

## *On the Forms of Mental Representation*

The human brain encodes, modifies, and stores information that is received through its various sense organs, transforms that information by the processes that are called “thinking,” and produces motor and verbal outputs of various kinds based on the stored information. So much is noncontroversial—only the most radical of radical behaviorists question it. What is highly controversial is *how* information is stored in the brain—in the usual terminology, how it is “represented”—or even how we can describe representations, and what we mean when we say that information is represented in one way rather than another.

There is not a single problem of representation; there are several, each referring to a different level of analysis. There is, of course, the basic physiological question of the nature of the “engram”; but I shall not be concerned with that question here. In talking about a computer memory, we do not need to concern ourselves with magnetic cores and integrated circuits; we can talk about memory structures at the information-processing level, in terms of symbol structures and the operators that act upon them. In different computers, these symbol structures and operators may be realized by quite different physical devices. In the same way, our interest here will lie in the symbol structures in the human mind and in the operators that act on them to transform them. A theory of mental representation at this symbolic level may be compatible with a variety of different biological realizations. That is fortunate, for very little indeed is known about the biological foundations of human memory.

*Note:* This research was supported by Research Grant MH-07722 from the National Institute of Mental Health.

Since we are concerned with representation at the symbolic or information-processing level, it will matter little whether the memory we are talking about resides in a human head or in a computer. Hence we may use what we know about computer memories to clarify our notions about human memory. This should be helpful since it is very much easier to find out in detail what is going on inside a computer than to find out what is going on inside a human brain. I shall base the discussion on evidence drawn from both artificial intelligence and experimental psychology.

### Equivalence of Representations

It is impossible to find an entirely neutral language in which to describe representations of information, for a language is itself a form of representation. We can avoid this difficulty, at least in part, by not attempting to describe representations directly, but by talking instead in terms of the equivalence of representations. It is essential that we define carefully what we mean by “equivalence,” and that we distinguish among the various kinds of equivalences that can hold between representations. Specifically, I shall talk about *informational* equivalence and *computational* equivalence.

*Informational Equivalence.* Two representations are informationally equivalent if the transformation from one to the other entails no loss of information, i.e., if each can be constructed from the other.

Thus, in an appropriate information-processing system, the statements “Distance equals average velocity times time” and “ $S = w * T$ ” are informationally equivalent. There is informational equivalence, also, between appropriately axiomatized formulations of Euclidean geometry and analytic geometry, respectively.

On the other hand, consider a system that receives information via a two-dimensional spatially extended “retina” but stores a particular three-dimensional interpretation of this information. What is presented to the system is, say, the familiar Necker Cube of Figure 1, and what is stored is the information about the vertices, edges, and faces of this cube, with the face ABCD labeled “front” and the face EFGH labeled “back.”

Now the important feature of the Necker Cube is that this is not the only possible three-dimensional interpretation of the presented




Figure 1. The Necker Cube

figure. It is equally consistent with that figure to label the face ABCD the “back” and the face EFGH the “front.” It is precisely because of this ambiguity that the Necker Cube is reversible—i.e., that a normal subject can alternate at a rate of several times per second between the two internal three-dimensional representations. Hence we must conclude that the two-dimensional stimulus and the internal representation are not informationally equivalent: the former can always be obtained univocally from the latter, but not vice versa.<sup>1</sup> These examples should suffice to illustrate the meaning of informational equivalence. We turn now to computational equivalence.

*Computational Equivalence.* Two representations are computationally equivalent if the same information can be extracted from each (the same inferences drawn) with about the same amount of computation.

For certain purposes, we may wish to weaken this definition slightly, replacing it by: Two representations are computationally equivalent *in the extended sense* if the same information can be extracted from each by the same amounts of computation, up to a factor of proportionality.

Thus, if we had two computers (like the IBM 709 and IBM 7090) that were essentially isomorphic at the information processing level, but one of which, for hardware reasons, had a much faster operation cycle than the other, all operations would take on the slower machine a fixed multiple of the time that they would take on the faster machine. If the processing of some particular representations on the two machines differed only by this fixed factor, we would say that

the representations were computationally equivalent in the extended sense. This weaker definition of equivalence is crucial in comparing computer simulations with human performance, since the representations in this case, having different hardware realizations, will have different processing times associated with the corresponding primitive operations.

Consider the following information:  $ABCD$  is a rectangle with  $AB$  twice as long as  $BC$ . It is bisected into two squares by the line  $EF$ , where  $E$  is the midpoint of  $AB$ , and  $F$  the midpoint of  $CD$ .  $ABCD$  is also bisected into two triangles by the line  $AC$ . Now we ask the following question of a system that has been given this information: Do  $EF$  and  $AC$  have a point of intersection inside the rectangle?

The information about the rectangle and the processes for answering questions about it might be represented in a variety of ways. (1) They might be represented as a set of propositions (more or less isomorphic with the verbal statement), together with an appropriate axiomatization of plane geometry. Then the question could be answered by proving the theorem that  $EF$  must intersect  $AC$ . (2) Or they could be represented as a set of coordinate pairs for the points  $A$  through  $F$ , together with an appropriate set of rules for constructing the equations of the lines joining these points and for finding points of intersection of these lines. Then the question could be answered by solving a pair of simultaneous linear equations. (3) Or they could be represented by a drawing on paper (see Figure 2), together with some visual scanning processes capable of noticing the intersection at  $G$ .




Figure 2. "Computation" with a visual image (by drawing the lines  $EF$  and  $AC$ , infer the existence of  $G$ ).

Although these three representations are informationally equivalent, they are not at all computationally equivalent. For human beings, at least, answering the question with the use of the first representation would generally be the most difficult, and answering it with the third representation, the least difficult. Most human beings, upon being given this problem verbally, would promptly translate it into an internal (“visual” or “pictorial”) representation that is more or less computationally equivalent with the third representation. But I am getting ahead of my story.

### Representation in Discrete Associative Memories

There is no close relation between a system’s “hardware” and the representations it may employ. The observation has already been made that computers that are quite different physically may use representations that are informationally, and even computationally, equivalent. As a matter of fact, virtually all computers of the so-called von Neumann type, that is to say, almost all general-purpose computers that have been built in the past 25 years, are very close to being computationally equivalent in the extended sense, so far as their basic machine-language representations are concerned.

The converse proposition holds also. Not only may the same representation be realized with different physical systems, but different representations may be realized with a single physical system—by superimposing the higher level representations upon the underlying physical scheme. There is, for example, a widespread belief today, based on a fair amount of evidence, that the left and right hemispheres of the human brain use different representations. In some versions of this theory, the left hemisphere employs “verbal” and “symbolic” representations, the right hemisphere, “pictorial” or “holistic” representations. Yet there is no known basic difference in the tissues of the two hemispheres, and indeed, if the left hemisphere of a child is damaged, the right hemisphere is often capable of taking over the usual functions of the other with little or no evidence of deficit.

Returning again from the obscurities of the human brain to computers, let us consider how specific representations are defined for a system, and how higher level representations are superimposed upon a hardware representation. Defining a representation means

(1) specifying one or more data types, and (2) specifying the primitive (i.e., “fast”) operations that can be performed on information in those data types.

A *data type* is some particular way of organizing information in memory. For example, among the data types that are commonly used in computing are lists and arrays. A list is an ordered set of symbols or symbol structures; an array is a set of symbols or structures to each of which has been assigned  $n$  coordinates, so that (in the case of a two-dimensional array) we can refer, say, to the symbol with coordinates  $(i, j)$ .

The declaration that information will be represented in lists or in arrays does not say anything about the physical location of the information in memory. Two adjacent items on a list need not occupy adjacent memory locations. Rather, the declaration of a data type means that certain *primitive operations* are available for accessing, storing, and modifying data. A list representation requires a primitive operation of *find.next*, which, given a particular symbol as its argument, will access the succeeding symbol in the list to which the first belongs. In a list representation, other primitive operations are available for inserting a symbol in a list (i.e., after symbol  $X$  has been inserted between  $Y$  and  $Z$ , execution of *find.next* ( $Y$ ) will yield  $X$ ); for deleting a symbol from a list; and so on.

We may indicate, for illustration, one way in which a list representation could be realized with standard computer hardware. At a hardware level, information is held, say, in the form of *words*, which are individually *addressable*. That is, with each word is associated a number, its address; so that a primitive operation, *find* ( $X$ ), will access the symbol whose address is the number  $X$ . To superimpose a list representation upon this hardware addressing scheme, a new level of primitive operations is defined (higher-level language), each element of which is a small algorithm of the original hardware-level primitive operations (machine language). The convention could be adopted, for example, that symbols are to be stored in the odd-numbered addresses of the machine-level locations, and in each succeeding even-numbered address is to be stored the address of the *next* symbol in the list. Thus, if  $X$  were stored in hardware location 17, the number “25” in location 18, and the symbol  $Y$  in location 25, then the list operation *find.next* (17) would access location 25,

i.e., the symbol  $Y$ . For further discussion of representations, the reader may consult Newell and Simon (1972, pp. 21-38), or for a more technical account, Galler and Perlis (1970) or Knuth (1968).

It is commonly thought that human memory is *associative*. In an associative memory, the basic *find* operation takes the form *find* ( $X, Y$ ); e.g., “find (the name of) the father of John.” With each symbol stored in such a memory is associated a set of labeled links. Given a symbol and one of the labels, we can find a second symbol that is associated with the first by that particular link.<sup>2</sup> It is a matter of considerable disagreement, however, whether the mental visual imagery of which humans appear capable can be accounted for in terms of such an associative structure. That question will provide a major theme for the next section of this paper.

### What Is Visual Imagery?

Modern computers have freed our thinking about representations. First, they have taught us how representations can be defined operationally by specifying data types together with their associated primitive operations. Second, they have taught us that representations can be superimposed on one another, in the manner that was described in the last section, so that hardware organization need not constitute a severe constraint on the representations that are available. Third, they have taught us that a multiplicity of representations may exist side by side, realized in the same hardware.

In the enjoyment of this freedom, we can now discuss representation in human memory without undue concern for our ignorance of hardware realization. We can ask, as an empirical question, how many representations are employed by the human mind, what kinds of data types they are built upon, and what kinds of primitive operations. Empirical evidence can be gathered by presenting human subjects with various kinds of information and instructing them to perform a variety of tasks that draw upon that information. The possibility or impossibility, ease or difficulty, of performing the tasks will inform us about the characteristics of the memory representation the subjects are using, and about the informational and computational equivalence of different representations.

In this section I would like to sketch out the evidence for the proposition that most of the phenomena associated with human

visual imagery can be realized in a representation that is superimposed on a basic associative system (see also Simon, 1972). Subsequently I shall expand the argument to claim that all of the known human memory representations—"pictorial," "propositional," or whatever—are subspecies of associative representation; but that inside the human head everything is lists and node-like structures of symbols.<sup>3</sup> This does not mean, of course, that all of these representations are computationally equivalent. It is precisely because they are not computationally equivalent that it is convenient for the system to have more than one of them.

In making my case, I shall refer only occasionally to formal experiments published in the psychological literature. Instead I shall rely mainly on informal experiments that you can perform, using yourself as subject, as you read about them. In fact, I have already asked you to perform two such experiments: one having to do with the Necker Cube (Figure 1), the other with a rectangle that was several times bisected by lines (Figure 2). Let me revert briefly to those examples, and then introduce some additional ones.

*The Necker Cube.* The Necker Cube can illustrate for us how a system with associational capabilities could picture a three-dimensional structure. The full scheme is fairly elaborate, so I can only sketch it here.<sup>4</sup> Vertices, edges, and faces will be represented in memory by nodes. These nodes will be connected by labeled links. For example, the labels "left.edge," "down.edge," and "back.edge" will associate the edges BA, BC, and BF, respectively, with the node B; the face ABCD will have the attribute "front," and EFGH, the attribute "back"; and so on.

The ambiguity of the Necker Cube (i.e., the fact that face ABCD could have the attribute "back," and the face EFGH the attribute "front," instead of vice versa) can arise in constructing the internal representation from the two-dimensional stimulus of Figure 1 as follows. Suppose a scanner were to begin by fixing on point B. The horizontal line segment BA would be interpreted as the right.edge of B. Similarly, the vertical segment BC would be interpreted as the down.edge of B. Since the system is programmed to interpret all angles as right angles unless a contradiction arises, BF must be labeled the forward.edge or the back.edge of B. Convention determines which label is chosen. Suppose it is the latter. Then, for consistency,



the edge DH will later have to be interpreted as the back edge of D and the forward edge of H, thus placing H behind B. If, on the other hand, the processing had begun at point H, the same convention would have placed H and F on the front face, and D and B on the back. If this hypothetical account has any relation to the ways in which people process and store the Necker Cube, then it should be possible to induce a figure reversal by shifting attention from point B to point H of the figure. The reader can verify for himself that he can indeed do this.

*The Dissected Rectangle.* The dissected rectangle of Figure 2 provides an argument against simple node-link processes as the *sole* means for processing and storing information about geometric figures. How is the point of intersection G created in the mental representation of the verbally described rectangle when AC is added to the description of the figure? If the annexing process had some implicit knowledge of plane topology, it would know that A could not be joined to C without penetrating the boundaries AEFD and EBCF. However, that information alone is not enough to infer an intersection of EF with AC; metrical information about the lengths of the segments and the angles, and guaranteeing the straightness of the lines, is also needed. If only topological information were processed, segments like DF and FE could be replaced by curved lines, and the intersection of AC with EF could not be inferred.

Considerations of this kind argue for the availability of a “mind’s eye” that has some of the capabilities of an analytic geometry representation, or a representation by means of a fine grid, to encode verbally presented descriptions into visual images that can be stored in memory.<sup>5</sup> On the other hand, if it seems difficult to account for the creation of such images in terms of an associational representation alone, such a representation may be fully adequate for storing the visual image once it has been generated (e.g., once the existence of the point G has been inferred).

The amount and kinds of information that can be stored in a visual image can be assessed by elaborations on the rectangle-generating task. Remembering that the rectangle of Figure 2 is twice as wide as high—that AB is twice AD—we now draw a diagonal in the right-hand square connecting B with F. BF will intersect AC at a point we will call H. Finally, we mark the midpoint of BF, calling it J.

Now we ask a subject who has been synthesizing this mental image from the verbal description (without Figure 2 before him) to tell us whether JH is longer or shorter than HF. If the answer is “shorter,” we then ask how much longer HF is than JH.

In informal experiments with this task, I have found that most, but not all, subjects can answer the first question correctly. (JH is shorter than HF.) Fewer than half of the subjects can answer the second question correctly. (HF is twice as long as JH.) This task seems to strain to the capacity point, or just beyond, the ability of most persons to assemble a mental picture of a geometric figure. Hence the task demonstrates that the total amount of information that can be held in such a picture is quite modest, so that it could readily be accommodated in an associational representation.<sup>6</sup>

*A Block Test.* Baylor has constructed a computer simulation, using an associational representation, of tasks involving mental visualization of blocks. A typical task is the following: Imagine a cube that is three inches on a side. Paint one face of the cube red, and two other faces, adjacent to the red face and opposite to each other, blue. Now slice the cube into one-inch cubes. How many of the one-inch cubes have exactly one red and one blue side?

Most people can answer this particular question correctly. (There are six such cubes.) It is also interesting that they can tell you *which* face of the cube they colored red, and which faces blue, although the instructions did not require them to determine that. Hence we can conclude that the representation-building process synthesizes the figure in a particular orientation.

Baylor's simulation of the processes used by subjects in this task is also able to answer the question. The symbol manipulation required, however, to create and modify the representation (e.g., to “slice” the cube) is quite extensive. With problems that are somewhat more difficult than this one, the simulation makes errors that are similar to typical human errors—it is unable to process the information in such a way as to draw all the inferences correctly. Hence Baylor's experiments and simulation add a little more credence to the idea that associational representations may be computationally equivalent, or approximately so, to the representations that people use to store visual images.<sup>7</sup>

*Summary of Evidence on Visual Representation.* The kinds of experiments we have been reporting illustrate how much information, and what kinds of information, people can store in visual images. It is clear from such experiments that visual images permit certain types of inferences to be drawn that would be difficult or impossible to draw by verbal or even mathematical reasoning without those images. On the other hand, these same experiments show that the amount of information that can be held in a visual image is modest. For a scene of any complexity, the image falls far short of a photograph in information content.

Finally, with one important exception (a phenomenon like the generation of the point G in Figure 2), the evidence seems to be compatible with the idea that the underlying representation is wholly composed of list structures in an associational memory. The exception, however, is not an isolated phenomenon, and cannot be ignored. It suggests the availability of somewhat more elaborate processes for operating on images, even if these processes are not needed for storing them, once generated. Baylor's evidence on this point is inconclusive. His simulation shows that accomplishing the block-slicing tasks requires fairly elaborate, but not inconceivably complex, information-processing capabilities if an associational representation is used.<sup>8</sup>

### How Many Representations in the Head?

Much of the information that a person receives comes to him in the form of natural language discourse, oral or written. As we saw in the last section, if the natural language sentences describe concrete objects, it is possible that they may be encoded as visual images, not computationally equivalent to linguistic strings. But much information derived from natural language is abstract and not readily encoded in pictures. How is such information stored, and in particular, how closely does the stored information resemble the linguistic strings that are input?

At least two experimental techniques are available for investigating the processing of memory contents that were derived originally from natural language inputs. Time measurements can be made of the speed of response (response latency) for different kinds of tasks

using the stored information; or subjects may be tested on their ability to remember which of a set of stimuli they have seen previously.

*A Chronometric Study.* An example of the first paradigm is the well-known experiment of Chase and Clark (1972) seeking to test how subjects determine whether a given sentence, such as, "The star is above the cross," is or is not true of a simple picture shown to them simultaneously with the sentence. To answer such questions, the subject must (1) have some way of translating the language of the sentence into the representation used for pictures, (2) perform the inverse translation from picture to language, or (3) translate both stimuli into some common representational form in which they can be compared. While the experiments do not choose unequivocally among these three alternatives, they tend to support the third.<sup>9</sup>

*A Recognition Study.* The second paradigm, first applied by Bransford and Franks (1971), has been used to show that the internal representation of information obtained through natural language inputs is not informationally equivalent to those inputs, having lost a substantial amount of syntactic information during the encoding and storage processes. Using this paradigm, Rosenberg and Simon (1977) gave subjects a sequence of stimulus sentences, one at a time. Subsequently they presented a second set of sentences, asking the subject each time whether he had seen precisely that sentence in the original set. When half the original sentences were in English and half in French (but on the same general subject matter in both languages), bilingual subjects frequently stated that they had previously seen sentences that they had in fact seen only in translation in the other language. The same experiment was performed in another variant in which half of the original stimuli were English sentences and half were simple drawings pertaining to the same subject matter. Again, subjects were frequently unable to distinguish between having seen a particular sentence and having seen a picture whose content was approximately equivalent to that sentence.

These experiments might be taken as evidence for the view that there is a *single* internal representation—a semantic or conceptual representation—into which all inputs, whatever their sensory modality or external coding, are translated. However, we must avoid overinterpreting the evidence. First, subjects did not make these errors in every case; sometimes they remembered syntactic details

that enabled them to distinguish sentence inputs from pictures or from sentences in another language. Moreover, these experiments used very simple sentences, and correspondingly, very simple pictures whose content was easily describable in language.

If the experiments do not demonstrate conclusively that there is a single form of internal representation, they at least argue for relatively simple and straightforward translatability from one internal representation to another. In this way, they lend some support to our hypothesis that underlying the modality-specific representations employed by the human processor there is a common associational memory in terms of whose primitives all the higher-level and more specific representations are coded.

*Dual Representations in Simulations.* Several problem-solving systems show us how two representations may be better than one in carrying out heuristic search. Although these systems must be classified primarily as experiments in artificial intelligence, the methods they employ are probably not far from the methods used by people in these same task environments.

An early example of an automatic theorem-proving system was Gelernter's geometry theorem-proving machine (1963). As its name indicates, this system was capable of discovering proofs for theorems in Euclidean geometry. It employed a dual representation of each problem, one symbolic and syntactical, the other in the form of a diagram of the geometric figure. Before the system attempted to prove syntactically that corresponding angles, say, or corresponding sides of a pair of triangles were equal, it first tested for approximate equality on the diagram. The space of the diagram therefore served as a planning space that prescreened proof attempts and saved effort in fruitless proof attempts. People probably use diagrams in geometry for the same purpose.

Recently Novak (1976) built a computer program (ISAAC) for solving physics (statics) problems presented in natural language. This program, after some initial syntactic parsing of the input sentences, generates a diagram-like representation of the problem situation, then uses the diagram in constructing the algebraic equations that translate the problem statements. The diagram contains sufficient information so that it can be used to generate pictures of the situation on a cathode-ray tube. Thus Novak's program uses a variety

of representations: the input language strings, the tree-like parsed sentences, the diagrams, the equations, and the output drawings on the CRT.

Both the Geometry Theorem Prover and ISAAC superimpose all of their representations on an associational, list-processing memory, although both have some capability for performing analytic geometry calculations on their pictorial images. Both demonstrate that “pictorial” representations can be realized on the basis of an underlying associational representation.

### Conclusion

In this paper I have tried to show that the question, “How are ideas represented in the mind?”, can be given a perfectly operational meaning and can be investigated experimentally. In order to study representations, we must first define what we mean by two representations being equivalent. We have seen that there are at least two distinct notions of equivalence, informational equivalence and computational equivalence, and that both are important for an understanding of representation.

The empirical evidence available today leads me, cautiously and tentatively, to believe that the brain operates basically as a system of labeled associations. This system may require supplementation from a specifically visual system (“the mind’s eye) to account for some of the mind’s capabilities for drawing inferences about the intersections and metrical properties of lines in geometrical figures.

Other specialized forms of representation in the brain may well be achieved by encoding higher level primitives at a level above the basic system of labeled associations. If one had to conjecture how many such specialized representations there might typically be, a plausible guess would be three: a verbal representation, perhaps not unlike the “deep structures” postulated by the transformational linguists; a visual representation, capable of holding, if not always generating, the information about spatial figures; and a conceptual representation, more “abstract” than the other two and particularly essential for handling abstract meanings.

But whether this is a correct description of the human representational system is not important. What is important is that we now have a variety of means for exploring representational issues empirically,

so that these questions can no longer be settled from the comfort of an armchair. Identifying the representations used by the brain is an operational, researchable task.

### Notes

1. Caution: in speaking here of the internal representation as "three-dimensional," I do not, of course, mean to imply that it occupies a three-dimensional region of the brain. I simply mean that it is described in terms of the attributes of a cube.

2. The need for a labeled or "colored" link to model human associative memory, rather than a simple "next" association, was first noted by the Würzburg psychologists early in this century. This requirement follows from a simple experimental datum. If, in an association experiment, a subject is given the stimulus "white," he is likely to respond "color" or "black," respectively, depending on whether he has previously been given the set to respond with superordinates or with opposites.

3. Presently I shall qualify my "everything" to "almost everything." Certain operations on visual images seem to call for processes that go beyond the usual basic primitives of the node-link scheme.

4. For a discussion of how a system like the one described here can account for the reversibility of the Necker Cube and for "impossible figures" like the drawings of Escher, see Simon (1967). Computer systems for "scene analysis," i.e., for three-dimensional interpretation of two-dimensional drawings, have been studied extensively by Guzman and Waltz, among others. See Waltz (1975).

5. Nothing mysterious is meant by the "mind's eye" metaphor; the same apparatus that is used to process the retinal image of an external stimulus might, for example, be used to process an image that had previously been stored in memory. All that is required is a set of processes to perform the inverse operations to those performed in encoding the retinal image in node-link form. See Novak (1976) for an artificial intelligence program that has precisely this capability of displaying the information from a node-link representation of a physics problem as a picture on a cathode ray tube. A similar view of imagery will be found in Chase and Clark (1972), pp. 224-230.

6. The reader may want to try a second task of the same nature. Imagine a square with sides labeled clockwise from the upper left corner, ABCD. Draw a circle with radius equal to AD and with D as center. Draw a second circle, with radius equal to one-half AB and with B as center. Do the two circles intersect?

7. See Baylor (1971). Additional evidence supportive of this hypothesis is provided by the works of Moran (1973) and Farley (1974).

8. Some well-known experiments by Shepard on the "mental rotation" of objects have been widely interpreted as denying the possibility of a discrete representation of mental images. (See Shepard & Metzler, 1971.) However, this interpretation has been challenged by the research of Just and Carpenter (1976), who provide a specific discrete processing system to account for the phenomena described by Shepard.

9. This experiment is just one of many directed at this issue, most of which support the conclusion reached here. Another good example is that of Potter and Faulconer (1975).

### References

- Baylor, G. W. Program and protocol analysis on a mental imagery task. *Proceedings of the Second International Joint Conference on Artificial Intelligence*, 1971.

- Bransford, J. D. & Franks, J. J. The abstraction of linguistic ideas. *Cognitive Psychology*, 1971, 2, 331-350.
- Chase, W. G. & Clark, H. H. Mental operations in the comparison of sentences and pictures. In L. Gregg (Ed.), *Cognition in learning and memory*. New York: John Wiley & Sons, 1972.
- Farley, A. M. *A visual imagery and perception system. The results of a protocol analysis*, Vol. 1 & 2. Pittsburgh, Pa.: Department of Computer Science, Carnegie-Mellon University, 1974. (NTIS No. A-004 090 for Vol. 1 and NTIS No. A-004 091 for Vol. 2)
- Galler, B. A., & Perlis, A. J. *A view of programming languages*. Reading, Mass.: Addison-Wesley, 1970.
- Gelernter, H. Realization of a geometry-theorem proving machine. In E. Feigenbaum & J. Feldman (Eds.), *Computers and thought*. New York: McGraw-Hill, 1963.
- Just, M. A. & Carpenter, P. A. Eye fixation and cognitive processes. *Cognitive Psychology*, 1976, 8, 441-480.
- Knuth, D. E. *The art of computer programming*, Vol. I, Chap. 2. Reading, Mass.: Addison-Wesley, 1968.
- Moran, T. P. The symbolic nature of visual imagery. *Proceedings of The Third International Joint Conference on Artificial Intelligence*, 1973.
- Newell, A., & Simon, H. A. *Human problem solving*. Englewood Cliffs: Prentice-Hall, 1972.
- Novak, G. S., Jr. Computer understanding of physics problems stated in natural language. Technical Report NL-30, Department of Computer Sciences, the University of Texas, Austin, Texas, 1976.
- Rosenberg, S. & Simon, H. A. Modeling semantic memory: Effects of presenting semantic information in different modalities. *Cognitive Psychology*, 1977, 9, 292-325.
- Shepard, R. N. & Metzler, J. Mental rotation of three-dimensional objects. *Science*, 1971, 171, 701-703.
- Simon, H. A. An information processing explanation of some perceptual phenomena. *British Journal of Psychology*, 1967, 58, 1-12.
- Simon, H. A. What is visual imagery? In L. W. Gregg (Ed.), *Cognition in learning and memory*. New York: John Wiley & Sons, 1972.
- Waltz, D. Machine vision, understanding line drawings of scenes with shadows. In P. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill, 1975.