In: M. N. Ozer (Ed.), *A cybernetic approach to the assessment of children: Toward a more humane use of human beings*. Boulder, CO: Westview Press, 67–113, 1979. 056

Cybernetics, Experience, and the Concept of Self

The mind organizes the world by organizing itself.
Jean Piaget, 1973

Encountering the term "cybernetics," most people tend to think of robots, computers, and electronics. This is not surprising, considering the noise that has been made about such gadgets and the hopes and fears associated with them. But there is another aspect to cybernetics that, in the long run, may turn out to be more important. When Norbert Wiener launched the term some 30 years ago, he defined it as "the study of control and communication in the animal and the machine" (Wiener, 1948). Since then, this study has led to a way of thinking about perception, behavior, and cognition that is revolutionary, not so much because of the problems it attacks, but rather because of the way in which it views them.

I shall not say very much here about cybernetics as a discipline, but I shall adopt a cybernetic attitude and develop some ideas on how a child forms certain basic concepts, among them the concept of self. Drawing on Piaget's analysis of cognitive development during the sensorimotor period, I shall try to show, on the one hand, that his theory is quite compatible with the cybernetic way of thinking and, on the other hand, that the cybernetic way of thinking may help to illuminate some of the darker corners of the theory.

There are three areas of cybernetic thought that are particularly germane to the study of early cognitive development: self-regulation, inductive learning, and the constructivist approach to experience and its organization. The following section explicates the notions and conceptions on which much of the subsequent discussion is based.

The second section, "Piaget from a Cybernetic Viewpoint," focuses on some of the most elementary conceptual operations that, from a logical-theoretical point of view, seem to be indispensable to such concepts as *sameness, identity, continuity, space, change, motion, cause, and time*. The section titled "Some Basic Constructs" then outlines a tentative approach to the construct of the *experiential self* and attempts to demonstrate why the locus of active experience, the entity that embodies the experiencer, must necessarily remain outside our picture.

Much of what is said in the second and third sections can make sense only if the reader keeps in mind the points raised in the last two paragraphs of the following section. The traditional view, both in psychology and epistemology, disregards the inevitable dichotomy between what can be said about observed organisms and what organisms might be able to say about their own experience. Insofar as the cyberneticist is a builder of models (physical or conceptual) that are supposed to regulate or govern themselves, he must remain aware of that dichotomy.

Feedback and Self-Regulation

Self-regulation, in cybernetics, usually refers to the principle of "negative feedback." Some practical applications of this principle were in use more than 2,000 years before its theoretical significance as an explanatory device in biology and psychology was discovered. In the simplest terms, control by means of "negative feedback" is an arrangement that enables a system (e.g., an animal or a machine) to gauge an activity according to its effect. Philon of Byzantium, in the third century B.C., built one of the earliest fully documented examples: an oil lamp in which the level of oil in the burner controlled the amount of oil fed into the burner from a reservoir (Mayr, 1970).

One of the most common examples today, a good deal more complicated in structure but embodying the same principle, is the thermostat. Here, there is a thermometer that senses the temperature in the area to be controlled. If that temperature rises beyond a preset value, a contact breaker closes and a cooling mechanism is switched on. If the temperature sinks below the set value, a heater is switched on.

To make the feedback principle quite explicit, we have to isolate its essential ingredients. There is, first of all, a sense organ that can indicate the actual temperature. Then there is the desired temperature or *reference value* that has been set by someone, and a comparator, where the sensed and the desired temperatures can be compared. As long as reference value and actual temperature are the same, the thermostat will do nothing. But if there is a discrepancy in one direction or the other, indicating either that the actual temperature in the controlled area has been *disturbed* or that the reference value has been changed, the thermostat sends an *error signal* to activate the cooling or heating machinery, as appropriate. If everything works as expected, the temperature in the controlled area will change in the right direction, the value indicated by the sense organ will adjust to the reference value, and the error signal will cease.

The theoretical importance of these gadgets springs from the fact that they provide an irrefutable demonstration of purposive, goal-directed behavior. With that, the concept of purpose is removed from the context of Aristotelian teleology that placed it out of bounds for the modern scientist. It now has a place in the design of functioning machines and can be reinstated as a legitimate, precise, and extremely useful explanatory concept.

As a result of its technological implementations we can now also discriminate two types of purpose that, formerly, seemed inextricably confused. We can clearly see that the thermostat has the purpose *of* maintaining the temperature in the controlled area close to the reference value, whereas it is some outside agent that sets the reference value as a purpose *for* the thermostat (Pask, 1969, pp. 22–24). This distinction is of particular significance if we want to use the feedback principle to explain living organisms. While the simple arrangement illustrated by the thermostat serves well enough as a model for the homeostatic functions that control single physiological conditions in the body, such as internal temperature, sugar level, and blood pressure (see Cannon, 1932), it is obviously insufficient to explain directed behaviors whose goals change from situation to situation and from context to context.

A more sophisticated system will have a hierarchical arrangement in which the reference values on one level are adjusted by a control system on another level. (See, for instance, Powers, 1973). In arrangements of this sort, it will be the goal *of* one level to set the goal *for* another.

More important is the fact that the feedback model as I have so far described it does not provide for any form of learning. There are different ways of learning that can be incorporated in cybernetic models (Powers, 1973; McFarland, 1971), and one of them is, in principle, an implementation of the age-old process of inductive inference. This was first suggested by Kenneth Craik in the early 1940s and then practically applied by Ross Ashby (1970).

Learning as a Process of Induction

The self-filling oil lamp, the thermostatic mechanism, and all similar devices that have the built-in purpose of maintaining some condition close to a pre-set reference value, are obviously the result of deliberate design. There was a designer who not only knew why he wanted a certain condition or quantity kept constant, but also how this could be achieved. The thermostat that controls the temperature in a building does not have to *learn* that it is the air conditioner that must be switched on when the sensory signal indicates a temperature above the reference value, and that it is the heating system that must be switched on in the opposite case. Since these connections are built in by the designer, the error signals that emanate from the comparator, indicating either "too hot" or "too cold," run along fixed lines to the appropriate machinery. The appropriateness of the two types of machinery, for heating and for cooling, is something that was decided by the designer on the basis of his prior experience and learning.

If the feedback model is to be of use in the study of the more complex forms of behavior we see in animals, and in humans, we shall have to give it some capability for learning. In Craik's words (1966, p. 59),

> We should now have to conceive a machine capable of modification of its own mechanism so as to establish that mechanism which was successful in solving the problem at hand, and the suppression of alternative mechanisms. Although this may seem a great demand, we can be comforted by the reflexion that animals and man can only modify their activity within the limits imposed by their anatomy, or the materials and machines available; though it is a great demand, it is not an infinitely great one.

In a sense, the solution lies in this very early statement of the problem. It consists in establishing and recording for every kind of error signal (problem) "that mechanism which was successful in solving the problem." In other words, if there are several kinds of *disturbance* and, consequently, several kinds of error signals, the system has to discover which of the activities in its behavioral repertoire is most likely to correct a particular error signal. On the simplest level this can be achieved only through inductive inference.

> A living system, due to its circular organization, is an inductive system and functions always in a predictive manner; what occurred once will occur again. Its organization (both genetic and otherwise) is conservative and repeats only that which works (Maturana, 1970, pp. 15–16).

The simplest learning system, thus, will have a repertoire of several different activities and at least one sense organ and one comparator that generates an error signal whenever the sensory signals do not match the reference value. What it has to *learn* (i.e., what is not determined by fixed wiring), is to make the error signal trigger the particular activity that is likely to reduce it.

There are several other assumptions that have to be made if the system is to work. First, there must be at least one activity in the systems' repertoire that can actually influence the condition represented by the sensory signals (e.g., the heating and cooling mechanisms in the thermostat). Second, the system must start out with something like a "tendency to act" whenever there is an error signal. Third, to learn, the system must be able somehow to keep track of whether or not a particular activity reduces a particular error signal; in other words, it must have some form of memory.

The first of these three assumptions is obvious and requires no explanation. The second may be questioned; but if we adopt the theory of evolution it should not be difficult to concede that organisms that, in the face of disturbance, will do *something,* have a better chance of survival than organisms that do not act at all. The third assumption is the most problematic, if only because we still have no adequate idea of what memory is and how it functions. On the other hand, we are not only certain that we humans remember past events, but we also know that there are experiences that leave some kind of a record in other animals.

Before discussing the general implications of the learning feedback model, there is one practical point to stress because, although it is implicit in any description of the system's functioning, its full import is rarely appreciated. The learning process necessarily begins with the random choice of an activity in response to an error signal. If that activity does not reduce the error signal, another activity will be tried, and so forth, until one is found that *does* lead to a reduction of the "disturbance." This trial-and-error procedure stops when the trial brings success. The connection between that activity and the particular error signal is then recorded and from then on, if there is no disruption, that error signal will "automatically" call up the activity that was successful. However, if there had been no disturbance and, consequently, no error signal, no learning could have taken place.

It seems quite possible, if not likely, that an organism with a fairly large repertoire of activities might have several that could reduce the same disturbance. Given the original random approach, however, the organism may not discover this. Since it has recorded that activity x was successful in eliminating a particular disturbance, it will enact this activity in response to that error signal as long as it continues to be successful, and there will be no motive to try others. One can say that such an organism will learn *only as a result of disturbance,* and it will give up or modify something it has learned only when this again leads to disturbance. This mode of functioning, as we shall see later, fits very well into the Piagetian conception of the complementary processes of *assimilation* and *accommodation.*

The Subject's Construction of Knowledge

The preceding paragraphs cover only a fraction of the work that has already been accomplished with the cybernetic approach to the analysis of regulatory functions in organisms. An excellent technical survey has been provided by McFarland (1971) and an integrated theory of behavior by Powers (1973). The epistemological aspects of the cybernetic approach have particularly interesting implications for the study of cognition and cognitive development.

The fact that we can build feedback systems capable of self-regulation within certain parameters makes it possible for us to visualize and exemplify the most elementary features of an organism's relation to its environment. For the first time we can not only ask but answer in a thoroughly operational way questions such as "What are the data of an organism's experience?"; "What can and does an organism associate when it is learning?"; and "What constitutes knowledge in an organism, and how does it relate to the outside world?"

A learning homeostat (Ashby, 1970) may be an abominably primitive organism compared to even an amoeba. It has only one function, against several dozen in the amoeba, and it certainly cannot reproduce itself or manifest any of the other characteristics of life. Nevertheless, insofar as it is self-regulatory, it is analogous to the self-regulatory functions of the amoeba. Even if this function is carried out by completely different elements in the two systems, the logical steps involved are equivalent.

If the organism's learning is inductive, it operates on the assumption (or belief) that there must be some regularity in its *experience:* "what occurred once will occur again." In fact, there can be no learning without that assumption, for, as Hume put it, "If there be any Suspicion, that the Course of Nature may change, and that the past may be no Rule for the future, all Experience becomes useless, and can give rise to no Inferences or Conclusions" (Hume, 1748/1963, p. 47).

Hume attributes the regularity to "the Course of Nature" and that is saying a good deal too much. The organism can afford to be more modest and assume merely some regularity in its experience, in the "data" or "signals" with which it operates and which, necessarily, are the only ones to which it has access. In the language of psychologists, they are the *proximal* data.

In our model the proximal data comprise the signals from the sense organ, the reference value, the error signal, and some kind of labels or signals that represent each activity in the organism's repertoire. About the last of these there are divergent views among cyberneticists, and there is certainly a need for differentiation according to the context of the activities; in one set of circumstances, for instance, one might want to speak of different *commands* that govern the activities; in another one might want to focus on *proprioceptive* signals from the parts of the organism that are involved in carrying out the activities.

As Powers (1973) has formulated it, an organism "behaves in order to control its perception." In more explicit terms, that means that an organism acts to modify a sensory signal towards a match with the reference signal, so that there will no longer be the error signal that triggers the activity. On the simplest level we may even say that an organism acts to eliminate error signals. And its learning consists in finding (and recording for future use) an activity that will do that. The trials with different activities will cease when the error signal ceases, and the successful connection that has “caused” the reduction of the error signal will be the new “knowledge.” The next time that same error signal comes from the comparator, the organism will “know” which activity to choose.

The point is that the organism has neither need nor use for what an *observer* of the organism calls its environment. Provided there is some recursion in the sequential conjunction of certain activities and certain modifications of sensory signals, the organism can learn to eliminate error signals. It needs no knowledge of distal data, of *environment,* or of an outside reality, and there seems to be no reasonable way for the organism to acquire such knowledge.

For an observer, of course, it may be plausible to establish all sorts of relations between the organism’s “output” (i.e. the effect of its activities on the environment) and its “input” (i.e. the environmental “stimuli” assumed to cause the organism’s sensory signals). But these items may not be quite as straightforward as they appear. I have elsewhere argued for a radical constructivist view of knowledge (von Glasersfeld, 1975, 1976, 1977; Richards and von Glasersfeld, 1978) on all levels of organization. Here I shall confine myself to pointing out that the kind of knowledge our simple organism acquires by installing connections between error signals and activities is, indeed, a form of *construction,* and since it deals exclusively with the proximal data of the organism’s own subjective experience, one would be justified in calling it wholly *subjective.* From there to Piaget’s statement that “intelligence organizes the world by organizing itself” (Piaget, 1954) may not be nearly as far as it seems.

Regularities, Rules, and Explanation

If the assertion that intelligent organisms selectively organize their experiential world were made by Piaget alone, one could perhaps brush it aside. Psychologists have more than once launched ideas that later turned out to be as absurd as they sounded**.** But while common sense and certain branches of science are still enmeshed in the realist faith of the 19th century, physics, the science we consider the “hardest,” the least speculative, and the most dependable when it comes to empirical tests, has moved away from the belief that the knowledge we gather from experience can or even should depict an objective reality. To express the idea behind tile quotation at the head of this *chapter,* I might just as well have chosen Einstein’s formulation: “It is the theory which decides what we can observe” (quoted in Heisenberg, 1971, p. 63), or a somewhat more direct and factual one from Heisenberg “The mathematical formulations (of physics) no longer depict Nature, but rather our knowledge of Nature” (Heisenberg, 1955, p.19).

If science can no longer be said to observe, explore, and eventually explain a “real” world, supposed to exist and to be the way it is, regardless of whether we are experiencing it or not, what then is science doing? “What science deals with is an *imagined* word” and it is “a construct, and some of the peculiarities of scientific thought become more intelligible when this fact is recognized” (Hebb, 1975, pp. 4 and 9). Scientists look for repetitive conjunctions among experiential (or experimental) data in the hope of establishing relatively reliable correlations or, better still, causal connections. They look for regularities in their experience that would allow the formulation of rules that could then be used, in the same old inductive fashion, to explain past experiences and to predict future ones.

In short, scientists seem to be involved in a process of learning that, *qua* process, is not at all unlike the learning of our ultrasimple model organism. Instead of establishing experiential regularities from which to derive rules of action to eliminate disturbances, they are searching for experiential regularities from which to derive rules of conceptualization for a homogeneous, internally consistent ordering of experience. In doing this, they encounter no shortage of disturbances that, as in the simple feedback model, must be eliminated. But the disturbances are now created by incompatibilities of rules and conceptualizations. And a closer look at history of science should convince anyone that scientists, in their quest for consistency and compatibility, are prepared not only to modify the conceptual relations by means of which they order experiential items, but also to restructure quite radically those items that they consider basic elements (see Hanson, 1958; Kuhn, 1970; Feyerabend, 1975).

The Use of Black Boxes

One of the early contributions of cybernetics to the theory of scientific analysis and investigation was the concept of a black box. This is obviously and exclusively an observer's concept. It is used for items that one suspects of performing some function, but that, for one reason or another, one cannot dismantle to see what is going on inside. Thus one might say, black boxes do not exist, but there may be many boxes that are black for someone. Living organisms are a case in point, especially the functions that, one suspects, constitute their intelligence. When one cuts open the organism, most of its interesting functions remain invisible or have ceased.

Cyberneticists in general, unlike the strict behaviorists who profess no interest in the internal machinery of a black box, are ready to make conjectures and to test their plausibility, much as scientists in other fields formulate and test hypotheses about things not directly observable. But there is one important difference. Other scientists believe that sooner or later they will be able to match their hypotheses with observations and find out whether they were right or wrong. A true cyberneticist knows that the "intelligent" functions he is investigating are never observable. All he can possibly check on is the material they start with and the results they produce: their *input* and *output*. Hence he tries to design a model that, given the same input, will always produce the same output as the black box. Although it is rewarding to design a physical model that actually works, cyberneticists much more often (for practical reasons or for lack of time or money) have to be satisfied with theoretical models that demonstrate at least that there is a logically feasible way of performing the function.

There are two aspects of the study of cognitive development that warrant the use of the black box concept. The first might seem almost trivial, if it were not for the widespread misunderstandings that certain models of cognitive development have generated. For the observing psychologist, the developing child is in many ways a very black box. But the observer, when he hypothesizes the child's intelligent or cognitive processes, must try to adopt a perspective *from inside the black box*. To say that he must try to see things from the child's point of view is to say far too little. It is *not* a question of trying to see the same things the observer sees but from a different angle. Instead, it is a question of hypothesizing how such a black box, whose cognitive processor has access to nothing but proximal data, or internal events, can possibly articulate and structure its experiential field to end up with a viable representation of an “external” world.

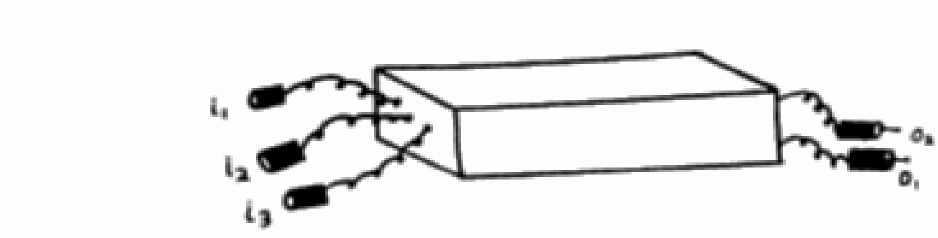

Figure 1. An item is considered a “black box” if, for one reason or another, the observer cannot investigate what goes on inside it. All he can see are inputs (i1, i2, i3, ...) and outputs (O1, O2, ...). On the basis of continued observation, both the behaviorist and tie cyberneticist will attempt to establish regularities in the input-output sequences. If this can be achieved, the behaviorist is satisfied because the observed regularities allow him to make certain probabilistic predictions about the item’s behavior. The cyberneticist goes on to ask what kind of mechanism inside the box could account for the observed regularities. Hence he will try to construct a “model” that produces the same input-output sequences as the black box.

The naive realist view, that what we experience has to be a more or less direct reflection of an independently existing reality in which everything is fully structured and fixed, has made insight into cognitive development impossible. On that basis, development seems an obligatory one-way street of maturation and learning—in the sense of “finding out” or “discovering” how things really are and how they work. The only theoretical puzzle would be that development so rarely leads to any adequate understand or wisdom.

This leads to the second use I want to suggest for the black box concept. If it is the experiencer’s intelligence or cognitive activity that, by organizing itself, organizes experience into a viable representation of a world, then one can consider that representation a model, and the “outside reality” it claims to represent, a black box. The moment we attribute to the learning homeostat (to use our original example) the capabilities of representation and hypothesis, it can begin to conjecture how it comes about that a certain activity regularly results in the modification of a certain sensory signal. It can begin to construct a representation of an external world with which it has two conceivable points of contact: “input” in the form of its effect on the outside, and “output” in the form of outside events that cause its own sensory signals. The representation, therefore, will have to be no more and no less than a hypothetical model of functions, entities, and events that could “explain” regularities in the organism’s experience. And as a cyberneticist would expect, there is no way to match the model against the “real” structure of the black box.

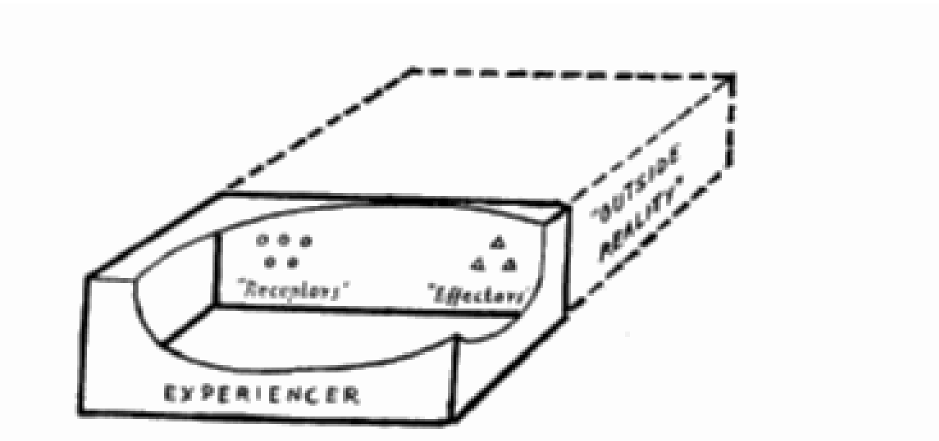

Figure 2. From the experiencer's point of view, "experience" basically consists of signals, which we may divide into "receptor" and "effector" or "sensory" and "motor" signals. His or her representation of an "outside reality" will necessarily be based on such regularities as they can establish in the experienced signal sequences. There is no way of investigating the "outside" except by observing what sensory signals follow upon certain effector signals which they categorized as "activities." Hence any representation of the outside reality will be a model of an inaccessible black box in which the input, registerd as effector signals, is systematically related to the output, registered as receptor signals.

Observer and Observed

To observe means to focus attention with particular care. When we use the verb transitively (which we usually do), it implies that a particular area of our experiential field is discriminable from the rest. Logically speaking, that part of our experience is the object we are observing. Psychologists, however, persistently speak of observing *subjects*. There are two reasons for that practice. First, psychologists presumably want to indicate at once that what they are observing is not a dead object but an active item that can play the role of subject in a number of activities. For the rigid behaviorist, of course, this is merely a conventional and somewhat misleading manner of speech, because in his view the rats he runs should be as passively determined by outside events as billiard balls or any other mindless objects. The cognitive psychologist, on the other hand, is ready to attribute subject-properties and even subjectivity to the items on which he focuses his attention In other words, he deliberately looks for cognitive organization and cognitive functions.

The second reason psychologists say they observe subjects is that they do not want to acknowledge their own activity as observers. By calling the observed item "subject," they hide the fact that it is their own observer's activity that determines the items they observe.

Usually, to *observe* comprises a little more than just focusing attention on discriminable items. It also involves an attempt to isolate some form of interaction between the part of the observer's experience that observers consider to be an organism and the rest of their experiential field, which then is considered that organism's environment.

In principle, this division between organism and environment is quite similar to the figure-ground division an artist makes when he draws the outline of an object on a sheet of paper. The point is this: just as the artist's figure and ground are both parts of the sheet of paper, so the organism and its environment are parts of the observer's field of experience. There can be no doubt that the division between an observed organism and its environment is both legitimate and extremely useful, provided we remain aware of who makes the division and where it is made (von Foerster, 1970). This awareness alone can help us avoid two traps that have generated enormous confusion in the past.

First, there is the tempting but logically erroneous idea that what we rightly call "environment" relative to an organism when both the organism *and* its environment are being observed by us, must also be *our* environment and can, therefore, be held causally responsible for what we ourselves experience. Second, there is the mistaken belief that the "environment" which is part of our experiential field has to be identical with the experiential field of the observed organism (von Glasersfeld, 1976).

The conception of an experiencer who, facing a black box reality, constructs a model of a world out of such regularities as he can establish in his experience provides us with a new perspective on some of the age-old problems of traditional epistemology. Similarly, the conception of an observer observing organisms that for him are black boxes with their own inaccessible fields of experience, provides a perspective on the study of cognitive development that only a few psychologists and educators have begun to appreciate. Among those few, of course, looms Piaget. The next section deals with his analysis of the most fundamental cognitive constructs.

Piaget from a Cybernetic Viewpoint

The Construction of Permanent Objects

One of the revolutionary findings of Piaget's research in cognitive development was that it takes the child almost all of his first two years to acquire the notion that objects have an "existence" of their own and can be presumed, under ordinary circumstances, to remain what they are, even when one is not actually perceiving them. The discovery was revolutionary because in our adult, common-sense world there is probably not a single everyday thought or activity that does not in some crucial way rely upon that notion of object permanence. It is perhaps the most deeply rooted notion after that of our own existence. Most of us, failing after a long, meticulous search to find a misplaced object, would doubt our wife's, husband's, or dearest friend's honesty, rather than give up the belief that the object must be *somewhere.* Hence it is quite a shock to be told that as children we did not start out with that notion but gradually and laboriously acquired it.

Piaget's analysis of that acquisition, the six stages into which he has articulated it, and the observational and experimental documentation he and others have accumulated during 50 years of research are well known. So, taking for granted that children can and do develop a concept of object permanence, I shall briefly examine some of the steps that seem indispensable for a hypothetical model of an organism to construct a similar concept.

There is hardly an introductory text of psychology today that does not refer to the child's development of the concept of object permanence. Looking at some of them, however, one gets the impression that the authors never read beyond the first hundred pages of Piaget's *The Construction of Reality in the Child* (1937/1971). In that first section of the book he does, of course, expound his theory of the genesis of the object concept, while the subsequent sections deal with the concepts of space, causality, time and, finally, the universe. Though he treats the construction of these concepts sequentially, he makes it very clear that he does not consider them sequential in the child's development. In his view, one conceptual construction gives rise to all these concepts. In other words, experiential objects, space, causality, time—and I would add the concepts of change, motion, substance, identity, and self—all stem from *one* common initial construction and are therefore *connate* and inextricably interrelated. Hence, mention of "steps" in subsequent paragraphs does not imply a chronological but a logical sequence. There are certain steps that are logically indispensable prerequisites for others. But the logic is our logic, an *observer's* logic, and as such it applies to a model the observer is building.

Establishing Sameness

Before anything like the notion of permanence can arise, there must be the possibility of separating items in experience and, then, of considering them the same. Some such ability, clearly, would be prerequisite in any organism that operates inductively—that establishes recurrences in its experience, considers them regularities, and draws inferences from them. Indeed, there could be no learning at all if there were not some notion of sameness. For the behaviorist, of course, it is simply a basic fact that organisms are capable of "stimulus generalization." Taking the ability for granted helps in avoiding theoretical problems, but not in constructing a viable model of organisms and their behavior. In our inquiry, to which the concepts of recurrence, regularity, and permanence are indispensable, we cannot avoid asking how an organism could act as though two experiences were *the same*.

Even if we wanted to believe that perception is nothing but an organism's internal replication of a ready-made external world in which objects are *given,* we would have to explain how such an internal replication is constituted. This question runs into the same difficulty regardless of whether it is asked by a realist or by a constructivist. Neither can disregard the simple fact that an "object," from the point of view of the experiencing organism (i.e., in terms of the organism's sensory experience) is never quite the same on different occasions.

For example, the visual experience that we consider an instance of a specific object is different every time. The object's shape changes according to the angle, and its size according to the distance from which it is seen. Its color changes according to the illumination, and other parameters are no less variable according to changes in the context. What, then, constitutes the invariant object which the organism recognizes? There seems to be no way around the assumption that, as far as the organism is concerned, an "object" must be a construct, actively abstracted from a number of experiences by holding on to a somewhat flexible constellation of characteristics and allowing each of them to vary within a certain range.

A closely related problem is that of the child's acquisition of names. Most psycholinguists agree that the child has to form some kind of concept before he can learn to associate a name with it, but there are different ideas on how that concept develops. Roger Brown (1958) summarized the two extremes and then provided a convincing synthesis.

Suppose a very young child applies the word *dog* to every four-legged creature he sees. He may have abstracted a limited set of attributes and created a large category, but his abstraction will now show up in his vocabulary. Parents will not provide him with a conventional name for his category, e.g., *quadruped*, but instead will require him to narrow his use of *dog* to its proper range…

The child who spontaneously hits on the category four-legged animals will be required to give it up in favor of dogs, cats, horses, cows, and the like … The schoolboy who learns the word *quadruped* has abstracted from differentiated and named subordinates. The child he was abstracted through a failure to differentiate. Abstraction after differentiation may be the mature process, and abstraction from a failure to differentiate the primitive.

Categorizing experiential objects in the particular way prescribed by the language the child has to acquire is, of course, not quite the same as deriving an object-concept from two or more experiences (see "Equivalence and Continuity," below). But in both cases differentiation and abstraction seem to play much the same role. In the naming of objects it is the conventional nature of linguistic expressions that compels the child to structure his concepts in a certain way. In the construction of recursively usable object-concepts it is the organism's active search for recurrence and regularity that compels it to create "sameness" by focusing on similarities and disregarding differences. In Piagetian terms, this active imposition of invariance on instances of experience that are always different in some way is the ubiquitous process of *assimilation*.

Assimilation and Accommodation

In the last paragraph I tried to show that an object-concept can only be the result of the experiencing organism's active construction. Even if an observer with realist convictions believes that the object is a ready-made "thing-in-itself" out there in reality waiting to be experienced, for an organism, every single experience of that object would nevertheless be a little different. So there is no immediate way of justifying the assumption that the concept of such an object could be a directly derived replica or representation of the "real" thing.

Giving up the realist epistemology, however, does not solve all problems. There remains, above all, the practical question how an organism could recognize an experience as the repetition of a previous one, when the two experiences are, in fact, not congruent in the sensory or experiential elements that compose them.

Piaget has resolved this difficulty by the introduction of the concept of *assimilation*. "To assimilate" means literally to make *like*, and Piaget uses the term quite literally.

> At the beginnings of assimilatory activity, any object whatever presented by the external environment to the subject's activity is simply something to suck, to look at, or to grasp.

> In its beginnings, assimilation is essentially the utilization of the external environment by the subject to nourish his hereditary or acquired schemata (Piaget, 1971, pp. ix and 396).

(Note that Piaget used the words "schema" and "scheme" differentially; here the translation should be "schemes.")

As observers, we may legitimately speak of the organism and its "external environment," but the organism cannot make that distinction with regard to itself; it merely has its own experience. Hence, from the organism's point of view, to assimilate means to modify a present experience so that it fits a hereditary or acquired scheme, i.e., a perceptual or motor pattern that already has, in some sense, the character of an invariant. In other words, invariants create repetition as much as repetition creates invariants. This may not be nearly as paradoxical as it sounds. The linguistic example of names may once more help to illuminate the point. Having established four-leggedness as the invariant critical feature of the complex experience associated with the word *dog*, the child focuses on four-leggedness and uses the word *dog* whenever that feature is available among the experiential material. That means that the child will assimilate all sorts of items—many of which he would later call *cat, horse, sheep* or *cow*—and in doing so, he will disregard the experiential elements that might distinguish them from the original experience associated with the word *dog*.

Needless to say, this is an extreme simplification. To my knowledge no child has ever assimilated a chair, a sofa, or a kitchen table to the concept named *dog,* and that should alert us to the fact that the child's concept of *dog*, by the time he begins to name things, must involve more criterial features than just four-leggedness. However, the principle is useful, and it helps us visualize not only how over-extension of a word or concept can develop, but also how the child eventually replaces it. As Brown pointed out, parents will require the child to narrow his use of "dog" to its proper range. That is, the child's misuse of the word will create disturbances in his experience there will be unexpected sequels. Such negative feedback can be eliminated only by a modified use of the word, in this case by its restriction.

In Piagetian terms, such a reduction of over-extension is a case of accommodation. Assilmilation we have said, is the application of an established invariant pattern or scheme to a present experience regardless of discrepancies. In accommodation, on the other hand, a discrepancy leads to the formation of a new pattern (either a modification of an old one or a novel assembly) that may then become a new invariant. The question that immediately arises is this: why should a discrepancy in experience sometimes lead to accommodation and thus to the creation of a new scheme, and at other tines, in assimilation, be disregarded? The answer is not too difficulty provided we view the organism as a fundamentally goal-directed system.

Piaget suggests this frequently by saying that the organism must be considered an active experiencer rather than a passive receiver of stimuli. In the excerpts quoted above, he is even more specific: the organism assimilates items *in order to* suck, look at, or grasp. These activities, like all others which the organism has or acquires, have a certain sequential pattern and usually lead to certain experiential results. They are procedures toward certain experiential goals. But to be attained, these goals require the support (i.e., the presence) of more or less specific elements of experience. And there may be occasions when the elements present could not be assimilated to conform to the expected results of the activity.

When an infant, for instance, assimilates some visual elements to the invariant pattern that, for him, constitutes a rattle, and grasps and shakes a piece of wood that happens to be within reach, then the absence of the auditory element expected to ensue may cause a discrepancy that cannot be eliminated by assimilation. In that case, attention is likely to be focused on any of the formerly disregarded visual or tactual elements by means of which the piece of wood could be discriminated from the rattle. Once the discrimination has occurred, the new elements, with or without some of the old ones, can be associated in an act of *accommodation* to form a novel scheme. This novel scheme, from then on, will serve as a relatively independent invariant for the assimilation of future experiences.

I hope this brief exposition of the complex interaction of assimilatory and accommodatory processes has indicated that this part of Piaget's theory is compatible with the cybernetic approach. To refer once more to the feedback model, one might say that assimilation, insofar as it adjusts sensory signals, reduces the generation of error signals. Accommodation, on the other hand, occurs only when there is a discrepancy or disturbance for which the organism does not yet have an established remedy.

Schemes and Conceptual Structures

In discussing the role of assimilation in the generation of "sameness," and that of accommodation in the generation of invariants, I have referred to "elements of experience" as though these elements were always perceptual data. That was an extreme simplification.

"The object is in the first instance only known through the subject's actions, and therefore must be itself constructed" (Piaget, 1972, p.82). For Piaget, early instances of "objects" are always subsections of an action scheme. They are the sensory schemes which, in conjunction with a motor scheme, constitute a sensorimotor activity. As such they are always a compound of perceptual as well as proprioceptive data. That is to say, they are a scheme composed not only of several sensory signals but also of signals in several sensory modes. Usually this means that they contain visual and tactual signals as well as proprioceptive signals deriving from the motor activity of the perceiver.

As the result of many acts of accommodation that added or removed particular experiential elements, an object-scheme becomes relatively invariant and may be used to assimilate new experience. But all this still takes place on the level of sensorimotor activities and, though it may serve as partial model for later developments, it does not entail the formation of concepts. Hence this use of an invariant scheme is by no means a manifestation of the concept of object permanence, because its invariance arises from and consists in the repetition of an activity and does not yet involve the invariance of an independent object.

> The growth of the human mind partly consists of the successive attainment or formation of cognitive *invariants*. As its name suggests, an invariant is something that remains the same while other things in the situation change or undergo various transformations. The identification of constant textures

> or invariants in the midst of flux and change is an absolutely indispensable cognitive activity for an adaptive organism, and it is particularly characteristic of human rationality (Flavell, 1977, p. 48).

Cognitive invariants may often have the same content as invariant sensory or motor schemes, but we attribute to them an additional feature that takes them out of the world of sensorimotor experience: we assume that they are available to the organism regardless of the sensory signals at the moment. They constitute part of the material on a second level of experience that is made possible by memory, a level of experience that we call *representation* to differentiate it from the level of perception and proprioception.

Introspectively we *know* that we can operate on a representational level. Empirically we know that people can solve quite complicated problems in their heads (i.e., without perceptual crutches, such as pencil and paper) and that some of them can even play several games of chess simultaneously without any visual aids whatever. Nevertheless I do not want to say that adults or children have representations as a matter of fact. All I intend is that the kind of model a cognitivist constructs to "explain" the functioning of such organisms must have that capability.

Once a system is equipped with this ability to maintain experiential compounds invariant and to use them independently of present sensory experience, all sorts of interesting things become possible. For the moment I shall mention only two. First, the term *invariant* acquires a new dimensions. In the context of sensorimotor assimilation, it was always a perceptual or motor activity that was called invariant because it did not change in the face of different sensory material. Now, if the invariant can be used on the representational level, without an activity, it becomes like a program or a subroutine that is invariant in that it is stored somewhere in a memory from which it can be retrieved. It is this change of status that gives rise to the concepts of permanence and of *identity,* a further step in the construction of permanent objects.

The second development made possible by the introduction of the representational use of invariants is that they can now be used as building blocks for conceptual constructions that move further and further away from the raw material of sensory or motor signals. This shift constitutes one of the salient characteristics of all the "higher," more sophisticated mental operations and it has consequences for epistemology far beyond the scope of this chapter. But, the principle of (1) learning to construct a composite in a certain way and out of certain elements, (2) storing the program or recipe of construction, and (3) retrieving it as a unit to combine with others of parallel origin and form a "higher-level" structure, *without* having to return to the "lower" level, has proven to be one of the most powerful in the construction of knowledge. It allows us to proceed much as a bricklayer, who can devote all his energy and attention to the creation of a wall or an arch, without ever stopping to ask where the bricks he is using came from or how they were made. And just as the characteristics of the bricks (e.g., shape and size) make it impossible for the bricklayer to build certain structures, so the ready-made conceptual building blocks impose constraints on any future construction.

The Ambiguity of Sameness

In everyday situations, when we say that an item is *the same,* we know what we mean and we are rarely misunderstood. If two or more items are present, there is no problem. We are saying that, in a way usually defined by context, we find no difference between the items and therefore consider them *equivalent*. If, however, only the one item is within our actual field or experience, the expression is ambiguous. It may be interpreted again as an assertion of equivalence, but it may also be interpreted as asserting that the present item is the *selfsame* individual that we have encountered at some other time.

The same ambiguity is inherent in the everyday use of the word *identical*. Although English provides a logically impeccable way of distinguishing the two meanings by the use of a different article (*an identical one* for equivalence, *the identical one* for individual entity), they are frequently interchanged in ordinary usage. As long as we are referring to fairly familiar items, this does not seem to lead to confusion.

In a discussion of conceptual invariants, however, the ambiguity of “sameness” and “identity” becomes a serious obstacle to understanding. Indeed, if we want to grasp the concept of object permanence, it is indispensable that we completely resolve that ambiguity.

Earlier, we asked the question, “What constitutes the invariant object that the organism *recognizes*?” If we take this question without context, “invariant” clearly could be interpreted in two radically different ways. On the one hand, it could be a prototype, or template, by which the organism categorizes certain experiences as exemplars of the class represented by the invariant. This is the sense of *object concept* and it was then illustrated by the example from psycholinguistics. On the other hand, the “invariant” could be interpreted as an object in its own right that remains unchanged because it “exists” and is recognized as the *selfsame* individual every time it enters the organism’s field of experience. This is the sense of “invariant” that corresponds to the conception of *object permanence*. Both the concept of the object as prototype, with regard to which experiences may be considered equivalent, and the concept of object permanence, as a result of which two or more experiences may be considered to derive from one identical individual, involve a form of invariance. But the invariance is certainly not the same in both cases. William James, who had an exceptionally keen eye for conceptual distinctions, said 70 years ago,

> Permanent “things” again; the “same” thing and its various “appearances” and “alterations”; the different “kinds” of thing … it is only the smallest part of his experience’s flux that anyone actually does straighten out by applying to it these conceptual instruments. Out of them all our lowest ancestors probably used only, and then most vaguely and inaccurately, the notion of “the same again.” But even then if you had asked them whether the same were a “thing” that had endured throughout the unseen interval, they would probably have been at a loss, and would have said that they had never asked that question, or considered matters in that light (James, 1907/1955, p. 119).

In this very limited, specific, but important respect, much of the literature on objects, identity, and the concept of self still seems to avoid that question and to shirk "considering matters in that light."

Equivalence and Continuity

The process of assimilation discussed earlier leads to practical, action-bound implementations of the notion of "the same again" in that it actively shapes a present experience to fit an available sensorimotor scheme. At that level one could speak of object-*schemes*. The term object-*concept*, however, should be reserved for constructs on the representational level, i.e., constructs that can be called up regardless of the sensory elements that are or are not available at the moment. If we accept this distinction (and if we want to attribute cognitive processes and operations to an organism, we must accept it), then also the *concepts* of recurrence, sameness, identity, and permanence will have to be constructed on the representational level. An attempt to map this construction in detail seems out of the question at present, but we can, I believe, outline some of the major steps.

No recurrence can possibly be established unless there are records of past experiences and the possibility of surveying them in some way. That requires not only memory and retrieval capabilities (which I shall take for granted), but that the experiencing organism can switch his attention from "present" items to the records of "past" items. It is only by switching from one item to another that absence of difference can be established, with the result that the two experiential items are the same. Eliane Vurpillot (1972) has elegantly documented the switching to and fro of children's eyes during visual comparison tasks. Eye movements indicate shifts of attention in the visual field. Shifts of attention, however, have also been observed when eye movement is eliminated by stabilizing the visual image (Pritchard, Heron, and Hebb, 1960; Zinchenko and Vergiles, 1972). Hence we may safely assume that attention can also shift between items when some or all of them are representational.

An analysis of the actual procedure of comparison also has conceptual implications. If the *direction* of comparison is from item A to item B, in the sense that the characteristics found in A are then checked in B, item B may be considered the same (in the likeness sense) if no difference is registered in the checked characteristics. But in that case B may or may not have characteristics which are not represented at all in A. Hence it should not be called "equivalent," let alone "identical." In order to establish *equivalence,* the comparison would have to be carried out in both directions. This distinction is of considerable practical importance, since it is all too easy to overlook the fact that in classifying or categorizing, as a rule, a one-directional comparison is all that is made.[*]

The direction of the comparison in classification, clearly, is from a representational prototype, or concept, to an experiential item. The nature of the class depends on precisely those characteristics that have been abstracted and combined to form the concept. (This could again be illustrated by the example of the child that abstracts, say, four-leggedness and spontaneous motion to form a concept and then—erroneously, from the adult point of view—associates the word dog with that concept.)

The procedural analysis also helps to throw some light on the still mysterious ambiguity of "sameness." To establish that two experiential items are equivalent, the simplest and most reliable way is to arrange them side by side in our visual field.[1] We can then switch our attention from one to the other and check out all the characteristics we consider relevant. This is precisely what Vurpillot demonstrated in her study, and she also noted the one peculiar feature in the procedure. “By convention, one difference is always excluded from the list of properties [to be compared] and that is the object’s location relative to the subject. Since they [the objects] can never appear at the same place at the same time, they will always be different from that point of view” (Vurpillot, 1972, p. 311; my translation).

In other words, when we say two items are “equivalent in every respect,” we understand that “in every respect” does not include spatial location. It is also clear that we accept the fact that the two items do *not* “appear at the same place at the same time” as evidence that they are not one and the same individual, even if we can find no other difference between them. This, I think, brings out very clearly the difference between the two meanings of “sameness.” The construct of equivalence involves two experiential items and comparison of a set of properties that does not yield a difference. Ideally, we have both items in our field of experience, separate but in a way that allows us to shift our attention from one to the other in *both* directions.[3] But if this ideal can not be achieved, because for some reason it is impossible to have both items in the experiential field, we may suspect that we are dealing with one individual only. There is, indeed, no foolproof way of ascertaining that this is not the case. That raises the question of how we can consider two experiential items as one, even when they are separated by other experiences.

At the core of this problem of individual identity is the conception of continuity. As long as an item remains within our field of vision and is marginally attended to, there is no problem. Continuity is the uninterrupted succession of signals from one source. If the item is a composite of signals from different sources, each of them will be constantly available to attention (which is *not* the case with alternating lights and, therefore, leads to the phi phenomenon; see note 2 at the end of this chapter).

But consider a case in which there is no continuous succession at all but, nevertheless, we are able to construe individual identity. A well-fed brother whom one has not seen for 20 years may be bald and scrawny when he returns; he may have a different accent, his likes and dislikes may have changed, and what he now says about politics, art, and women may be incompatible with what one remembers of him. Yet one could still accept him as the self-same individual. How do we construct continuity across such enormous experiential gaps? I believe we acquire the ability in small steps.

The first step is to assume continuity of a composite whole on the strength of an experientially continuous part. We do this every time we watch a moving object that for a moment partially disappears and then comes into full sight again. In an infant’s early life, that is a frequent experience, since there are nearly always some visual obstacles in the immediate environment behind which parts of people disappear. Visual tracking is manifest very early and soon enables the infant to follow an item even when it wholly disappears for a moment (Bower, 1974). In that case it cannot be a visual part of the experiential item. Rather it is the proprioceptive signals generated by the tracking motion that supply the continuity. The essential feature, however, is the experiential continuity of *some* signal sequence that connects the percept that disappears with the percept that reappears, and that can hold the child’s attention so that no other item comes into focus. If there is no such sequence and, consequently, there is a refocusing of attention in the interval, the two experiential items will not be construed as one individual, no matter how similar they may be as percepts. For many five-year-olds, for instance, the sun today and the sun yesterday are not yet one and the same individual (Piaget, 1971, p. 87).

As long as the linear sequence of attention focused on sensory signals is the only dimension of the child's experience, it is logically impossible to connect two experiential items across an interval during which *none* of the signals constituting them is continuous. Such a connection has to be created *outside* the ongoing experiential sequence, so that it can subsist, as it were, in parallel and is not broken by the actual sensory experiences that occur during the internal. This second dimension is the representational one, and hand in hand with its development goes the process that Piaget has called *externalization*. He speaks of a "miniature Copernican revolution" at the end of the sensorimotor period, as a result of which the child begins to see himself as a permanent object among other permanent objects "in a universe that he has gradually constructed himself, and which hereafter he will experience as external to himself" (Piaget, 1967, p. 9).

It is easy to see that this externalization of items has a momentous effect on the way we thenceforth think and speak of experience and knowledge. It accomplishes a conceptual revolution that is almost impossible to undo. Perhaps the most remarkable thing is that the logical analysis that inexorably reduces regularities and continuity in our experience to our own constructive activities comes to a halt in the case of the returning brother. For if we do accept him as the selfsame individual who left 20 years ago, it is not on the basis of equivalences we establish on the level of immediate sensory experience, but only because he is able to recollect and communicate particular experiential items, such as objects, situations, and events, that the person of 20 years ago had occasion to construct. His individual identity, in the last analysis, depends not on the permanence of what we usually call physical characteristics, but on the permanence of his recollections or, if you will, his mind.

The Problem of Realty

The outside realm into which we place items *as we represent them to ourselves* when we are not experiencing them—a kind of limbo where ready-made objects await entry into the process of our assimilatory experience—inevitably turns into what philosophers call "ontological reality." As soon as we conceive of continuity as an inherent property of things, we have laid the foundation for a world that "exists," a world that is "there" whether we happen to perceive it or not, a world that, ultimately, is wholly detached from the experiencing subject.

Instead of becoming aware of ourselves as creators of continuity and thus of recurrence and regularity in the flow of our proximal signals, we begin to attribute continuity and permanence to our constructs. They become an independent distal "reality," and our acts of experience take on the character of exploration and discovery of things that are already there quite apart from ourselves. As a result of our detachment we find ourselves facing that most peculiar, unanswerable question of traditional epistemology—how we, irrevocably tied to our ways and means of experiencing, can ever transcend that limitation and acquire "true knowledge of the real world."

Some Basic Constructs

The preceding pages have laid out some implications of Piaget's theory of early cognitive development that are rarely emphasized because they require a rather drastic modification of our common-sense ideas. To conclude this section, I shall try to indicate very briefly how some of the concepts indispensable in our construction of reality derive from the basic assumptions in this early development.

Space. The early continuity the child constructs, for example, on the basis of his own tracking motion, is the continuity of a scheme and does not involve a conception of the disappearing object continuing its path behind the visual obstacle. Once that conception is formed by externalizing a representation, however, the external permanent object requires a place, somewhere *beyond* the experiential field, where it can "exist" when it is not being experienced. Such a place, before it is particularized by spatial relationships to other items or to the experiencer, constitutes the basic concept of space. It's one and only feature is that it allows an arrangement of perceptual items such that attention can shift from one to the other in both directions and at the discretion of the perceiver. When this shift of attention is actually carried out in perception, it is accompanied by proprioceptive signals from eyes, hand, or some other part of the body (indicating what an observer would call "motion") and this addition of signals leads to the concept of extension.

Change. Once the continuity of an item is established by means of an experientially continuous part or through the continuity of related proprioceptive signals, a comparison of all the characteristics of the present item with those of the record from the time before the interval may yield a difference. If we then want to maintain the item as a continuous one, we shall consider it *changed* but, nevertheless, the same individual. The more characteristics are changed, the more difficult it becomes to maintain the continuity of the individual, and we are driven to anchor the continuity in more and more abstract (i.e., less sensory) characteristics. An ice cube that has melted, or a log that has turned to ashes in the fireplace, tax our conception of continuity. In cases where there is even less to hold on to, we try to make do, as did Aristotle, with the paradoxical construct of "substance," which supports individuality without having itself any properties at all. The search for this mysterious substance must forever fail because, having no properties, it cannot be found on the level of sensory signals.

Motion. A very particular kind of change is the one that involves not a characteristic of the permanent object itself, but some relationship with another item in the experiential field. According to whether relationships (e.g., inclusion, contiguity, proximity) are being terminated, maintained, or newly established, we get *change of location* or *change of extension* of the object. Any repetitive change that involves the forming and ending of relationships with a succession of different items will be considered some form of *motion*.

Since the origin of these concepts is, under all circumstances, a change of relationships, the attribution of motion to a particular item always requires a further point of reference or other sensory signal.

Cause. When there is change and we explicitly register a difference between the two experiences, it is obviously more difficult to maintain the item's individual identity than if there is no change. Indeed, we are construing the two experiences as one individual on the one hand, and as different on the other. I believe it is this apparent clash which later drives us to search so persistently for causal explanations.

The practice of causation originates with the infant's "circular reactions" (Piaget, 1971, p. 351), the *instrumental* repetition of certain actions because they have, in the past, led to satisfying consequences. The infant shakes the hand holding the rattle because this has been followed by somehow pleasing sensory signals. Similarly, the infant learns to pull a cloth in order to bring an object that is on it within reach. Once some of the sensory compounds involved in these situations have been externalized as permanent objects, the pleasing events can be seen as *changes* and the activity that regularly precedes such a change can be conceived as causing the change. On the conceptual level, then, once a cause is found, the clash between the maintained identity and the registered change is neutralized, because the change noted in the second experience can now be balanced by the cause as adjunct to the first experience. From there to the general conception that every change can be reduced to *some* cause, and that this cause need not be an activity of the experience does not seem a verse great step.

Time. The earlier statement that continuity derives from a succession of sensory signals from one source was simplified by leaving out a difficult but important point: a succession of similar signals cannot be perceived as a sequence unless the experiencer's attention takes in something else *between* the signals. If there is no such break, the signals, because they are all the same, cannot constitute a plurality and will be registered as one. On the other hand, if attention is interrupted by a total shift to something else, the signal before the break and the one after it will be perceived as separate, and there will be no *sensory* basis for the conception of continuity. The answer to this problem lies in the possibility of marginal (or, perhaps, divided) attention. We know this in the visual field, where attention can dart to the periphery *without* motion (i.e., without the addition of proprioceptive signals) and we can experience it also tactually by keeping track (in a dark room, for example) of the edge of a table with one hand, while we mainly focus on exploring the table's surface with the other. That is to say, we can, in fact, operate in two channels.[3] *Continuity* and sequence, thus, both spring from the juxtaposition of two successions of signals that are separate in the experiential field but interrelated by attention; the one is continuous relative to the other, the other is sequential relative to the first.

The concept of *time* arises when the sequential items are mapped or projected on the continuous one. Indeed, in order to construct the temporal dimension, we always need two processes—the ticking of the clock as one and, as the other, the hum of the traffic, the glow of a light, or anything else that we experience as continuous. With the abstraction of our *self as* experiencer from the flow of experience we create the most continuous of all continuous items on which we can project any succession of experiential elements as a temporal sequence. But, as we shall see in the next section, this continuous self is destined to remain a purely conceptual construct that we cannot encounter on the sensory level.

The Construction of Self

Self Observing Itself

The concept of "self" seems simple enough when we refer to it in an accustomed context and in ordinary language. As a rule people do not object if one makes statements such as "That's typical of James" or "Well, I can't help it, I *am* like that." Even the rather peculiar expressions "You are you" and "I am I" do not seem as peculiar as Gertrude Stein's "A rose is a rose." What we apparently have in mind when we make any such statement is the individual identity or continuity of a person. However, as soon as we attempt to analyze what precisely it is that constitutes the continuity of our "selves," we run into difficulties and get the impression that there is an ambiguity. The "self" seems to have several different aspects.

First of all, there is a self that is part of one's perceptual experience. In my visual field, for instance, I can easily discriminate my hand from the writing pad and the table, and from the pencil it is holding. I have no doubt that the hand is part of *me,* while the pad, the table, and the pencil are not.

Second, if I move my eyes, tilt my head, or walk to the window, I can isolate my "self" as the locus of the perceptual (and other) experiences I am having. This self as the "experiencer" appears to be an active agent rather than a passive entity. It *can*, in fact, move my eyes, tilt my head, change location—and it can also attend to one part of the visual or experiential field rather than to another. This active self can decide to look or not to look, to move or not to move, to hold the pencil or not to hold it and, within certain limits, to experience or not to experience.[4]

Beyond these, there are still other aspects of the concept of "self." There is, for instance, the *social* self: *qua* experiencers, we enter into specific relations to other experiencers, and *qua* actors, we adopt specific patterns or roles that eventually come to be considered characteristic parts of what we call our "selves." But this chapter focuses on the early development of concepts, and I shall disregard the social aspects of the self because, although they, too, have some roots in the sensorimotor period, their main development seems to take place during adolescence.

In the paragraphs that follow, then, I shall be concerned exclusively with the self as perceptual entity and as experiencer.

The Self as Visual Percept

Perhaps the most serious obstacle that has impeded traditional psychology from arriving at a plausible analysis of the concept of "self" is the assumption that the dichotomy between an organism and its environment is basically the same as the dichotomy between an experiencing subject and what it experiences. As argued in the section "Observer and Observed," the distinction between an organism and its environment can be made only by an observer of the organism. The organism itself has no access to distal data, to items outside itself. But with the construction of permanent objects, the organism *externalizes* some of the invariants it abstracts from its experience and treats them, from then on, as independent external items (see "Equivalence and Continuity," above). This externalization, as we have seen, goes hand in hand with the establishing of *internal* representations or concepts, and this dual development of objects, which are "perceptual," and concepts, which are "representational," leads to a sharp division between two forms of experience, one "external" and the other "internal." (There are, of course, illusions, dreams, and hallucinations that, from the subject's point of view, blur that division.) Both the internal and the external, however, are explicitly *experience,* and the division between them, therefore, is a division between two types of experience and *not* the division between an experiencing subject and the objects it experiences.

Wapner and Werner 1965, p. 10) are aware of the problem and speak of two "complementary notions," a holistic one and a polar one.

> Our theoretical-experimental approach has focused on two characteristics of this relationship between one's own body and environmental objects. First of all, we assume that there can be no perception of objects "out there" without a bodily framework and, conversely, we assume that there can be no perception of the body-as-object without an environmental frame of reference. Thus, one basic feature of this "body:object" relationship pertains to the interaction constantly going on between them. The central notion here is that the appropriate unit to be dealt with is not the organism per see, but rather, the organism in its environmental context, this conceptualization, then, is that the variability or stability of the biological unit, "body:environment," reflects itself in body perception as well as in object perception.
>
> Second, complementary to this holistic notion of the biological unit composed of body and environment is a feature which is seemingly in contradiction to it, viz., the feature of oppositeness, or separateness, or polarity between these two elements. Such oppositeness is characteristic of the normal adult insofar as he experiences the world and himself as standing at polar distinction in each other.

I have quoted the two authors at some length because they seem to be more aware than any others I have read of the experiential origin of the "body percept" but they still represent the general confusion between the organism's perception of its own body and the conception of the "self" as the locus of experience.

The first step that leads to that confusion is the failure to separate what *only an observer* of organisms can say about organisms and their environments from what an organism itself may say about its experience. (See "Observer and Observed," above.)

When I visually distinguish a hand from the writing pad and the table on which it lies, I carry out exactly the same kinds of operations as when I distinguish the coffee cup from the table on which it stands, or the picture from the wall on which it hangs, or the cardinal outside my window from the branch on which it happens to be perched and from the rest of the landscape. In all these cases I am recognizing certain objects to which I have attributed relative consistency (closure) and permanence. Having successfully externalized permanent objects, I am now experiencing them as parts of "distal reality." From the purely visual point of view, the operations by means of which I separate objects from the rest of the visual field or "ground" are always the same kind. And the observer's distinction between an organism and its environment is *normally* made in the visual field (which is not to say that such a distinction *could* not be made in the tactual mode). Thus, although we can visually distinguish birds, coffee cups, tables, and hands from the rest of the visual field and from one another, it seems clear that a naive organism (i.e., an organism such as an infant that does not yet have a great deal of intermodally coordinated experiences) cannot *visually* discriminate between *a* hand and his *own* hand.

Trevarthen et al. (1975) have shown that, contrary to previous assumptions, infants in the very early stages of reaching and grasping do *not* make use of visual feedback concerning the hand's position. This seems likely because it takes an infant some time to relate the visual image of the hand to the "self" that has the motor command over it.

If one accepts the conclusion that there are no a priori *visual* features that differentiate visual experience of one's own body from visual experience of other items, one is at once compelled to raise the question how that differentiation could be made. As we shall see, there are several levels on which the coordination of sensory signals yields invariants that contribute to the constitution of a "permanent" entity corresponding to what we may call our "body percept." Before surveying these levels, however, I want to emphasize once more that, no matter how successful our analysis of the body percept may be, it cannot possibly tell us anything about the "self:world polarity" mentioned by Wapner and Werner (see quotation above), because insofar as the perception of one's self is the result of perceptual experience, it belongs to the world that is being experienced and not to the experiencer.

Levels of Self-Perception

If the Piagetian approach to the notion of object permanence is a viable one, it should not surprise us that the notion of self as a constant perceptual entity cannot be derived from vision alone. According to that theory, permanent objects are the result of the coordination of signals from more than one sensory source.[5] Since the body percept obviously *does* achieve the status of permanent object, it must be multi-modal, and the question now is: what kinds of signals and coordinations of signals would enable an organism to differentiate one permanent object—his own body—from all the other permanent objects that have been constructed and externalized?

The answer is extremely complex; there are many factors that contribute to the differentiation and isolation of the body percept. In this summary exposition I shall sketch out a few of the points that seem crucial. Such an analysis is necessarily made by an observer who can only hypothesize what goes on in the black box we call an observed organism. (See "The Use of Black Boxes," above.) The indispensable limitation of this hypothesizing is that the organism can operate only with its own proximal data, i.e., with signals that can be supposed to originate within it rather than with "information" originating in what from the observer's point of view is the organism's environment. I would also like to emphasize that this analysis is provisional and lays no claim to being definitive, let alone exhaustive.

One of the primary factors seems to be the experience that there is motor control over certain visual items. When a discriminable item moves partially across an organism's visual field and the organism follows it with its gaze, the organism can correlate certain visual signals with the proprioceptive signals from its own tracking motion (eye and head movements). If the visual item happens to be, say, the organism's own hand, another kind of signal can also be correlated with the visual and tracking signals; namely, the proprioceptive signals generated by the active motion of the hand. This additional correlation can be used to discriminate moving objects that are parts of the organism's own body from objects that are not.

This very elementary distinction is, of course, strongly reinforced by the fact that the organism is able to generate and control (i.e., direct, speed up, stop, etc.) the motion of its own limbs. This particular coordination of motor control on the basis of visual feedback, for instance in hand movements of human infants, is a difficult task and, as a rule, is not mastered until an age of six or seven moths (Bower, 1977).

There are various experiential phenomena that confound these basic coordinations. If the moving hand happens to hold an object, it must be discriminated from the hand on the basis of *other* factors, such as tactual signals from the fingers because, for the coordination of visual and proprioceptive signals, the held object is initially indistinguishable from the hand. Once the organism has a certain amount of experience, it will be able to discriminate its hand from other items on the basis of visual signals alone. This ability, however, consists in recognizing the hand as the one categorized as its *own* by nonvisual means.

The interplay with tactual signals is presumably essential for the evolution of a primitive visual and proprioceptive scheme of the body percept. Every contact with other items that gives rise to tactual signals is an indication of the limits of the body. The progressive coordination of these "contact signals" with the accompanying visual signals is, in fact, the essential element in the organism's mastery of locomotion and other motor skills (Held, 1965/1972).

Tactual signals are also involved in another aspect of the body percept's ontogenesis. When an organism touches some part of his body with his hand, tactual signals are generated on both sides of the point of contact. This allows the organism to distinguish with great reliability between touching other things and touching himself. If the touch is vigorous, it may even give rise to pain. Kittens chase, catch, and occasionally bite not only their mother's and littermates' tails, but also their own. There is no question that they quickly learn to distinguish their own tail when it comes to biting, and to that extent they come to have a notion of "self'.

Similarly, the baby that sucks his thumb as well as the nipple of mother or bottle already has the sensory material to make a distinction that will be a primary source for the construction of both the body percept and the concept of self. And the same could be said of the grooming behaviors that nonhuman primates apply to their companions and themselves.

One further situation involving tactual signals may well supply the operational prototype for construction of the concepts of motion and time, and the sensory model of the continuously experiencing self. In the paragraph on basic concepts, I said that any repetitive change that involves the forming and ending of relationships with a succession of different items will be considered a form of *motion*. Under the heading *time*, I said that continuity and sequence both spring from the juxtaposition of two successions of signals that are separate in the experiential field but interrelated by attention. The one is continuous relative to the other, the other is sequential relative to the first. Take a finger of your right hand and run it along your left forearm: the tactual signals originating in your finger will be a homogeneous "continuous" succession because the receptors from which they come remain the same; the tactual signals originating in your left arm, instead, will constitute a sequence of different signals because they come from different receptors. If you consider this second set of signals as a sequence of different locations with which your finger establishes and terminates contact, you will conceive of your finger as *moving*. If you consider them equivalent units linked into sequence by the continuous signals from your finger, you will conceive of them as points or "moments" in *time*. In this second case, the finger of your right hand supplies what is perhaps the closest sensory-motor analogy to the continuity of the experiencing subject that we call our ""self."

The Image in the Mirror

On a later and more sophisticated level, once tactual and proprioceptive elements have contributed to a protoconcept of her own body, the child will be capable of visual recognition of her own hand or limbs. This sets the stage for a considerably more complex experience of the physical self: the child's recognition of her own shadow, her reflection on a shiny surface, and her image in a looking glass.

Gallup (1977), in a survey of research on self-recognition in primates, comes to the conclusion that only the great apes have the ability to recognize their mirror image as their own. Monkeys and, as many of us have observed in our homes, cats and dogs, quickly learn to discriminate their shadows, reflections, and mirror images from *other* moving objects or animals, but do not appear to relate them in any way to themselves [6]. The simple synchrony of movement between, say, a paw and its shadow or reflection does not seem sufficient to establish the link. It may be that a causal connection must be constructed from a deliberate motion to its reflected counterpart, and that it is this connection which differentiates the motion of a mirror image from the motion of another object or organism.

The child who stands in front of a looking glass, sticks out his tongue, and contorts his face into all sorts of grimaces gets a constant confirmation of this causal link. The mirror image is as obedient as his own limbs and can, thus, be integrated with the body percept, expanding it by providing visual access to otherwise invisible aspects. And like the body image, it is a visual percept, an item that is experienced not the item that does the experiencing. This central item, the experiencer himself, remains mysterious. Without ever perceiving it, we know that it is at the heart of whatever continuity or invariant we construct in our perceptual world. As teenagers, at one time or another, many of us stood in front of a looking glass and wondered: where am I? it is a question we are still unable to answer when we are adults.

The Cybernetic Metaphor

Sketchy and incomplete though they are, I hope the preceding paragraphs have shown that there are several relatively independent sources from which facets of the self percept can be developed. Much remains to be worked out; above all, the detailed analysis of the process by which these facets become integrated into what we so strongly feel to be a unitary concept of our self. While we can work out a plausible model for the self as an entity of our sensorimotor world of experience, this model cannot throw any light on what we feel to be our self as *experiencing entity*. The reason lies in the very structure of our conception of knowledge. In the Western tradition of science and rational explanation, knowledge by its very nature requires a dichotomy between the knower and the things he knows. In other words, we can come to know only what we consider to be in some sense separate from our knowing selves. By questioning something, by the very act of asking what it *is*, we have already set our self, the questioner, apart.

In the realist view, the self we perceive, by being perceived, becomes the object of a perceiving subject. In the constructivist view, the self we conceive is necessarily the product of an active subject that remains outside the construction. It may be a viable construct in that it appears compatible with *what* we experience, but it does not and cannot incorporate that primary act of constructing itself. Berger and Luckmann (1967, p. 50) express this very neatly:

> On the one hand, man is a body, in the same way that this may be said of every other animal organism. On the other hand, man has a body. That is, man experiences himself as an entity that is not identical with his body, but that, on the contrary, has that body at his disposal. In other words, man's experience of himself always hovers in a balance between being and having a body, a balance that must he redressed again and again.

The two aspects seem wholly incompatible. The paradox of the self experiencing itself, from the logician's point of view, is analogous to the paradox of self reference. The logical paradox has recently been approached with great success through a novel interpretation of the concept of recursion (Varela, 1976), and it is surely no accident that the very same concept of recursion has opened an equally novel path towards the logical interpretation of "permanent objects" (von Foerster, 1976). An exposition of the formal intricacies of these achievements would be beyond the scope of both this chapter and my competence. One general point, however, brings this discussion back to the place where it began.

Even a very simple system that regulates itself by negative feedback (such as Philon's oil lamp of two dozen centuries ago) must be characterized not by a description of its component parts but by a description of their circular interaction. The essence of the system, its individual identity that perpetuates it, cannot be ascribed to any particular part or property, nor can it be located in a particular point. The identity of such a system resides exclusively in the invariant that is the result of mutually balanced changes. The characteristic feature of Philon's lamp is not that its burner always remains full, but rather the paradox that lowering the oil level causes the oil level to be raised.

As a metaphor—and I stress that it is intended as a metaphor—the concept of an invariant that arises out of mutually or cyclically balancing changes may help us to approach the concept of self. In cybernetics this metaphor is implemented in the "closed loop," the circular arrangement of feedback mechanisms that maintain a given value within certain limits. They work towards an invariant, but the invariant is achieved not by a steady resistance, the way a rock stands unmoved in the winds but by compensation over time. Wherever we happen to look in a feedback loop, we find the present act pitted against the immediate past, but already on the way to being compensated itself by the immediate future. The invariant the system achieves can, therefore, never be found or frozen in a single element because, by its very nature, it consists in one or more relationships—and relationships are not *in* things but between them.

If the self, as I suggest, is a relational entity, it cannot have a locus in the world of experiential objects. It does not reside in the heart, as Aristotle thought, nor in the brain, as we tend to think today. It resides in no place at all, but merely manifests itself in the continuity of our acts of differentiating and relating and in the intuitive certainty we have that our experience is truly ours.

Notes

* I nevertheless continued to use the term “equivalence” in the sense of: equivalent with regard to the properties examined in a comparison.

1. I am here limiting the discussion to the visual aspects but it should be clear that it applies equally to all sensory modalities, singly or in combination.
2. Note that in the phi phenomenon, where two lights flash in quick alternation, the viewer’s attention, though shifting in both directions, cannot do so at its own rate but is obliged to keep pace with the lights, this destroys the “twoness” and results in the perception of one moving light.
3. A manifestation of “parallel processing” in the auditory mode is found in musicians who in a fugue, for example, are perfectly able to keep track of two rhythmically different sequences. Of great interest in this regard is also the recent work on levels of awareness by Hilgard (1974).
4. While Oriental philosophy has always cultivated this autonomy of the experiencer, the Western world, in defense of its traditional belief in an objective reality, has tended to consider experience as obligatory, inevitable and rather passive.
5. In accepting this view it is essential to realize that there are certain intermediary phenomena between the simple association of, say, visual signals on the basis of their contiguity, and the fully fledged scheme of a permanent object composed of signals of multimodal origin. Permanent two-dimensional shapes, for instance, are the result of program-like patterns in which visual signals are linked by a continuous pathway of proprioceptive data or moments of attention. (See Ceccato, 1960; Shumaker. 1977; von Uexküll, 1933/1970.)
6. This must also be so for all wild animals, since they do not take fright when they lower their heads to drink from a water hole.

References

Ashby, W. Ross. Learning in a Homeostat. *Symposium on Artificial Intelligence*. Knoxville: University of Tennessee, 1970.

Berger, P.L. and Luckmann, T. *The Social Construction of Reality*. Garden City, N.Y.: Doubleday, l966 (Anchor paperback, 1967).

Bower, T.G.R. *Development in Infancy*. San Francisco: Freeman, 1974.

Bower, T.G.R. *The Perceptual World of the Child*. Cambridge, Mass.: Harvard University Press, 1977.

Brown, R. How Shall a Thing be Called? *Psychological Review, 65*:14–21. 1958.

Cannon, W.B. *The Wisdom of the Body*. New York: Norton, 1932.

Ceccato, S. Operational Linguistics and Translation. In Ceccato, et al. *Linguistic Analysis and Programming for Mechanical Translation*. Milan: Feltrinelli, 1960; and New York: Gordon and Breach, 1962.

Craik, K.J.W. *The Nature of Psychology*. Cambridge, England: The University Press, 1966.
Feyerabend, P. *Against Method*. Atlantic Highlands: Humanities Press, 1975.
Flavell, J.H. *Cognitive Development*. Englewood Cliffs, N.J.: Prentice-Hall, 1977.
Gallup, G.G. Self-Recognition in Primates. *American Psychologist, 32*:329–338, 1977.
Hanson, N.R. *Patterns of Discovery*. Cambridge, England: The University Press, 1958.
Hebb, D.O. Science and the World of Imagination. *Canadian Psychological Review 16*:44–11, 1975.
Heisenberg, W. *Das Naturbild der Heutigen Physik*. Hamburg: Rowohlt, 1955.
Heisenberg, W. *Physics and Beyond*. New York: Harper and Row, 1971.
Held, R. Plasticity in Sensory-Motor Systems. *Scientific American*. November 1965 (Reprinted in *Perception Mechanisms and Models*. San Francisco: Freeman).
Hilgard, E.R. Toward a Neo-Dissociation Theory: Multiple Cognitive Controls in Human Functioning. *Perspectives in Biology and Medicine, 17*:301–316, 1974.
Hume, D. *An Enquiry Concerning Human Understanding* (1742). New York: Washington Square Press, 1963.
James, W. *Pragmatism* (1907). New York: Meridian Books, 1955.
Kuhn, R.A. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 2nd edition, 1970.
Maturana, H.R. Neurophysiology of Cognition. In P.L. Garvin (Ed). *Cognition: A Multiple View*. New York: Spartan. 1970.
McFarland, D.J. *Feedback Mechanisms in Animal Behavior*. New York: Academic Press, 1971.
Mayr, O. *The Origin of Feedback Control*. Cambridge, Mass.: M.I.T. Press. 1970.
Pask, G. The Meaning of Cybernetics in the Behavioral Sciences. In J. Jose (Ed). *Progress of Cybernetics*. New York: Gordon and Breach, 1969.
Piaget, J. *La Construction du Réel Chez l'Enfant*. Neuchâtel: Delachaux et Niestlé, 1937 (English translation, New York: Ballantine, 1971).
Piaget, J. *Six Psychological Studies*. New York: Random House, 1957.
Piaget, J. *The Principles of Genetic Epistemology*. New York: Basic Books, 1972.
Powers, W.T. *Behavior: The Control of Perception*. Chicago: Aldine, 1973.
Pritchard, R.M., Heron, W. and Hebb. D.O. Visual Perception Approached by the Method of Stabilized Images. *Canadian Journal of Psychology, 14*:67–77, 1950.
Richards, J., and von Glasersfeld, E. The Control of Perception and the Construction of Reality. *Dialectica, 33*:37–58, 1979.
Shumaker, N.W. Conceptual Analysis of Spatial Location as Indicated by Certain English Prepositions (Unpublished dissertation). Athens, Georgia: University of Georgia, 1977.
Trevarthen, C., Hubley, P. and Sheeran, L. Les Activités Innées du Nourrisson. *La Recherche, 6*, no. 56, pp. 447–458, 1975.
Varela, F. The Arithmetic of Closure. *Proceedings of the Third European Meeting on Cybernetics and Systems Research*. Vienna, Austria, 1975.
von Foerster, H. Thoughts and Notes on Cognition. In P.L. Garvin (Ed). *Cognition: A Multiple View*. New York,: Spartan, 1970.
von Foerster, H. Objects: Tokens for (Eigen) Behaviors. Piaget Jubilee Meeting, Geneva, 1976. Reprinted in *Cybernetics Forum, 8*, nos. 3 and 4, pp. 91–96, 1976.

von Glasersfeld, E. Radical Constructivism and Piaget's Concept of Knowledge. In F.B. Murray (Ed). *The Impact of Piagetian Theory*. Baltimore: University Park Press, 1978.

von Glasersfeld, E. Cybernetics and Cognitive Development. *American Socciety for Cybernetics Forum, 8*, nos. 3 and 4, pp. 115–120.

von Glasersfeld, E. The Concepts of Adaptation and Variability in a Radical Constructivist Theory of Knowledge. Seventh Annual Meeting of the Jean Piaget Society. Philadelphia, 1977. In I.E. Sigel, D.M. Brodzinsky and R.M. Golinkoff (Eds). *Piagetian Theory and Research*. Hillsdale, N.J.: Erlbaum, 1981.

von Uexküll, Jakob. *Streifzüge durch die Umwelten von Tieren und Menschen*. Frankfurt am Main: Fischer, 1933/1970.

Vurpillot, E. *Le Monde Visuel du Jeune Enfant*. Paris: Presses Universitaires de France, 1972.

Wapner, S. and Werner, H. An Experimental Approach to Body Perception from the Organismic-Developmental Point of View. In S. Wapner and H. Werner (Eds) *The Body Percept*. New York: Random House.

Wiener, N. *Cybernetics*. Cambridge, Mass.: M.I.T. Press, 1948/1965.

Zinchenko, V.P. and Vergiles, N.Y. Formation of Visual Images. *Special Research Report*. New York: Consultants Bureau, 1972.
