How Do You Know?

FRED ATTNEAVE *University of Oregon* [1]

Epistemology is traditionally classified as one of the fields of philosophy, and some psychologists have been content to leave it to the philosophers. In my opinion, the problem of how we know is an absolutely basic concern of psychology, but whether you consider my remarks to be psychology or philosophy I really do not care, as long as you are willing to grant me that the problem is a fundamental one.

Naively, it seems to us that the outside world, the world around us, is a given; it is just there. I look out and see you sitting in front of me; around you I see walls that enclose the room and stop me from seeing farther. But my world, the world I live in, does not seem to stop at the walls; beyond them, in the same continuous space, there are cities, roads, rivers, oceans, all of which have some determinate loci in my picture of the world. We all feel as if our experiencing of the world around us were quite direct. However, the apparent immediacy of this experience *has* to be more or less illusory because we know that every bit of our information about external things is coming in through our sense organs, or has come in through our sense organs at some time in the past. All of it, to the best of our knowledge, is mediated by receptor activity and is relayed to the brain in the form of Morse code signals, as it were, so that what we experience as the "real world," and locate outside ourselves, cannot possibly be anything better than a *representation* of the external world. (Epistemologists can argue about whether it is even that, but I am willing to take for granted the existence of a physical world that is being represented.) The afferent nerve impulses that link the representation to the reality are extraordinarily dissimilar to either, however formally considered. This is the point of Brunswik's (1952) *lens analogy:* information about the world is represented in a very diffuse way at the receptor level, but is brought to a "focus" within the nervous system where objects and events are presumably represented in a form more "like" the outside world than they are at the receptor surface.

The statement that the world as we know it is a representation is, I think, a truism—there is really no way it can be wrong. It has some fairly interesting corollaries, however, that you may or may not have contemplated. One of them is that I have got to be at least as complex as the world as I know it, and psychologists who are attracted by simplistic theories might do well to ask themselves whether they are taking this fact seriously.

We can say in the first place, then, that knowing necessarily involves representation. Now what does "representation" mean? To say that one system represents another, we must at least mean that certain parts and relations between parts of one system must, via some sort of transformation, *correspond to* certain parts and relations between parts of the other system. This is not necessarily to say, however, that the correspondence has to be part for part and relation for relation. Another important possibility, which I shall discuss presently, is that what is a relation in one system may be a part, or an element, in the other. It just makes for a more complicated transformation if that is true.

Let us consider two representational systems that we know something about, and that exist objectively, outside the head where we can examine them. The first is spoken or written language. Think of the *Encyclopedia Britannica.* An encyclopedia *represents* considerable portions of the world. Note that in the case of language, the English language in which the *Encyclopedia* is written, the relations in the world being described are not preserved as relations on the printed page. Relations in the world are represented almost entirely by *things,* that is, *words,* in the language system. The word *chair* and the word *inside* are both simply elements of language, and if you do not know English, you do not know that one is a noun

[1] This article was the presidential address delivered at the meeting of the Western Psychological Association, Anaheim, California, April 13, 1973. It was prepared with support from the National Institute of Mental Health Grant MH 20-449-02 for Studies on Spatial Representation.

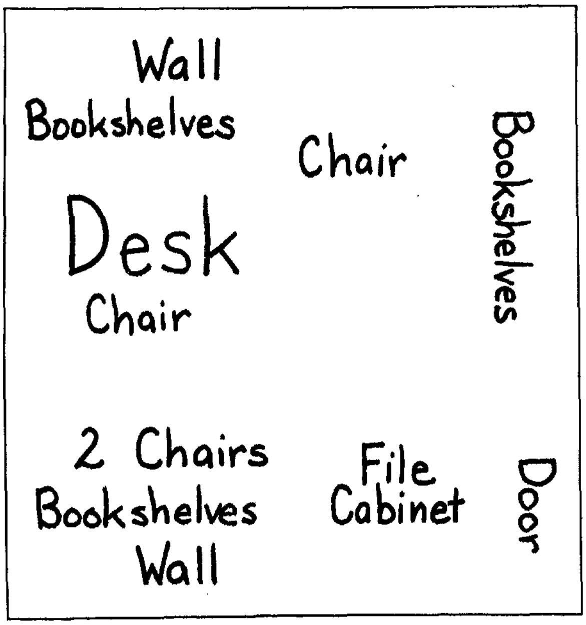

Figure 1. A hybrid (digital–analogue) representation of an office.

and the other a preposition. This substitution of things for relations is one of the most fundamental of all linguistic devices. The English language and, in fact, all other natural languages employ it quite effectively in describing large portions of what we call reality.

Another kind of representational system that we are all familiar with is the *map*. A map is quite different from language in that it is an analogue representation, whereas language is a digital representation. What this means, at least for our present purposes, is that relations among the items being mapped, items like cities, lakes, rivers, and so on, are represented by relations on the map, whereas they are not in language. In fact, in the special case of the map, space is represented by space; within certain well-understood limits, the map is just a minification of the terrain that is being mapped. It is not, of course, a completely literal minification. On the map, towns are usually represented by black circles, though they do not look like black circles from an airplane; on this microlevel the analogy is not maintained.

Digital systems of representation and analogue systems may coexist in many combinations. The *Encyclopedia Britannica contains* pictures and maps. Likewise, a map contains words to tell the names of towns, states, lakes, and rivers. Furthermore, we can have all sorts of hybrid systems, as in Figure 1, in which gross spatial relationships *between* objects are shown in an analogue manner—a maplike manner—but spatial relationships *within* objects are not shown. The objects are merely symbolized by arbitrary words.

Now, I think nobody would particularly object to my saying that both the map and the encyclopedia contain knowledge. Do we really want to say, though, that the *Encyclopedia Britannica knows* the various facts that it contains? I suspect you would be uncomfortable with that proposition, and, just to exacerbate your discomfort a little, let me ask you to imagine a person who knows nothing about the English language, let us say a Chinese, who is nevertheless gifted with such uncommon eidetic imagery that he can look through the encyclopedia and retain a precise visual image of every page, which he can call up at will and, given enough time, even reproduce. Are we willing to say that he *knows* what is in the *Encyclopedia Britannica?* Again, this does not seem right, and the reason it does not is that "knowing" at least connotes some kind of utility, some kind of consequences for behavior. Let us sidestep the discussion of "meaning" that suggests itself at this point and merely ask, very explicitly: Why is representing the world a good thing biologically—that is, for survival? It is not too hard to give a commonsense answer to this question, and I think the commonsense answer is quite right. If it is not obvious that all knowledge is important for survival, it should be at least moderately obvious that *knowing how* is biologically important. It appears to me that *knowing how* is indeed both the beginning and the end of knowledge.

The basis for *knowing how* can be represented in a somewhat oversimplified way as an SRS linkage, like this:

$$S_1, R \rightarrow S_2 .$$

If situation S_1 obtains at a given time, and I do R, then situation S_2 results. If I *know* this, I know how to change situation S_1 into situation S_2. The beginning of knowledge, I think, is to be found in the fact that we live in a lawful world, in which propositions of this SRS type have some continuing validity from one day to the next. If we did not live in a world in which lawful relationships of this form existed, learning would simply not be worthwhile to the organism at all—nor, for that matter, would innate reflexes or tropisms have any value. It is only by doing things that change one situation into a more favorable situation that the organism can possibly affect his chances of survival by his own behavior. Likewise, the end of knowledge, that is, the utility of knowledge, lies in knowing how to change things for the better and to avoid changing them for the worse. This SRS notation is shorthand for something that is a lot more complex, of course. If we wanted to get a little more fancy in one fairly obvious way, we should formulate this in terms of a conditional probability matrix, because S_2 may not always follow S_1 and R; it may merely be a good bet; also, S_2 may happen regardless of R, and so on. Without pursuing this probabilistic aspect, let me just point out that the organism will be better off to the degree that he can slice up the world into chunks such that situations described in terms of those chunks yield more determinate SRS relationships than if the world were partitioned otherwise. In other words, good organization in the representational system is organization that makes for relatively determinate SRS linkages; poor organization is that which makes for more indeterminate linkages.

Although I am proposing that *knowing how* is the beginning and end of knowledge, I am not proposing that it is necessarily the middle. The middle may take quite a different form. In fact, there may be considerable advantage in depersonalizing the R portion of the SRS linkage. To take one such step: The R may be, not what *I* do, but what *one* does, what some generalized person might do in that situation. Thus, if I see another person perform the act with a determinate result, I may imitate him; I may say that if he can do it, I can. There are further steps of depersonalization in which the representation can simply take the form of causal linkages between events. If I understand how the world works, if I effectively represent what Tolman and Brunswik (1935) called "the causal texture of the environment," then the opportunity exists for me to intervene in the causal network and make things come out my way.

I think the utility of knowledge has to be explained in terms such as these. But let me add that we have many items of knowledge, the utility of which we would be hard put to demonstrate individually. Accumulating knowledge, at least at the human level, is a lot like saving pieces of string, for which the classical justification is that "It may come in handy some day." So it is with knowledge.

My argument, then, is not that representation is *reducible* to SRS linkages but merely that it must embody information from which SRS linkages can be derived conveniently when the occasion demands. Phenomenally, the set of representations we are talking about is completely identifiable with "the external world as we know it." But the scientific problem is to characterize this system in more objective terms, looking at it from the outside, that is, *as* a system. How can we *represent* the representational system?

There is one thing that the system is certainly *not:* It is not just a lot of stimulus objects (or surrogates thereof) tied together by simple, uniform associative bonds. Objects in the world as we know it are indeed associated, but they are associated *by particular relations*, and relations have to be differentiated one from another. Relations have quite as much fine structure as objects and classes of objects have. The situation that is described by the sentence "John gave the ball to Mary" cannot be represented by undifferentiated associative linkages between John, Mary, and ball. The particular act of giving, of exchange of possession, is involved here, and it takes a particular direction—the ball passes *from* John *to* Mary. The relations that we wish to represent vary tremendously one from another. For example, the relationship that is described by the word *give* is of a completely different *family* from the relationship that is described by the words *is a*—as in "John is a boy," or "John is a dog," as the case might be. (If I say "John is a dog," that is essentially a statement of how I can transfer knowledge. Whatever I know about dogs, I can apply to John.) If only one kind of relationship between the elements of a representational system is possible—that is, a simple associative bond—the only way such a system can be made to represent the numerous and diverse relations it has to cope with is by transforming these relations into *things*, or elements—just as they are transformed into words in natural language—in which case the associative bond will have a status comparable to that of serial order in language.

Now, should we take one further step and suppose that language *is* the representational system, for man? This position has been favored by a good many behaviorists in the past, who have supposed that anything they could comfortably call knowledge does not come into being until the human level, when language becomes available to subserve that knowledge.

It is not difficult to refute this hypothesis, at least in its most naive form. For example, Bartlett (1932) showed some years ago that memory for stories is not in terms of the words that make up the stories, but in terms of the ideas. Words, in other words, are attached to representational structures that are more fundamental than language. These structures can be given various names. We can call them concepts; we can call them logogens, as Morton (1969) did: In common parlance, they are the *ideas* to which the words are attached. We are all clearly aware that the word *tree* is not the same as a real tree—even Premack's (1970) chimpanzee Sarah knows this difference. But now let me remind you, once again, that what you know as a real tree is not the real tree out there at all, but your representation of a tree, so there has to be some level of representation that is not verbal. Furthermore, I do not believe for a moment that concepts come into being only when language is acquired. It is much more plausible to suppose that subhuman animals—certainly the higher mammals—have concepts and concept structures that are not ridiculously different from ours, and that when these concept structures reach a certain level of evolution, they are ready to have language grafted onto them. I think a comparative psychology of semantics is very much a possibility and that something of the sort is coming into being at this time, particularly with the work of people like Gardner and Gardner (1971) and Premack (1970), who are teaching language to chimpanzees. It seems moderately clear from their work that lower animals do have concepts just sitting there waiting for words to be attached to them. Premack emphasized this point and gave some evidence for it. The Gardners found that their chimpanzee Washoe would say "dog" either to a real dog or to a picture of a dog, or to the sound of a dog barking. Herrnstein and Loveland (1964) got similar results, even with pigeons. They found that pigeons would give a particular response to human figures generally and would even give the same response to parts of human beings, such as hands. So it seems by no means unreasonable to suppose that lower animals, in many cases, may have already evolved languages of their own—languages peculiar to their own nervous systems. Covert languages, or languagelike processes, may be entirely possible in the absence of language responses. Let me suggest, then, that some of the structures that we find in natural languages may be a guide, albeit a feeble one, to the nature of languagelike processes at the logogen or concept level. One can certainly imagine that there are unitary events in the nervous system that correspond to objects, or classes of objects; that there is a "dog" neuron somewhere and a "table" neuron somewhere else. These would be essentially what Konorski (1967) suggested as *gnostic units*. But there may likewise be unitary processes, even in the case of the lower animals, that are the precursors of relational words like *inside, before, give,* and so on. Relations, like objects, may be represented in digital terms as in fact they are in language.

There are many psychological functions, however, that give every indication of entailing *analogue* representation, that is, representation in which relations are represented by relations. To take an example that is by no means the most obvious: People seem to represent numbers in an analogue fashion. Moyer and Landauer (1967) found that if a subject was asked to tell which of two numbers was greater, his reaction time was shorter the greater the difference between the two numbers. This is what you would expect from an analogue representation, but it is exactly the opposite of what you would expect if, for example, the subject were counting from one number to the other. More generally, I think that any psychophysical continuum to which metric considerations apply, that is, any continuum that is potentially scalable, is essentially an analogue system, or subsystem. The very fact that a subject is able to squeeze a hand dynamometer as hard as a light is bright (Stevens, Mack, & Stevens, 1960) is exceedingly difficult to understand at all, except in terms of analogue processes. Some people have tried to conceive psychophysical continua as the axes of a huge multidimensional space in which objects may be represented as points. I think that in general this is the wrong approach. What is more typically the case is that the psychological continuum functions as a medium for the representation of relations—a medium *in which* relations occur (see Attneave & Olson, 1971).

Perhaps the most important case of all in which analogue representation seems to be necessary is that of the system for representing physical space—the tridimensional space around us. If anything seems immediate to us, in our perception of the world, it is this space in which objects are *located* and in which we move about. The three dimensions of physical space are related to one another in quite a unique way. They are *isotropic:* the same metric applies to all three. You can use a yardstick to measure the height of something, or the width, or to measure in any oblique direction you like. Moreover, a yardstick *looks* just about a yard long regardless of its distance away and the angle at which it is viewed. Physical entities are all represented as having determinate locations in this space, and this is true of imagined space as well as of perceived space. The animal *knows where* the water hole is.

Now, it is true that we have linguistic terms for spatial relationships: *above, behind, in front of, inside, outside,* and so on. It nevertheless seems exceedingly likely to me that these words refer to relations in an analogue model of space. This assumption of a tridimensional analogue model of physical space has some important theoretical advantages. For example, the Gestalt psychologists (e.g., Koffka, 1935) and Ernst Mach (1886/1959) before them argued that in perceiving objects we represent them in the simplest way that is consistent with the constraints of the input, the stimulation. This is the Gestalt principle of *Prägnanz,* and there is a great deal to recommend it. Several years ago, Robert Frost and I (Attneave, 1972; Attneave & Frost, 1969) did some experiments, and the results were quite consistent with the *Prägnanz* account of space perception. Thus, we were led to ask: What does one have to assume about the nature of a system that is capable of *finding* a simplest representation? It appeared to us that the minimum assumptions involved some kind of analogue tridimensional model in which objects *could* be represented either more or less simply and in which more complicated representations could be smoothly transformed into simple ones (i.e., by continuous, as opposed to discontinuous, processes) under the guidance of some feedback signal from a primarily digital descriptive system.

An arrangement of this sort is suggested in Figure 2. On the left, we have a tridimensional modeling medium, in which any representation consistent with the stimulus constraints *might* be constructed. This representation is *described* in the center box in Figure 2, and if it changes from one moment to the next, any resulting change in the complexity of the description is then fed back as a hot-cold signal into the tridimensional system, thereby guiding it into a *simplest* representation as

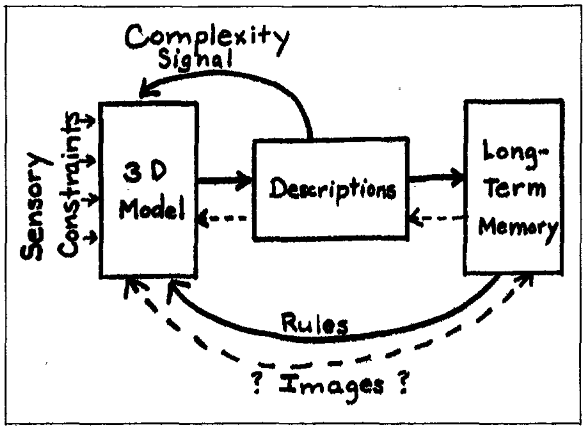

Figure 2. Outline of a system for achieving economical representations.

a stable state. This is essentially a hill-climbing machine, and the analogue medium provides a smooth terrain, so to speak, for the hill climber to operate on.

An analogue stage like this also makes a great deal of sense in terms of the identification or categorizing of objects, because the descriptive machinery is taking its descriptions, or its defining features of objects, not from the flat picture on the retina but from a model of the tridimensional world. It is describing solid objects rather than plane projections of the objects.

Evidence for the analogue modeling of physical space comes from quite diverse sources. A few years ago I became intrigued with an old study by H. H. Corbin (1942) on apparent movement, in which he found that the time interval between two flashing lights necessary to produce apparent movement varies with the distance between the lights (as Korte's third law of apparent movement says it should) but that the distance in question is not the retinal distance, not the visual angle between the lights, so much as their *physical* separation. He dissociated the two by presenting the lights on a slanting board. Gene Block and I (Attneave & Block, 1972) have recently done a series of similar experiments in which we have fairly thoroughly verified this finding. In fact, we found that the time interval necessary for apparent movement varies with the phenomenal or perceived distance between the light stimuli even when this distance is varied by purely illusory means, that is, by placing the lights at different apparent depths in a picture. What this shows, or what we believe it shows, is that apparent movement must be occurring in a tridimensional representational medium in which the isotropy of physical dimensions is fairly well preserved.[2]

However, the most compelling evidence I know for the existence of a tridimensional modeling medium in the head comes from the work of Roger Shepard and his students. Shepard and Metzler (1971), in a study that many of you know, presented the subject with two pictures of tridimensional objects and asked: "Can one of these be rotated into the other? Are they the same except for orientation?" In the case of positive responses, that is, when one could be rotated into the other, the reaction time turned out to be a very precise linear function of the angular difference between the objects. Furthermore, it made no difference whether the rotation was in the frontal plane or in depth: The reaction time function was almost identical for the two types of rotation. This and other studies from Shepard's laboratory, which I do not have time to review, show beyond any reasonable doubt that when one rotates a mental image from one aspect to another, the representation of the object is in fact going through all of the intermediate aspects in a continuous manner. I have no idea how anybody could possibly account for these results without postulating an analogue representational medium.

This brings up the business of imagination and images, and I would like to suggest that this tridimensional modeling medium can be used not merely to represent the current input but also to represent images that are taken from memory, that is, that imaginary scenes can be reconstructed in space and that the organism can then proceed to use this as a work space in which he tries out things and sees what happens. He can engage in vicarious manipulation; he can engage in vicarious locomotion. He can try out the results of particular forms of behavior before he commits himself to them in practice. This highly developed facility for handling spatial information may be used in various ways. Consider, for example, the popularity of graphs, in psychology and other sciences, in which nonspatial continua are mapped onto spatial coordinates in order to make functional relationships more easily apprehended.

Now, if images in this space can be generated from memory, the question immediately arises, In what form do they exist in memory? We might suppose, of course, that they are *stored* in an imagelike, or picturelike, form. Alternatively, we could suppose that they are stored as languagelike descriptions, which are reversible in the sense that images can be reconstructed from them. Pribram (1971) has a third alternative, or perhaps an in-between alternative: He believes that images are stored in the form of neural holograms of some sort. This is a very real possibility, but I am more inclined to favor the notion of reversible descriptions. In a remarkable article published a few years ago, the Dutch psychologist Leeuwenberg (1971) developed a powerful descriptive language in terms of which highly complex tridimensional visual forms can be represented. These descriptions are about as economical as they can be; that is, they take into account virtually all of the internal regularities of the forms being described to eliminate redundancy. Moreover, they are reversible: They can be used as instructions for rebuilding the form. Leeuwenberg believes, as I do, in the necessity of a tridimensional analogue representational medium from which such descriptions can be taken and into which images can be "projected" from descriptions.

In any case, if we suppose the existence of languagelike representational structures, the elements of which have a wordlike status, whether in memory or in consciousness, we need to consider how and where these structures get their *meaning*. How does the nervous system understand its own language? One place to look for the meaning would certainly seem to be in the imagery that the description can generate. In other words, what has the function of a *logogen*, looked at one way, may also have the function of an *iconogen*, looked at the other. But there is another aspect of meaning that I think is even more essential. I can imagine going for a walk and encountering a dog, or a bird, or a wildcat in the woods, but how the scenario progresses beyond that point is highly dependent on *which one* of the three I imagine meeting. The rules of the game are by no means the same for dogs, wildcats, and birds; the SRS linkages that are plausible are extremely different

[2] I have just learned that the Japanese psychologist Jiei Ogasawara (1936) did a study of this type that antedated Corbin's by several years. His results were quite consistent with those mentioned in the text: The apparent movement threshold was highly dependent on apparent separation when viewing distance was varied with retinal separation held constant. I am grateful to Tadasu Oyama for bringing Ogasawara's article to my attention and for providing me with a copy and an English summary of it.

in the three cases. The real point or utility of identifying anything is in order to access the rules by which it interacts with other things and, more particularly, with us. We may be able to suppose that the rules of geometry, and perhaps even some very simple principles of causation, are embedded in the spatial representational medium, but a great many other rules are certainly dependent on *what* it is that is represented in that medium.

However, these two aspects of meaning that I have just suggested—that is, connection with imagery on the one hand, and access to rules on the other—are not as distinctly different as they may seem at first. I am thinking particularly of deGroot's [3] work with chess masters who are able to look at some chess position from a real game for five seconds and afterwards reconstruct it perfectly. It becomes quite evident from their reports that they are not just assigning particular pieces to particular squares on the board, but rather remembering *functional relationships* between pieces—what piece is threatening another, what piece is guarding another piece against a threat, and so on. In other words, the rules of the game turn out to be quite essential to the way the position is remembered and reconstructed.

It should be obvious at this point that I do not have a beautiful, elegant, comprehensive answer to the question that I posed in the title of this presentation. But let me leave you with this suggestion. Let me ask you to imagine—use your spatial representational system to imagine—a 2×2 table. On one axis put "digital, languagelike processes" versus "analogue, maplike processes," and on the other axis put "consciousness" versus "memory." Now, if we can ever figure out what belongs in each of these four cells, and how each of the four is related to the other three, then I think we shall possibly know what knowing is.

---

[3] deGroot, A. D. "Perception and Memory in Chess: An Experimental Study of the Heuristics of the Professional Eye." In preparation, 1974.

REFERENCES

Attneave, F. Representation of physical space. In A. W. Melton & E. J. Martin (Eds.), *Coding processes in human memory.* Washington, D.C.: V. H. Winston, 1972.

Attneave, F., & Block, G. Apparent movement in tridimensional space. *Perception & Psychophysics,* 1973, **13**(2), 301–307.

Attneave, F., & Frost, R. The determination of perceived tridimensional orientation by minimum criteria. *Perception & Psychophysics,* 1969, **6**, 391–396.

Attneave, F., & Olson, R. K. Pitch as a medium: A new approach to psychophysical scaling. *American Journal of Psychology,* 1971, **84**, 147–166.

Bartlett, F. C. *Remembering: A study in experimental and social psychology.* London: Cambridge University Press, 1932.

Brunswik, E. The conceptual framework of psychology. *International Encyclopedia of Unified Science,* 1952, **1**(10).

Corbin, H. H. The perception of grouping and apparent movement in visual depth. *Archives of Psychology,* 1942, Pt. 1, No. 273.

Gardner, B. T., & Gardner, R. A. Two-way communication with an infant chimpanzee. In A. Schrier & F. Stollnitz (Eds.), *Behavior of nonhuman primates.* Vol. 4. New York: Academic Press, 1971.

Herrnstein, R. J., & Loveland, D. H. Complex visual concept in the pigeon. *Science,* 1964, **146**, 549–551.

Koffka, K. *Principles of Gestalt psychology.* New York: Harcourt, Brace, 1935.

Konorski, J. *Integrative activity of the brain: An interdisciplinary approach.* Chicago: University of Chicago Press, 1967.

Leeuwenberg, E. L. J. A perceptual coding language for visual and auditory patterns. *American Journal of Psychology,* 1971, **84**, 307–350.

Mach, E. *The analysis of sensations.* New York: Dover, 1959. (First German edition published in 1886.)

Morton, J. Interaction of information in word recognition. *Psychological Review,* 1969, **76**, 165–178.

Moyer, R. S., & Landauer, T. K. Time required for judgments of numerical inequality. *Nature,* 1967, **215**, 1519–1520.

Ogasawara, J. Effect of apparent separation on apparent movement. *Japanese Journal of Psychology,* 1936, **11**, 109–122.

Premack, D. A functional analysis of language. *Journal of the Experimental Analysis of Behavior,* 1970, **14**, 1–19.

Pribram, K. H. *Languages of the brain.* Englewood Cliffs, N.J.: Prentice-Hall, 1971.

Shepard, R. N., & Metzler, J. Mental rotation of three-dimensional objects. *Science,* 1971, **171**, 701–703.

Stevens, J. C., Mack, J. D., & Stevens, S. S. Growth of sensation on seven continua as measured by force of handgrip. *Journal of Experimental Psychology,* 1960, **59**, 60–67.

Tolman, E. C., & Brunswik, E. The organism and the causal texture of the environment. *Psychological Review,* 1935, **42**, 43–77.