space. e) Modern prokaryotic and eukaryotic cells are typical autopoietic systems in the molecular space, and because their autopoiesis is not the result of their being composed by more basic autopoietic subsystems, I call them first order autopoietic systems. I call second order autopoietic systems systems whose autopoiesis is the result of their being composed of more basic autopoietic unities; organisms as multicellular systems are such. Yet, organisms may also "be," and I think that most of them actually are, first order autopoietic systems as closed networks of molecular productions that involve intercellular processes as much as intracellular ones. Accordingly, an organism would exist as such in the structural intersection of a first order autopoietic system with a second order one, both realized through the autopoiesis of the cells that compose the latter. This happened originally with the eukaryotic cell as this arose through the endosymbiosis of prokaryotic ones (Margulis 1981). f) An organism as a second order autopoietic system is an ectocellular symbiont composed of cells, usually of common origin but not always so, that constitute it through their coontogenic drift. An organism as a first order autopoietic system, however, is not composed of cells even though its realization depends on the realization of the autopoiesis of the cells that intersect structurally with it as they constitute it in their coontogenic drift. The first and second order autopoietic systems that intersect structurally in the realization of an organism, exist in different nonintersecting phenomenal domains.

viii) Phylogenic structural drift. Reproduction is a process in which a system gives origin through its fracture to two systems characterized by the same organization (class identity) that characterized the original one, but with structures that vary with respect to it (Maturana 1980). A reproductive phylogeny or lineage, then, is a succession of systems generated through sequential reproductions that conserve a particular organization. Accordingly, each particular reproductive lineage or phylogeny is defined by the particular organization conserved through the sequential reproductions that constitute it. Therefore, a reproductive phylogeny or lineage lasts only as long as the organization that defines it is conserved, regardless of how much the structure that realizes this organization in each successive member of the lineage changes with each reproductive step (see Maturana 1980, and Maturana and Varela 1987). It follows that a reproductive phylogeny or lineage as a succession of ontogenic drifts, constitutively occurs as a drift of the structures that realize the organization conserved along it. It also follows that each of the reproductive steps that constitute a reproductive phylogeny is the occasion that opens the possibility for a discrete, large or small, change in the course of its structural drift. As such, a reproductive phylogeny or lineage comes to an end through the structural changes of its members. And this occurs either because autopoiesis is lost after the last of them, or because through the conservation of autopoiesis in the offspring of the last of them, a particular set of relations of the drifting structure begins to be conserved through the following sequential reproductions as the organization that defines and starts a new lineage. This has several general implications of which I shall mention only a few: a) A member of a reproductive phylogeny either stays in structural coupling (conserves adaptation) in its domain of existence until its reproduction, and the phylogeny continues, or it disintegrates before then and the phylogeny ends with it. b) A living system is a member of the reproductive phylogeny in which it arises only if it conserves through its ontogeny the organization that defines the phylogeny, and continues the phylogeny only if such organization is conserved through its reproduction. c) Many different reproductive phylogenies can be conserved operationally embedded in each other, forming a system of nested phylogenies, if there is an intersection of the structural realization of the different organizations that define them. When this happens there is always a fundamental reproductive phylogeny whose realization is necessary for the realization of all the others. This has occurred in the evolution of living systems in the form of the phylogenic drift of a system of branching nested reproductive phylogenies in which the fundamental reproductive phylogeny is that in which autopoiesis is conserved (see Maturana 1980, and Maturana and Varela 1987). Thus, the system of branching phylogenies defined by the conservation of autopoiesis through reproductive cells in eukaryotic organisms, has carried embedded in it, through the structural intersection of their realizations, many staggered nested organizations that characterize the coincident lineages conserved through it. This circumstance we recognize in the many nested taxonomic categories that we distinguish in any organism when we classify it. For example, a human being is a vertebrate, a mammal, a primate, a _Homo_, and a _Homo sapiens_--all different categories corresponding to different systems of partially overlapping phylogenies that are conserved together through the conservation of the human being's autopoiesis. d) The ontogenic drifts of the members of a reproductive phylogeny take place in reciprocal structural coupling with many different, and also continuously changing, living and nonliving systems that form part of the medium in which they realize their niches. As a result, every individual ontogeny in living systems follows a course embedded in a system of coontogenies that constitutes a network of cophylogenic structural drifts. This can be generalized by saying that evolution is constitutively a coevolution, and that every living system is at any moment where it is, and has the structure that it has, as an expression of the present state of the domain of operational coherences constituted by the network of cophylogenic structural drifts to which it belongs. As a result, the operational coherences of every living system at every instant necessarily entail the operational coherences of the whole biosphere. e) The observer as a living system is not an exception to all that has been said above. Accordingly, an observer can only make distinctions that, as operations in his or her praxis of living, take place as operations within the present state of the domain of operational coherences constituted by the network of coontogenic and cophylogenic structural drifts to which he or she belongs.

ix) Ontogenic possibilities. The ontogeny of every structure determined system starts with an initial structure that is the structure that realizes the system at the beginning of its existence in its inception. In living systems such initial structure is a cellular unity that may originate either a) as a single cell or as a small multicellular entity through a reproductive fracture from a cellular maternal system whose organization it conserves, or b) as a single cell *de novo* from noncellular elements. In every living system the system's initial structure constitutes the structural starting point that specifies in it what an observer sees as the configuration of all the courses of ontogenic drifts that it may undergo under different circumstances of interactions in the medium. As a result, what constitutes a lineage in living systems is the conservation through their reproduction of a particular initial structure that specifies a particular configuration of possible ontogenic drifts; and what constitutes the organization conserved through reproduction that specifies the identity of the lineage is that configuration. Accordingly, a lineage comes to an end when the configuration of possible ontogenic drifts that defines it stops being conserved. The configuration of possible ontogenic drifts that specifies a lineage through its conservation I call the ontogenic phenotype of the lineage. In each particular living system, however, only one of the ontogenic courses deemed possible in the ontogenic phenotype by the observer, is realized as a result of its internal dynamics under the contingencies of the particular perturbations that it undergoes in its domain of existence with conservation of organization and adaptation. Consequently, and in general, it is only within the domain of possibilities set by their different or similar initial structures that different composite unities may have different or similar ontogenic structural drifts under different or similar histories of perturbations in their domains of existence. Indeed, nothing can happen in the ontogeny of a living system as a composite unity that is not permitted in its initial structure. Or, in other words, and under the understanding that the initial structure of a living system is its genetic constitution, it is apparent that nothing can happen in the ontogenic structural drift of a living system that is not allowed in its genetic constitution as a feature of its possible a living system, conservation of living (conservation of autopoiesis and of adaptation) constitutes adequate action in those circumstances, and, hence, knowledge: _living systems are cognitive systems, and to live is to know_. But, by showing this I have also shown that any interaction with a living system can be viewed by an observer as a question posed to it, as a challenge to its life that constitutes a domain of existence where he or she expects adequate action of it. And, at the same time, I have also shown, then, that the actual acceptance by the observer of an answer to a question posed to a living system, entails his or her recognition of adequate action by the living system in the domain specified by the question, and that this recognition of adequate action consists in the distinction of the living system in that domain under conditions of conservation of autopoiesis and adaptation. In what follows I present this general explanatory proposition under the guise of a particular scientific explanation:

a) The _phenomenon to be explained_ is adequate action by a living system at any moment in which an observer distinguishes it as a living system in action in a particular domain. And I propose this as the phenomenon to be explained in the understanding that the adequate actions of a living system are its interactions with conservation of class identity in the domain in which it is distinguished.

b) Given that structural coupling in its domain of existence (conservation of adaptation) is a condition of existence for any system distinguished by an observer, _the generative mechanism for adequate action_ in a living system as a structurally changing system, is the structural drift with conservation of adaptation through which it stays in continuous adequate action while it realizes its niche, or disintegrates. Since a system is distinguished only in structural coupling, when an observer distinguishes a living system he or she necessarily distinguishes it in adequate action in the domain of its distinction, and distinguishes it as a system that constitutively remains in structural coupling in its domain of existence regardless of how much its structure, or the structure of the medium, or both, change while it stays alive.

c) Given the generative mechanism proposed in (b), _the following phenomena can be deduced_ to take place in the domain of experiences of an observer: i) the observer should see adequate action taking place in the form of coordinated behavior in living systems that are in coontogenic structural drift while in recurrent interactions with conservation of reciprocal adaptation; ii) the observer should see that living systems in coontogeny separate or disintegrate, or both, when their reciprocal adaptation is lost.

d) _The phenomena deduced in_ (c) are apparent in the domain of experiences of an observer in the dynamics of constitution and realization of a social system, and in all circumstances of recurrent interactions between living systems during their ontogenies, in what appears to us as learning to live together. One of these cases is our human operation in language.

The satisfaction of these four conditions results: a) in the validation, as a scientific explanation, of my proposition that cognition as adequate action in living systems is a consequence of their structural drift with conservation of organization and adaptation; b) in showing that adequate action (cognition) is constitutive to living systems because it is entailed in their existence as such; c) in entailing that different living systems differ in their domains of adequate actions (domains of cognition) to the extent that they realize different niches; and d) in showing that the domain of adequate actions (domain of cognition) of a living system changes as its structure, or the structure of the medium, or both, change while it conserves organization and adaptation.

At the same time, it is apparent from all of the above that what I say of cognition as an explanation of the praxis of living takes place in the praxis of living, and that to the extent that what I say is effective action in the generation of the phenomenon of cognition, what I say takes place as cognition. If this sounds strange, it is only because we are in the habit of thinking about cognition in the explanatory pathway of objectivity without parentheses, as if the phenomenon connoted by the word cognition entailed pointing to something whose existence can be asserted to be independent of the pointing of the observer. I have shown that this is not and cannot be the case. Cognition cannot be understood as a biological phenomenon if objectivity is not put in parentheses, nor can it be understood as such if one is not willing to follow all the consequences of such an epistemological act.

Let us now treat human operation in language as one of the phenomena which take place as a consequence of the operation of cognition as adequate (or effective) action. It is particularly necessary to proceed in this manner because our operation in language as observers in the praxis of living is, at the same time, our problem and our instrument for analysis and explanation.

ii) Language. We human beings are living systems that exist in language. This means that although we exist as human beings in language and although our cognitive domains (domains of adequate actions) as such take place in the domain of languaging, our languaging takes place through our operation as living systems. Accordingly, in what follows I shall consider what takes place in language as language arises as a biological phenomenon from the operation of living systems in recurrent interactions with conservation of organization and adaptation through their coontogenic structural drift, and thus show language as a consequence of the same mechanism that explains the phenomenon of cognition:

a) When two or more autopoietic systems interact recurrently, and the dynamic structure of each follows a course of change contingent upon the history of each's interactions with the others, there is a coontogenic structural drift that gives rise to an ontogenically established domain of recurrent interactions between them which appears to an observer as a domain of consensual coordinations of actions or distinctions in an environment. This ontogenically established domain of recurrent interactions I call a domain of consensual coordinations of actions or distinctions, or, more generally, a consensual domain of interactions, because it arises as a particular manner of living together contingent upon the unique history of recurrent interactions of the participants during their coontogeny. Furthermore, because an observer can describe such a domain of recurrent interactions in semantic terms, by referring the different coordinations of actions (or distinctions) involved to the different consequences that they have in the domain in which they are distinguished, I also call a consensual domain of interactions a linguistic domain. Finally, I call the behavior through which an organism participates in an ontogenic domain of recurrent interactions, consensual or linguistic according to whether I want to emphasize the ontogenic origin of the behavior (consensual), or its implications in the present state of the ongoing interactions (linguistic). Similarly, I speak of coordinations of actions or coordinations of distinctions, according to whether I want to emphasize what takes place in the interaction in relation to the participants (coordinations of actions), or what takes place in the interactions in relation to an environment (coordinations of distinctions).

b) When one or more living systems continue their coontogenic structural drift through their recurrent interactions in a consensual domain, it is possible for a recursion to take place in their consensual behavior resulting in the production of a consensual coordination of consensual coordinations of actions. If this were to happen, what an observer would see would be that the participants of a consensual domain of interactions would be operating in their consensual behavior making consensual distinctions upon their consensual distinctions, in a process that would recursively make a consensual action a consensual token for a consensual distinction that it obscures. Indeed, this process is precisely what takes place in our languaging in the praxis of living. Accordingly, I claim that the phenomenon of language takes place in the coontogeny of living systems when two or more organisms operate, through their recurrent ontogenic consensual interactions, in an ongoing process of recursive consensual coordinations of consensual coordinations of actions or distinctions (Maturana, 1978). Or, in other words, I claim that such recursive consensual coordination of consensual coordinations of actions or distinctions in any domain, is the phenomenon of language. Furthermore, I claim that objects arise in language as consensual coordinations of actions that operationally obscure for further recursive consensual coordinations of actions by the observers the consensual coordinations of actions (distinctions) that they coordinate. Objects are, in the process of languaging, consensual coordinations of actions that systems. This is arbitrary since what I have said in relation to existence applies to every entity brought forth through an operation of distinction. Therefore, I make this distinction only because I am speaking of living systems and the word cognition is historically bound to them through us. Within this restriction we as observers can say that there are as many domains of cognition as there are domains of existence specified by the different identities that living systems conserve through the realization of their autopoiesis. These different cognitive domains intersect in the structural realization of a living system as the living system realizes the different identities that define them as different dimensions of simultaneous or successive structural couplings, orthogonal to the fundamental structural coupling in which the living system realizes its autopoiesis. As a result, these different cognitive domains may appear or disappear simultaneously or independently according to whether the different structurally intersecting unities that specify them integrate or disintegrate independently or simultaneously (see section 7 vi). Thus, when a student graduates, the cognitive domain specified by the operation in the domain of structural coupling that defines the identity "student" disappears together with the disintegration of the student, or, when a bachelor marries, the cognitive domain that the identity "bachelor" defines as a domain of operational coherences in structural coupling, disappears together with the disintegration of the bachelor. Conversely, when a student graduates and a bachelor marries, the identities "graduate" and "husband" appear with the corresponding cognitive domains specified by the operational coherences that these identities entail.

It follows, therefore, that a living system may operate in as many different cognitive domains as there are different identities that the different dimensions of its structural coupling allow it to realize. It also follows that the different identities that a living system may realize are necessarily fluid, and change as the dimensions of its structural coupling change with its structural drift in the happening of its living. To have an identity, to operate in a particular domain of cognition, is to operate in a particular domain of structural coupling.

iii) <u>Language is the human cognitive domain</u>. Human beings as living systems operating in language operate in a domain of recursive reciprocal consensual perturbations that constitutes their domain of existence as such. Therefore, language as a domain of recursive consensual coordinations of actions is a domain of existence, and, as such, a cognitive domain defined by the recursion of consensual distinctions in a domain of consensual distinctions. Furthermore, human beings as living systems operating in language constitute observing, and become observers, by bringing forth objects as primary consensual coordinations of actions distinguished through secondary consensual coordinations of actions in a process that obscures the actions that they coordinate. Human beings, therefore, exist in the domain of objects that they bring forth through languaging. At the same time, human beings by existing as observers in the domain of objects brought forth through languaging, exist in a domain that allows them to explain the happening of their living in language through reference to their operation in a domain of dynamic reciprocal structural coupling.

iv) Objectivity. Objects arise in language as consensual coordinations of actions that in a domain of consensual distinctions are tokens for more basic coordinations of actions, which they obscure. Without language and outside language there are no objects, because objects only arise as consensual coordinations of actions in the recursion of consensual coordinations of actions that languaging is. For living systems that do not operate in language, there are no objects; or in other words, objects are not part of their cognitive domains. Since we human beings are objects in a domain of objects that we bring forth and operate upon in language, language is our peculiar domain of existence and our peculiar cognitive domain. Within these circumstances, objectivity arises in language as a manner of operating with objects without distinguishing the actions that they obscure. In this manner of operating, descriptions arise as concatenations of consensual coordinations of actions that result in further consensual coordinations of actions which, if performed without distinguishing how objects arise, can be distinguished as manners of languaging that take place as if objects existed outside of language. Objects are operational relations in languaging.

v) Languaging: operation in a domain of structural coupling. To the extent that language arises as a consensual domain in the coontogenic structural drift of living systems involved in recurrent interactions, the organisms that operate in language operate in a domain of reciprocal coontogenic structural coupling through reciprocal structural perturbations. Therefore, to operate in language is not an abstract activity, as we usually think. To language is to interact structurally. Language takes place in the domain of relations between organisms in the recursion of consensual coordinations of consensual coordinations of actions, but at the same time language takes place through structural interactions in the domain of the bodyhoods of the languaging organisms. In other words, although languaging takes place in the social domain as a dance of recursive relations of coordinations of actions, interactions in language as structural interactions are orthogonal to that domain, and as such trigger in the bodyhoods of the participants structural changes that change as much the physiological background (emotional standing) on which they continue their languaging, as the course that this physiological change follows. The result is that the social coordinations of actions that constitute languaging, as elements of a domain of recursive operation in structural coupling, become part of the medium in which the participant living systems conserve organization and adaptation through the structural changes that they undergo contingent to their participation in that domain. Thus, although the domain of coordinations of actions and the domain of structural change of the participants in language do not intersect, their changes are coupled orthogonally through the structural interactions that take place in language. As the body changes, languaging changes; and as languaging changes, the body changes. Here resides the power of words. Words are nodes in coordinations of actions in languaging and as such they arise through structural interactions between bodyhoods; it is through this interplay of coordinations of actions and changes of bodyhood that the world that we bring forth in languaging becomes part of the domain in which our ontogenic and phylogenic structural drifts take place.

vi) Language is a domain of descriptions. Language is a system of recursive consensual coordinations of actions in which every consensual coordination of actions becomes an object through a recursion in the consensual coordinations of actions, in a process that becomes the operation of distinction that distinguishes it and constitutes the observer. In these circumstances, all participants in a language domain can be observers with respect to the sequences of coordinations of actions in which they participate, constituting a system of recursive distinctions in which systems of distinctions become objects of distinction. Such recursive distinctions of distinctions in the happening of living in language that bring forth systems of objects, constitute the phenomenon of description. As a result, all that there is in the human domain are descriptions in the happening of living in language which, as happenings of living in language, become objects of descriptions in language. Descriptions, however, do not replace the happening of living that they constitute as descriptions; they only expand it in recursions that follow its operational coherences. Accordingly, scientific explanations, as systems of descriptions, do not replace the phenomena that they explain in the domain of happening of living of the observer, but bring forth operational coherences in that domain that allow for further descriptions in it.

vii) Self-consciousness arises with language. For a living system in its operation as a closed system, there is no inside or outside; it has no way of making the distinction. Yet, in language such a distinction arises as a particular consensual coordination of actions in which the participants are recursively brought forth as distinctions of systems of distinctions. When this happens, self-consciousness arises as a domain of distinctions in which the observers participate in the consensual distinctions of their participations in language through languaging. It follows from this that the individual exists only in language, that the self exists only in language, and that self-consciousness as a phenomenon of self distinction takes place only in language. Furthermore, it also follows that since language as a domain of consensual coordinations of actions is a social phenomenon, self-consciousness is a social phenomenon, and as such it does not take place within the anatomical confines of the bodyhood of the living systems that generate it; on the contrary, it is external to them and pertains to their domain of interactions as a manner of coexistence.

viii) History. The significance or meaning of any given behavior resides in the circumstances of its enaction, not in the characteristics of the dynamics of states of the behaving living system or in any particular feature of the behavior itself. In other words, it is not the complexity of the inner states of a living system or of its nervous system, nor any aspect of the behavior itself, that determines the nature, meaning, relevance, or content of any given behavior, but rather its placement in the ongoing historical process in which it arises. The higher human functions do not take place in the brain: language, abstract thinking, love, devotion, reflection, rationality, altruism, etc., are not features of the dynamics of states of the human being as a living system or of its nervous system as a neuronal network; they are social historical phenomena. At the same time, history is not part of the dynamics of states of a living system because this latter takes place only in the present, instant after instant, in the operation of its structure in changes that occur out of time. History, time, future, and past--as well as space--exist in language as forms of explanation of the happening of living of the observer, and partake of the involvement of language in this happening of living. Therefore, it is in the explanation of the happening of living through the coherences of language that an observer can claim that the structure of a living system that determines its changes of state in the present always embodies its history of interactions because it continuously arises in the present in a structural drift contingent to such history.

ix) The nervous system expands the domain of states of the living system. For living systems to operate in language, the diversity and plasticity of their internal states must match the diversity of the changing circumstances generated in their recursive consensual coordinations of actions. In other words, although language does not take place within the bodyhood of the living system, the structure of the living system must provide the diversity and plasticity of states required for language to take place. The nervous system contributes to the fulfillment of this requirement by expanding the domain of states of the organism through the richness of its dynamics as a closed network of changing relations of neuronal activities (see Maturana 1983), and by expanding in the organism the domain of its changes of states that follow in it a course contingent upon both its own changes of states and its interactions in the medium. And this the nervous system does: a) by admitting the interactions of the organism as orthogonal perturbations from the medium, a condition that makes its structural drift as a cellular network, as well as the structural drift of the organism and its participation in the generation of behavior, contingent upon the history of those interactions; and b) by admitting orthogonal interactions from the components of the organism, a condition that makes its structural drift as a cellular network, as well as the structural drift of the organism and its participation in the generation of behavior, recursively contingent upon the dynamics of structural changes of the organism. The result of all this for the organism (including its nervous system) is the possibility of the recursive involvement of its dynamics of states with the ongoing flow of its own dynamics of states through its behavior, if it has sufficient plasticity in the nervous system and participates in a sufficiently large domain of recurrent interactions with other organisms. Indeed, this recursive involvement is what permits the production of language as this arises when the internal recursiveness of the dynamics of states of the nervous system couples with the recurrence of social consensual coordinations of actions, giving rise to the recursion of consensual coordinations as an ongoing process in the generation of social behavior.

The ongoing recursive coupling of behavioral and structural changes that gives origin to language, is possible because a structure determined system exists in two nonintersecting phenomenal domains realized through orthogonally dependent structures, namely, its domain of states and its domain of interactions. It is our basic double existence as structure determined systems in two nonintersecting but orthogonally coupled phenomenal domains that permits us in our operation in language to generate endless orthogonally interdependent and yet nonintersecting phenomenal domains in the happening of our living.

x) Observing takes place in languaging. The nervous system is a closed network of interacting active neuronal elements (neurons, effectors, and receptors) that are structurally realized as cellular components of the organism. As such, it operates as a closed network of changing relations of activity between its components; that is, it is constitutive to the organization of the nervous system that any change of relations of activity between its components leads to further changes of relations of activity between them, and that in that sense it operates without inputs or outputs. Therefore, any action upon an environment that an observer sees as a result of the operation of the nervous system, is a feature of the structural changes that take place in the nervous system as a cellular network, and not a feature of its operation as such. Indeed, the operation of the nervous system and the actions of the organism take place in nonintersecting phenomenal domains realized by orthogonally related structures. Similarly, any perturbation of the medium impinging upon the organism is a perturbation in the structure of the nervous system, not an input into the nervous system's dynamics of states, and if this dynamics of states changes, it does so because the structure of the nervous system changes in a manner contingent to the perturbation, not because it admits an input to its operation. As a result, all that takes place in the nervous system is a dance of changing relations of neuronal activities that in the domain of structural coupling where the observer beholds the organism appears as a dance of changing configurations of effector-sensor correlations. An observer that sees an effector-sensor correlation as an adequate behavior does so because he or she beholds the organism in the domain of structural coupling in which the distinguished behavior takes place in the flow of its conservation of adaptation. The organism in its operation does not act upon an environment, nor does the nervous system operate with a representation of an environment in the generation of the adequate behavior of the organism; the environment exists only for an observer (see section 6, paragraph xiii), and as such it is a phenomenon of languaging.

That the nervous system should operate as a closed network of changing relations of activity between its components, and not with representations of an environment, has two fundamental consequences: a) For the operation of the nervous system, everything is the same. Or, in other words, all that takes place in the operation of the nervous system are changes of relations of activity between its components, and it does not distinguish in its operation whether its changes of state arise through its internal dynamics or as a result of structural changes triggered in it through what an observer sees as external structural perturbations. b) For the observer, the organism operates in many different domains of structural coupling which intersect operationally in the domain of states of the nervous system through the structural perturbations triggered in it by the interactions of the organism in these different domains. As a result of this circumstance, several things happen that are relevant for the understanding of the domains of reality that the observer brings forth (see the following sections). Firstly, an observer can always treat a state of activity of the nervous system (a configuration of changes of relations of activity) that arises as a result of a particular interaction of the organism, as a representation of that interaction, and can do so by constituting the domain of descriptions as a meta phenomenal domain in which both the organism and the circumstances of its interactions are distinguished together. Secondly, different states of activity of the nervous system that for an observer represent interactions of the organism in nonintersecting phenomenal domains (different domains of structural coupling), can affect each other and give rise to behaviors of the organism that constitute meta domains of relations between the phenomena that take place in those nonintersecting phenomenal domains. Thirdly, the meta domains of relations established through their operational intersection in the domain of states of the nervous system of otherwise nonintersecting phenomena that arise in the operation of the organism in its different domains of structural coupling, constitute, through the behaviors that these intersections generate, new domains of structural coupling of the organism that do not intersect with the others. And, fourthly, the operational intersection of the different domains of interactions (different domains of structural coupling) of an organism in the operation of its nervous system, allows it to operate in recurrent interactions with other organisms in the continuous recursive generation of meta domains of relations which become phenomenal domains in their own right in the ongoing flow of those recurrent interactions. The result of all this intersection of domains of relations in the closed operation of the nervous system through its coupling to the interactions of the organism, is the possibility of the arising of self observing as the closed operation of the nervous system becomes recursive when it couples to the dyanmics of observing as two or more organisms generate a recursive domain of coordinations of actions. That is, the operation of the nervous system as a closed network of interactions permits observing and the observer to arise as operations in language brought forth through the operational coherences of languaging. Or, in other words, since the operation of the nervous system appears in the domain of operation of the organism as sensory-effector correlations, observing is coordinations of bodyhoods of observers through their generation of a choreography of interlaced sensory-effector correlations, because all that there is for the operation of the nervous system of the observer in observing is its closed dynamics of changing relations between its neuronal components. It is only for an observer who sees two or more interacting organisms in his or her praxis of living, that the sensory-effector correlations of these organisms appear recursively involved with each other in a network of recursive sensory-effector correlations constituted through the orthogonal interactions of their nervous systems. And, finally, it is only for an observer that such a network of recursive sensory-effector correlations becomes language and constitutes a meta domain (with respect to the operation of the nervous system) where explanations and observing take place, when the organism's recurrent interactions become a recursive system of consensual coordinations of consensual coordinations of actions.

10. THE DOMAIN OF PHYSICAL EXISTENCE

A domain of existence is a domain of operational coherences entailed by the distinction of a unity by an observer in his or her praxis of living. As such, a domain of existence arises as the domain of the operational validity of the properties of the unity distinguished if it is a simple unity, or as the domain of validity of the properties of the components of the unity distinguished if the unity distinguished is a composite one. As a consequence, the distinction of a unity entails its domain of existence as a composite unity that includes the distinguished unity as a component. Therefore, there are as many domains of existence as kinds of unities an observer may bring forth in his or her operations of distinction. In these circumstances, since the notion of determinism applies to the operation of the properties of the components of a unity in its composition (see sections 6 ix, and 7 iy), all domains of existence, as composite entities that include the unities that specify them, are deterministic systems in the sense indicated above. This has certain consequences for us living systems existing in language, and for the explanations that we generate as such beings. The following are some of these consequences:

i) Our domain of existence as the composite unities that we are as molecular autopoietic systems, is the domain of existence of our component molecules, and entails all the operational coherences proper to the molecular existence. Therefore, our existence as autopoietic systems implies the satisfaction of all the constraints that the distinction of molecules entails, and our operation as molecular systems implies the determinism entailed in the distinction of molecules.

ii) If we distinguish molecules as composite entities, they exist in the domain of existence of their components, and as such their existence implies the satisfaction of the determinism that the distinction of the latter entails. The same applies to the decomposition of the components of molecules, and so on recursively. Since unities and their domains of existence are brought forth and specified in their distinction in the happening of living of the observer, the only limit to the recursion in distinctions is the limit of the diversity of experiences of the observer in his or her happening of living (praxis).

iii) Since the observer as a living system is a composite entity, the observer makes distinctions in his or her interactions as a living system through the operation of the properties of his or her components. If the observer uses an instrument, then his or her distinctions take place through the operation of the properties of the instrument as if this were one of its components. The result of all this is that an observer cannot make distinctions outside its domain of existence as a composite entity.

iv) Descriptions are series of consensual distinctions subject to recursive consensual distinctions in a community of observers. Observers operate in language only through their recursive interactions in the domain of structural coupling in which they recursively coordinate consensual actions as operations in their domains of experiences through the praxis of their living. Therefore, all interactions in language between observers take place through the operation of the properties of their components as living systems in the domain of their reciprocal structural coupling. Or, in other words, we as human beings operate in language only through our interactions in our domain of existence as living systems, and we cannot make descriptions that entail interactions outside this domain. As a consequence, although language as a domain of recursive consensual distinctions is open to unending recursions, language is a closed operational domain in the sense that it is not possible to step outside language through language, and descriptions cannot be characterizations of independent entities.

v) Since everything said is said by an observer to another observer, and since objects (entities, things) arise in language, we cannot operate with objects (entities or things) as if they existed outside the distinctions of distinctions that constitute them. Furthermore, as entities in language, objects are brought forth as explanatory elements in the explanation of the operational coherences of the happening of living in which languaging takes place. Without observers nothing exists, and with observers everything that exists exists in explanations.

vi) As we put objectivity in parentheses because we recognize that we cannot experientially distinguish between what we socially call perception and illusion, we accept that existence is specified by an operation of distinction: <u>nothing pre-exists its distinction</u>. In this sense, houses, persons, atoms or elementary particles, are not different. Also in this sense, existence as an explanation of the praxis of living of the observer, is a cognitive phenomenon that reflects the ontology of observing in such praxis of living, and not a claim about objectivity. Therefore, with objectivity in parentheses, an entity has no continuity beyond or outside that specified by the coherences that constitute its domain of existence as this is brought forth in its distinction. The claim that the house to which I return every evening from work is the same that I left in the morning, or that whenever I see my mother I see the same person that gave birth to me, or that all the points of the path of an electron in a bubble chamber are traces left by the same electron, are claims that constitute cognitive statements that define sameness in the distinction of the unity (house, mother, or electron) as this is specified in the operation of distinction that brings it forth together with its domain of existence. Since according to all that I have said, cognitive statements are not, and cannot be, statements about the properties of independent objects, sameness is necessarily always a reflection by the observer in the process of observing in the domain of existence that he or she brings forth in his or her distinctions. Furthermore, since no entity can be distinguished outside its domain of existence as the domain of operational coherences in which it is possible, every distinction specifies a domain of existence as a domain of possible distinctions; that is, every distinction specifies a domain of existence as a versum in the multiversa, or, colloquially, every distinction specifies a domain of reality.

vii) A scientific explanation entails the proposition of a mechanism (or composite entity) that, if realized, would generate the phenomenon to be explained in the domain of experiences (praxis or happening of living) of the observer (see section 4). The generative character of the scientific explanation is constitutive to it. Indeed, this ontological condition in science carries with it the legitimacy of the foundational character of the phenomenal domain in which the generative explanatory mechanism takes place, as well as the legitimacy of treating every entity distinguished as a composite unity, asking for the origin of its properties in its organization and structure. And because this is also the case for our common sense explanations in our effective operation in our daily life, it seems natural to us to ask for a substratum independent of the observer as the ultimate medium in which everything takes place. Yet, although it is an epistemological necessity to expect such a substratum, we constitutively cannot assert its existence through distinguishing it as a composite entity and thereby characterize it in terms of components and relations between components. In order to do so, we would have to describe it, that is, we would have to bring it forth in language and give it form in the domain of recursive consensual coordinations of actions in which we exist as human beings. However, to do so would be tantamount to characterizing the substratum in terms of entities (things, properties) that arise through languaging, and which, as consensual distinctions of consensual coordinations of actions, are constitutively not the substratum. Through language we remain in language, and we lose the substratum as soon as we attempt to language it. We need the substratum for epistemological reasons, but in the substratum there are no objects, entities, or properties; in the substratum there is nothing (no-thing) because things belong to language. In other words, nothing exists in the substratum.

viii) Distinctions take place in the domain of experiences, in the happening or praxis of living of the observer as a human being. For this reason, the domain of operational coherences that an observer brings forth in the distinction of a unity as its domain of existence, also occurs in his or her domain of experiences as a human being as part of his or her praxis of living. Therefore, since language is operation in a domain of recursive consensual coordinations of consensual coordinations of actions in the domain of experiences of the observers as human beings, all dimensions of the domains of experiences of the observers exist in language as coordinations of actions between observers. As such, all descriptions constitute configurations of coordinations of actions in some dimension of the domains of experiences of the members of a community of observers in coontogenic structural drift. Physics, biology, mathematics, philosophy, cooking, politics, etc., are all different domains of languaging, and as such are all different domains of recursive consensual coordinations of consensual coordinations of actions in the praxis or happening of living of the members of a community of observers. In other words, it is only as different domains of languaging that physics, biology, philosophy, cooking, politics, or any cognitive domain exists. Yet, this does not mean that all cognitive domains are the same; it only means that different cognitive domains exist only as they are brought forth in language, and that languaging constitutes them. We talk as if things existed in the absence of the observer, as if the domain of operational coherences that we bring forth in a distinction would operate as it operates in our distinctions regardless of them. We now know that this is constitutively not the case. We talk, for example, as if time and matter were independent dimensions of a physical space. Yet, it is apparent from my explanation of the phenomenon of cognition that they are not and cannot be. Indeed, time and matter are explanations of some of the operational coherences of the domains of existence brought forth in the distinctions that constitute the ongoing languaging in the praxis of living of the members of a community of observers. Thus, time--with past, present, and future--arises as a feature of an explanatory mechanism that would generate what the observer experiences as successive nonsimultaneous phenomena; and matter arises as a feature of an explanatory mechanism that would generate what he or she experiences as mutually impenetrable simultaneous distinctions. Without observers nothing can be said, nothing can be explained, nothing can be claimed,... in fact, without observers nothing exists, because existence is specified in the operation of distinction of the observer. For epistemological reasons, we ask for a substratum that could provide an independent ultimate justification or validation of distinguishability, but, for ontological reasons, such a substratum remains beyond our reach as observers. All that we can say ontologically about the substratum that we need for epistemological reasons, is that it permits what it permits, and that it permits all the operational coherences that we bring forth in the happening of living as we exist in language.

ix) As we operate in language we operate in a domain of reciprocal structural coupling in our domain of existence as composite unities (molecular autopoietic systems), that is, we operate in the domain of existence of our components. Therefore, anything that we say, any explanation that we propose, can only entail distinctions that involve the operation of our components in their domain of existence as we operate as observers in language. Accordingly, it is in the domain where we exist as composite entities, that we distinguish molecules, atoms, or elementary particles as entities that we bring forth in language through operations of distinction that specify them as well as the operational coherences of their domains of existence. If what we call the physical domain of existence is the domain where physicists distinguish molecules, atoms, or elementary particles, then we as living systems specify the domain of physical existence as our limiting cognitive domain as we operate as observers in language, interacting in the domain of existence of our components as we bring forth the physical domain of existence as an explanation of the happening of our living. We do not exist in a pre-existing domain of physical existence; we bring it forth and specify it as we exist as observers. The experience of the physicist, be this in classic, relativistic, or quantum physics, does not reflect the nature of "the universe"; it reflects the ontology of the observer as a living system as he or she operates in language bringing forth the physical entities and the operational coherences of their domains of existence. Einstein made the assertion that scientific theories (explanations) are free creations of the human mind; and then, in what seemed to reveal a paradox, he asked the question, "How is it, if that is the case, that the universe is intelligible through them?" In this article I have shown that there is no paradox if one reveals the ontology of observing and the ontology of scientific explanations through putting objectivity in parentheses. Indeed, I have shown that a scientific explanation entails: a) the proposition of a phenomenon to be explained, brought forth as such *a priori* in the praxis of living (domain of experiences) of the observer; b) the proposition of an *ad hoc* generative mechanism, also brought forth *a priori* in the praxis of living of the observer, that if allowed to operate would generate the phenomenon being explained as a consequence to be witnessed by the observer in her or his praxis of living; c) the operational coherence of the four operational conditions that constitute its criterion of validation, as they are realized in the praxis of living of the observer; and d) the superfluity and impertinence of the assumption of objectivity. From all this it follows that the explanatory mechanism proposed in a scientific explanation is constitutively "a free creation of the human mind" because it is brought forth constitutively _a priori_ in the praxis of living of the observer, that is, without any other justification than the _ad hoc_ generative character of the phenomenon explained. It also follows from all this, that a scientific explanation constitutively explains the universe (versum) in which it takes place because both the explanatory mechanism and the phenomenon being explained occur, in a generative relation, as nonintersecting phenomena of the same operational domain of the praxis of living of the observer. Or, in other words, it also follows from all this, that since the operation of distinction specifies the entity distinguished as well as its domain of existence, a scientific explanation constitutively explains the universe (versum) in which it takes place because it brings with it the domain of operational coherences (the versum of the multiversa) of the praxis of living of the observer that it makes intelligible. Strictly, then, there is no paradox: scientific explanations do not explain an independent world or universe; they explain the praxis of living (the domain of experiences) of the observer, making use of the same operational coherences that constitute the praxis of living of the observer in languaging. It is here that science is poetry.

11. REALITY

The word "reality" comes from the Latin noun _res_ that means "object" (thing), and as it is commonly used signifies objectivity without parentheses. The real, and sometimes the really real, is meant to be that which exists independently of the observer. Now we know that the concepts entailed in this way of speaking cannot be sustained. Objects, things, arise in language when a consensual coordination of actions, by being consensually distinguished in a recursion of consensual coordinations of actions, obscures the actions that it coordinates in the praxis of living in a consensual domain. Since according to this circumstance, an object, a unity, is brought forth in language in an operation of distinction that is a configuration of consensual coordinations of actions, when an object is distinguished in language its domain of existence as a coherent domain of consensual coordinations of actions becomes a domain of objects, a domain of reality, a versum of the multiversa such that all that is in it is all that is entailed in the consensual coordinations of actions that constitute it. Every domain of existence is a domain of reality, and all domains of reality are equally valid domains of existence brought forth by an observer as domains of coherent consensual actions that specify all that is in them. Once a domain of reality is brought forth, the observer can treat the objects or entities that constitute it both as if they were all that there is and as if they existed independently of the operations of distinction that bring them forth. And this is so because a domain of reality is brought forth in the praxis of living of the observer as a domain of operational coherences that requires no internal justification.

It follows from all this, that an observer operating in a domain of reality necessarily operates in a domain of effective actions, and that another observer claims that the first one commits a mistake or has an illusion only when the first observer begins to operate in a domain of reality different from the one that the second observer expected. Thus, if we specify the operation of distinction "ghost," then ghosts exist, are real in the domain of existence brought forth in their distinction, and we can do effective actions with them in that domain, but they are not real in any other domain. Indeed, everything is an illusion outside its domain of existence. In other words, every domain of reality as a domain of operational coherences brought forth in the happening of living of the observer in language, is a closed domain of effective consensual actions, that is, a cognitive domain; and conversely, every cognitive domain as a domain of operational coherences is a domain of reality. What is uncanny, perhaps, is that although different domains of reality are seen by an observer as different domains of coordinations of actions in an environment, they are lived by the observer as different domains of languaging which differ only through their ongoing transformation in the different circumstances of recursion in which they arise. We as observers can explain this now by saying that, as we operate in language through our consensual interactions in the happening of living of a community of observers, our structural drift in the happening of our living becomes contingent upon the course of those consensual interactions, and that this takes place in a manner that keeps the transformation of the happening of our living congruent with the domain of reality that we bring forth in that community of observers, or we disintegrate as members of it. It is this that makes us observing systems systems capable, through language, of an endless recursive generation of new cognitive domains (new domains of reality) as new domains of praxes of observing in our continuous structural drifts as living systems.

12. SELF-CONSCIOUSNESS AND REALITY

The self arises in language in the linguistic recursion that brings forth the observer as an entity in the explanation of his or her operation in a domain of consensual distinctions. Self-consciousness arises in language in the linguistic recursion that brings forth the distinction of the self as an entity in the explanation of the operation of the observer in the distinction of the self from other entities in a consensual domain of distinctions. As a result, reality arises with self-consciousness in language as an explanation of the distinction between self and non-self in the praxis of living of the observer. Self, self-consciousness, and reality exist in language as explanations of the happening of living of the observer. Indeed, the observer as a human being in language is primary with respect to self and self-consciousness, and these arise as he or she operates in language explaining his or her experience, his or her praxis of living as such. That the entities brought forth in our explanations should have an unavoidable presence in our domain of existence, is because we are realized as observers as we distinguish these entities in the domain of operational coherences that they define as we distinguish them. We do not go through a wall in the praxis of living because we exist as living systems in the same domain of operational coherences in which a wall exists as a molecular entity, and a wall is distinguished as a composite entity in the molecular space as that entity through which we cannot go as molecular entities.

The observer is primary, not the object. Better, observing is a given in the praxis of living in language, and we are already in it when we begin to reflect upon it. Matter, energy, ideas, notions, mind, spirit, God,... are explanatory propositions of (about) the praxis of living of the observer. Furthermore, matter, energy, ideas, notions, mind, spirit, or God, as explanatory propositions entail different manners of living of the observer in recursive conservation of adaptation in the domains of operational coherences brought forth in their different distinctions. Thus, when the observer operates with objectivity without parentheses, he or she operates in an explanatory avenue that entails neglecting the experiential indisinguishability between what we call perception and illusion, and when he or she operates with objectivity in parentheses he or she operates in an explanatory avenue that entails accepting this indistinguishability as a starting point. In the explanatory path of objectivity without parentheses, the observer, language, and perception cannot be explained scientifically because in this explanatory path it is assumed that the observer can make reference to entities that exist independently of what he or she does, an assumption which is in contradiction with the structural determinism of living systems; while in the explanatory path of objectivity in parentheses there is no such contradiction. At the same time, as one operates within any given domain of reality one can operate with objectivity without parentheses without contradiction, but when a disagreement arises with another observer, and one thinks that it is not a matter of a simple logical mistake, one is forced to claim a privileged access to an objective reality to resolve it, and to deal with errors as if they were <u>mistakings</u> of what is. If in similar circumstances one is operating with objectivity in parentheses, one finds that the disagreeing parties operate in different domains of reality, and that the disagreement disappears only when they begin to operate in the same one. Furthermore, one also finds that errors are changes of domain of reality in the operation of an observer that he or she notices only <u>a posteriori</u>. Finally, by operating in the explanatory pathway of objectivity without parentheses we cannot explain how an observer operates in the generation of a scientific explanation because we take for granted the abilities of the observer. Contrary to this state of affairs, if we operate in the explanatory pathway of objectivity in parentheses, scientific explanations and the observer appear as components in a single closed generative explanatory mechanism, in which the properties or abilities of the observer are shown to arise in a different phenomenal domain than the one in which its components operate.

We human beings exist only as we exist as self-conscious entities in language. It is only as we exist as self-conscious entities that the domain of physical existence exists as our limiting cognitive domain in the ultimate explanation of the human observer's happening of living. The physical domain of existence is secondary to the happening of living of the human observer, even though in the explanation of observing the human observer arises from the physical domain of existence. Indeed, the understanding of the ontological primacy of observing is basic for the understanding of the phenomenon of cognition. Human existence is a cognitive existence and takes place through languaging; yet, cognition has no content and does not exist outside the effective actions that constitute it. This is why nothing exists outside the distinctions of the observer. That the physical domain of existence should be our limiting cognitive domain does not alter this. Nature, the world, society, science, religion, the physical space, atoms, molecules, trees,... indeed all things, are cognitive entities, explanations of the praxis or happening of living of the observer, and as such, as this very explanation, they only exist as a bubble of human actions floating on nothing. Every thing is cognitive, and the bubble of human cognition changes in the continuous happening of the human recursive involvement in coontogenic and cophylogenic drifts within the domains of existence that he or she brings forth in the praxis of living. Every thing is human responsibility.

The atom and the hydrogen bombs are cognitive entities. The big bang, or whatever we claim from our present praxis of living gave origin to the physical versum, is a cognitive entity, an explanation of the praxis of living of the observer bound to the ontology of observing. That is their reality. Our happening of living takes place regardless of our explanations, but its course becomes contingent upon our explanations as they become part of the domain of existence in which we conserve organization and adaptation through our structural drifts. Our living takes place in structural coupling with the world that we bring forth, and the world that we bring forth is our doing as observers in language as we operate in structural coupling in it in the praxis of living. We cannot do anything outside our domains of structural coupling; we cannot do anything outside our domains of cognition; we cannot do anything outside our domains of languaging. This is why nothing that we do as human beings is trivial. Everything that we do becomes part of the world that we live as we bring it forth as social entities in language. Human responsibility in the multiversa is total.

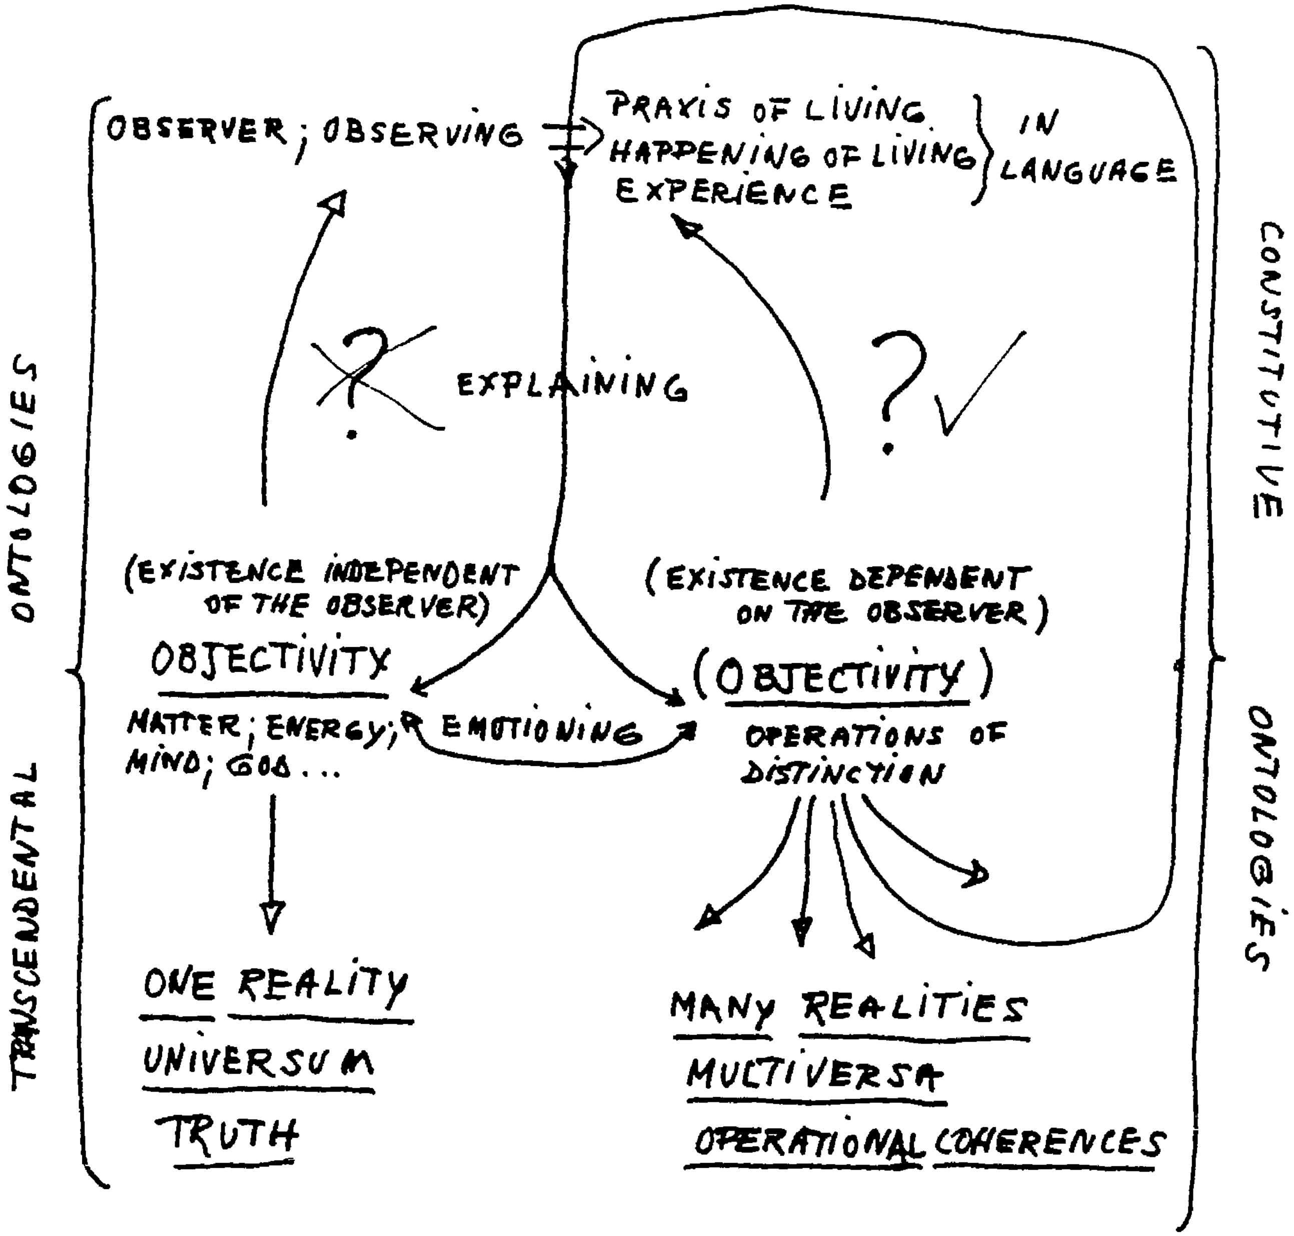

ONTOLOGY OF OBSERVING

BIBLIOGRAPHY

Maturana H. R. (1970) "Biology of Cognition", in Autopoiesis and Cognition by Maturana H. R. and F. G. Varela. Reidel 1980.

Maturana H. R. and F. G. Varela (1973) "De maquinas y seres vivos" in english "Autopoiesis: the organization of the living", in Autopoiesis and Cognition by Maturana H. R. and F. G. Varela, Reidel 1980.

Maturana H. R. (1975) The organization of the living: a theory of the living organization. Internat. J. of Man-Machiene Studies 7; 313-332

Maturana H. R.(1978) "Biology of language: epistemology of reality", in Psychology and biology of language and thought. Editors: G. A. Miller and Elizabeth Lenneberg. Academic Press.

Maturana H. R.(1980) "Autopoiesis: reproduction, heredity and evolution", in Autopoiesis, dissipative structures and spontaneous social orders. Edited by Milan Zeleny. AAAS Selected Symposium 55. Westview.

Maturana H. R. and F. J. Varela (1987) The tree of knowledge. New Science Library. Shambhala.

Margulis Lynn (1981) Symbiosis in cell evolution. Freeman.

NOTES

NOTES

NOTES

NOTES

NOTES