

THE NATURE OF PERCEPTUAL REPRESENTATION:
AN EXAMINATION OF THE
ANALOG/PROPOSITIONAL CONTROVERSY

Stephen E. Palmer
Department of Psychology
University of California, Berkeley

In recent years a number of theorists have proposed that perceptual information can be represented in terms of propositional (or relational) data structures (e.g., Baylor, 1971; Palmer, 1975; Winston, 1973). This development has provoked considerable controversy among psychologists and computer scientists over the type of representation most suitable for perceptual information, especially in the case of perceptual imagery (e.g., Bower, 1972; Kosslyn & Pommerantz, in preparation; Pylyshyn, 1973). The arguments typically pit some form of "analog" representation against some form of "propositional" representation. The tack often taken is to contrast a reasonable version of the favored type against an unreasonable version of the other type. This is particularly easy to do since there is considerable latitude within the categories of "analog" and "propositional" representations. While the arguments between the opposing schools of thought have been provocative and useful as an initial enterprise, their continuation in the present vein seems unlikely to produce eventual agreement on the nature of perceptual representation.

It is clearly in the spirit of this conference to analyze the problem in terms of the underlying issues to reach consensus at more basic levels. In this paper I discuss briefly a number of issues relevant to representing perceptual knowledge: is perceptual representation uninterpreted or interpreted, complete or partial, implicit or explicit, holistic or atomic, quantitative or qualitative, absolute or relative? I do not pretend that the problems considered here are exhaustive. (See Bobrow (1975) for some other issues to be resolved.) I have chosen to discuss the problems I see as being most central to the current analog/propositional controversy. For each topic I make a few theoretical, empirical, and methodological observations that seem to merit consideration. The conclusion I reach on many issues is that some kind of synthesis or compromise is required. Throughout the discussion I present my current approach to perceptual representation (see also Palmer, 1975) which illustrates how synthesis and compromise can be achieved within a single formal system.

Structural and Parametric Information

Before turning to the issues, I want to make a distinction between two fundamentally different types of information: structural (or organizational) and parametric (or dimensional) information. Structural information refers to the organization of perceptual elements into groups.

Figure-ground and part-whole relationships are paradigm examples of structural information. Parametric information refers to the values of the stimulus along various perceivable dimensions. Color, size, shape, position, orientation and so forth are examples of parametric information.

Perhaps the most basic distinction between these two types of information is that parameters are properties of both the perceptual representation and the stimulus itself, while structure is a property of the representation alone. If the parameters of two perceptual representations are different, they necessarily correspond to two physically different stimuli. If the structure of two perceptual representations are different, they might correspond to the very same stimulus. Figure 1 illustrates this point by showing a case in which structure and parameters are combined orthogonally.

I do not want to give the impression that structural and parametric information are independent. Clearly, the parameters of a stimulus affect its perceived structure. This fact is manifest in the work of Gestalt psychologists and codified in their laws of organization. Context and world knowledge can also affect the assignment of structural organization -- e.g., the figure 13 will be structured as a single unit in the context of 130Y and as two units in the context of 1357.

It is also true that perceived structure can affect the representation of parameters. Someone perceiving the "parallelogram" organization of the form in Figure 1 might perceive it as taller and thinner than someone who perceived the same form with the "triangle" organization. Once again, context and world knowledge can affect the parameters of a perceptual representation -- e.g., the height of two people in an Ames room.

It is important to make clear the distinction between structural and parametric information because failure to do so can lead to mistaken conclusions about the nature of perceptual representation. A case in point is the frequency with which the results of mental rotation experiments (e.g., Cooper, 1975; Cooper & Shepard, 1973; Shepard & Metzler, 1971) are cited as evidence that perceptual representations and processes are "analog" and continuous in all respects. When the process of rotation is considered in light of the structural/parametric distinction, it is seen that it deals only with changes in parameters -- namely, the orientation of the figure in space. The structure of the figure presumably remains constant, and thus does not contribute to the measured reaction times.

To demonstrate the operation of structural variables in perceptual and imaginal processes, I have developed a "mental synthesis" task. In mental synthesis, subjects are required to imagine synthesizing two spatially separated parts

moving the right part onto the left part (see Figure 2), such that they know what the resulting figure looks like. The data of interest are the times required to synthesize two "good" parts of the to-be-synthesized figure (parts A1 and A2 in Figure 2) versus two "bad" parts (B1 and B2). The results are unequivocal. Synthesizing the bad parts took more than twice as long as synthesizing the good parts, even when the to-be-moved parts were virtually identical, as in Figure 2. I do not yet know exactly what is happening in the mental synthesis task, but I submit that it cannot be accounted for by the simple parametric transformation of translational motion. Once the parametric change in position is performed, structural manipulations (e.g., local and global coupling of segments) must be done in order to construct the perceptual representation of the synthesized figure. I suggest that structural and parametric operations are fundamentally different because the types of information upon which they operate are fundamentally different. As a result, structural and parametric information will be considered separately at appropriate points in the discussion that follows.

The Uninterpreted/Interpreted Issue

One fundamental issue which often arises in debates over analog versus propositional representation is that of interpretation. Are the contents of perceptual representations "raw" sensory data or "interpreted" conceptual structures? Before such a question can be answered, we must be clear about what "interpretation" means.

I view the interpretation of sensory data as having several levels, each being characterized by additional inferences which add meaning to the data being processed. At the lowest level there is the raw sensory data -- the stimulation of the retinal mosaic, if you wish. Some minimal interpretation is added when these raw data are organized into areas and low-level perceptual properties are extracted; this blue, rectangular portion of the scene is some sort of unit as distinct from that green oval portion nearby. (I do not wish to imply that the verbal description given is the same as the evolving perceptual representation. The representation of the regions has information about the particular colors and particular shapes.) These low-level interpretations in terms of colors, shapes, sizes, and so forth provide information from which the observer identifies the areas as objects; the blue area is identified as a car and the green area as a tree.

I think at this level we would all agree that the sensory data have been "interpreted," but there can be yet further interpretations. For example, suppose the car is identified as that of the observer's friend, and it is parked in the observer's driveway. The observer will very likely make the further inference that the friend

has come to visit. One might now object that this is no longer a "perceptual" interpretation, but the fact is that it will influence subsequent perceptions. The person standing at the door will more quickly be identified as the friend than if the car had not been noticed.

Thus, there can be many levels of interpretation and for each there can be a representation. Which one is called the perceptual representation is a matter of personal preference and may be responsible for some of the misunderstandings between those who argue about the nature of perceptual representation. My own preference is to consider them all perceptual representations, but at different points along a sensory-cognitive dimension. The representations are increasingly modality-specific at the sensory end and increasingly modality-independent at the cognitive end. For example, the inference that the friend had come to visit might be equally well made by recognizing the particular sound the friend's car made as it came into the driveway, or by recognizing the friend's distinctive knocking pattern on the front door.

It should be clear that I am advocating an integrated approach to sensation, perception, and cognition (see also Norman & Bobrow, 1975) which is somewhat at variance with modularized, box-theory approaches. Within this framework, the nature of perceptual representation is different at different interpretive levels. At the sensory end, I view the representation as being very much analog, even in the strong form. At the cognitive end, I view the representation as being very much propositional, also in the strong form. But it is about the middle levels that most of the arguments seem to take place, and here neither pure analog or pure propositional representation seem to make much sense. There must be a transition between the two which is likely to be some sort of hybrid. In the discussion that follows, it should be assumed that I am considering a level of perceptual representation in which the stimulus has already been operated upon by low-level perceptual processes, but has not yet been conceptually categorized.

The Complete/Partial Issue

The picture-metaphor is usually characterized as a complete representation. It is complete in the sense that an analog image contains (explicitly or implicitly) all the information present in the stimulus. The quasi-linguistic description is usually a partial representation, since it would be cumbersome (at best) if it described all the information in the stimulus.

I think it can be stated flatly that there are very few (if any) objects or scenes about which people have complete perceptual knowledge. A complete representation of some object implies that its perceptual representation can be discriminated from that of any different

object, regardless of their similarity, context, or any other factors. This simply is not the case. For example, I performed an informal study for perceptual memory of a building in which all of the subjects had worked for at least two years. They could not discriminate an accurate drawing of the building from three qualitatively similar drawings which differed substantially in the overall height-to-width ratio and the number of windows per floor (see Norman, 1975). In fact, the most preferred drawing was the least accurate one, and the least preferred drawing was the most accurate one.

The real issue is not whether perceptual knowledge is complete or partial, but what part of the information is encoded and why it is selected. I favor an approach based on pragmatic considerations of contextual factors. The building mentioned above stands among a group of other gray, concrete buildings similar in architectural style, but quite different in overall size, height, specific detail, and (of course) position. What information is required to discriminate this building from the others? Relative position alone would be sufficient, or height, or size, or any combination of these three parameters. One simply does not need to know details like the exact height-to-width ratio or the number of windows per floor to find it.

The strongest form of the pragmatic view implies that only discriminative information will be represented at any given level. The models developed by Quillian (1968) and Winston (1973) for representing categories and their instances follow this general principle. I am suggesting that it might be fruitful to extend this strategy to contextually dependent descriptions (cf., Norman & Bobrow, 1975; Bobrow & Norman, 1975). That is, a person may not encode the perceptual description of a building in terms of its similarities and differences relative to all other buildings, but relative to just those buildings from which it must be discriminated in its given context or set of contexts.

This strongly pragmatic view leads to certain testable predictions. Suppose that a subject is presented with a designated target stimulus, S_t , in a set of other stimuli, and is told that he or she will later be asked to pick the target item from that set. Performance on subsequent recognition tests for S_t and various distractor items should depend strongly on the context in which they originally appeared. For example, the discriminability of distractors that differ from S_t on some dimension should depend whether that dimension was a discriminative feature of S_t within the presented contextual set. Reaction time measures might provide some insight into what information was stored explicitly (in the representation of the target item) and what information was stored implicitly (in the contextual representation).

This is but one example of the kind of experiments that might provide insight into the rules that govern the specific information content of perceptual representations in context.

The Explicit/Implicit Issue

One of the thorniest problems faced by anyone dealing with perceptual representation is that of explicit versus implicit representation of information. In large measure, the difficulty is that of separating structure from process. If a person can perform a task on a perceptual representation requiring certain information, it cannot be readily determined whether the information was stored explicitly and simply retrieved, or whether it was stored implicitly and then made explicit by some inferential process that would not otherwise have been invoked. The analog "picture-metaphor" constitutes a representation in which virtually all information (except point locations) is implicit. The propositional quasi-linguistic description is a representation in which a good deal of information (but not all) is explicit.

Given the astonishing flexibility of the perceptual information processing system, I doubt that many (if any) hard-and-fast generalizations can be made about what information is represented explicitly and what implicitly. It depends strongly on the requirements of the task being performed. In reading coherent, connected discourse, for example, it may be that only global information is represented explicitly because expectations derived from linguistic context are strong enough to allow identification without detailed processing of low-level components. In comparing two novel figures for a same/different judgment, however, the low-level components may be represented specifically because their use is required for the task. Perhaps the best one can hope for is to draw conclusions about the representation of information in a given contextual situation with certain analyzable task requirements and possible strategies.

The Holistic/Atomic Issue

The issue of whether perception is holistic or composed of smaller components has a long history in psychology, most notably in the confrontations between the structuralist and Gestalt movements. This problem has been resurrected to some extent in the analog/propositional controversy. The strong form of analog representation is decidedly holistic, while the strong form of propositional representation is atomic (or componential).

Structural information. Consider the simple form shown in Figure 3. It is quite easily decomposed into a triangle and rectangle in a particular arrangement. These parts, in turn, have lower-level components in the angles and lines of which

they are composed. Recent data (Palmer, 1974; Reed, 1974) have demonstrated that people can find a "good" or "natural" part within a figure (i.e., a subset of the figure corresponding to a single structural unit such as triangle ABE) much more quickly and accurately than a "bad" or "unnatural" part (i.e., a subset which crosses structural boundaries, such as segments AB, BE, and E). I find it difficult to explain such results without positing some representation of component structure.

The Gestalt maxim "the whole is more than the sum of its parts," however, cannot easily be denied. Placing a series of points in the configuration of a line, for example, adds new dimensions to the figure -- the "emergent" properties of length and orientation which are undefined for the individual points that comprise it. Similarly, arranging three lines to form a triangle produces properties like closedness, area, and symmetry which are not properties of the lines alone. The thrust of this argument is that if emergent properties are to be given explicit representation, there must be a mechanism for encoding both parts and wholes.

A simple formalism for doing this is the hierarchical network (c.f. Baylor, 1971; Palmer, 1975; Winston, 1973). At each level a node representing the global unit subminimates its local parts. An example is given in Figure 3. The node [ABCDE] represents the whole figure, nodes [ABE] and [BCDE] represent the triangle and rectangle parts, nodes [AB], [BC], [CD], etc. represent the lines comprising the parts of the triangle and rectangle, and nodes [A], [B], [C], etc. represent the individual points. There are several aspects of the representation worth noting. The first is that if all the points were shown, then the lowest level of the structure would correspond to an "analog" representation. It is not until the second level that any interpretive information is represented -- e.g., the grouping of points into lines. Second, the structure is a network rather than a tree. A given unit can be part of more than one higher order unit -- e.g., segment [BE] is an element of both the triangle and the rectangle. Third, we have not yet specified the nature of the structural units or the connectors between them. The nodes might stand for feature lists, propositional structures, or even mini-templates. The connections might be uni- or bi-directional associations, explicit part/whole relations, or parametric relations such as relative position, orientation and size.

Parameters. Any structural unit can be characterized by values along certain dimensions -- its size, shape, color, position, and so forth. In the picture-metaphor, all such dimensions are represented holistically and implicitly in the image. They are encoded integrally and can only be separated by some set of processes which extract this information from the holistic image. In the language-metaphor representation, most

dimensional information is encoded componentially and explicitly -- e.g., SIZE (OBJECT, LARGE), COLOR (OBJECT, RED), and so forth. Here the parameters are represented separably and can only be integrated by some process that combines them into a composite unit (such as a point in a multidimensional space). What evidence is there for the integrality or separability of perceptual parameters.

Both personal experience and psychological results suggest that people sometimes can remember, say, the location and shape of an object without remembering its orientation (cf., Frost & Wolf, 1973). At least under certain circumstances, then, some kinds of parametric information must be separable, since there is no reasonable mechanism by which components of integrally stored information can be differentially forgotten.

There are certain dimensions, however, that seem to function integrally. Perhaps the clearest example is that of color. Hue, saturation, and brightness function as a package which is only "unpacked" into its components under unusual circumstances. (The reader is referred to Garner (1974) for a clear and elegant presentation of the differences between integral and separable dimensions in terms of experimental tasks.) There is also some evidence that distantly related parameters (e.g., color, size, and shape) can function integrally in simple same/different tasks. Egeth (1966) and Nickerson (1967) have reported reaction-time results with multidimensional stimuli which are consistent with a matching process that first evaluates all parameters integrally and then evaluates each component separately (see Reed, 1973, pp. 58-61).

Given the evidence, I conclude that parametric information should have both holistic (integral) and atomic (separable) levels of representation, with certain constraints. For example, Figure 4 shows a possible representation for a rectangle of a given size and color. The node representing the whole object has parametric information associated with it -- color and size. Thus, these dimensions are integrated at the object node, such that it could be evaluated as a point in a multidimensional space. Color and size, however, are separated at the next level, and could be processed separately if required by the task. At this level, the components of color (hue, saturation, and brightness) and size (length and width) are represented as integral units, such that each could be evaluated as a point in different multidimensional spaces. These components are then separated at the lowest level. Thus they could be processed separately if necessary.

I do not know whether this formulation is consistent with all of the data, but the general thrust seems right. Certain types of parametric information (e.g., length and width) are represented integrally at one level and separately at a lower level. Other types of information (e.g., length and hue) are never represented as a

self-contained integral unit. To the extent that this is a reasonable scheme for representing parameters, we must begin to examine the structural aspects of parametric information.

The Qualitative/Quantitative Issue

Another recurrent theme in debates over the relative merits of analog and propositional representation is that of whether qualitative or quantitative information is encoded. The typical claims are that analog representations are quantitative while propositional representations are qualitative.

Structure. Structural relationships themselves -- e.g., PART OF (EYE, FACE) -- are qualitative in nature. If perceptual structure is to be represented explicitly, there must be a qualitative mechanism for doing so. I am not certain, however, that it does not have quantitative aspects. A particular perceptual element might be grouped within two different units with different "strengths", especially during the initial process of assigning structure.

I have proposed that the "goodness" of a set of components as a structural unit is determined by context-sensitive associations between its elements (Palmer, 1974). Suppose there is a pattern composed of elements A, B, C, D, E, and F within which A, B, and C are a possible structural unit. The more strongly A, B, and C are associated with each other, according to Gestalt principles of organization, the "better" they are as a part within the figure. The more weakly A, B, and C are associated with D, E, and F, the "better" they are as a part within the figure. This idea has much in common with hierarchical clustering programs. I conceive of the process of organization as one in which quantitative, low-level associations between elements interact with each other -- strengthening and weakening various groupings -- until a stable organization determines "the" structure perceived.

A similar situation may occur in categorical assignments. People are able to make discrete decisions about whether a robin is a bird. However, current evidence suggests that quantitative processes are at work prior to the decision (see Smith, Shoben & Rips, 1975). Certain birds are rated as "better" exemplars than others (Rosch, in press) and the "better" the bird, the faster it can be categorized (Rips, Shoben, & Smith, 1974). The point is that the qualitative discrete result of a cognitive process may be reached by quantitative processing of quantitative information.

Parameters. There are two separate issues for parametric information with regard to quantitative versus qualitative information. The first concerns the nature of the dimensions themselves; the second concerns the nature of the values along the dimensions.

Clearly, the perceptual dimensions (or "qualities") of the representation are qualitative. Color is undeniably different in nature from location or shape or any other perceptual parameter. If this dimensional information is to be encoded explicitly, the representation of each dimension must be qualitatively distinct.

The nature of dimensional values is a more complex problem. Is the color of some object to be represented by a quantitative value or is it to be represented by a qualitative category such as COLOR (OBJECT, RED)? The question is really about the nature of generalization along dimensions. Suppose I show someone a patch of a color between red and orange, but enough toward red that the observer would call it "red" if forced to choose. What will the subsequent recognition errors look like? If the color is represented by coordinates in the color space, then the errors should be distributed as a simple function of distance from the color presented. If the color is represented by a qualitative, categorical description, then the errors should be distributed either uniformly within the category boundaries or as a simple function of distance from the categorical prototype, depending on whether the categories are defined by boundaries or prototypes. Although I do not know of such an experiment, I expect that the results would be more like the quantitative prediction than the qualitative-categorical one. It seems likely, though, that if there were a drift or skew in the errors, it should be toward the center of the category into which the presented stimulus fell. To some extent, I favor quantitative representation of parametric values because it seems simpler and more natural. Metric operations on parameters, for example, are simpler to accomplish with quantitative values. Of course, one can approximate a quantitative scale by having many small categories, but beyond some reasonably small number, the categories lose their qualitative nature and begin to look more like a quantitative dimension. That is, the categories become functionally quantitative.

The Absolute/Relative Issue

Can the information in a perceptual representation be best characterized as encoded in an absolute form or an encoded relative to some other information? Given the emphasis I have placed on contextual factors, it should come as no surprise that I strongly favor the relative approach. But once again, there are constraints which limit the generalizations that can be made.

Structure. The structural organization of perceptual representation is clearly a relative concept. It would be foolish to think that any particular set of perceptual elements will always be grouped together. The Gestalt demonstrations of "hidden figures" are unequivocal evidence that organization depends on (in fact, is largely determined by) contextual factors. Similarly, my own research on the

part-structure of figures shows that a given set of segments will only be perceived as a structural unit within certain contextual segments (Palmer, 1974). In my view, no reasonable case can be made for absolute determination of structural organization.

Parameters. The issue is more complex for parametric information. Since people can recognize most objects at various orientations, perspectives, and distances, relative encoding seems desirable. But it is a fact that people are not equally good at recognizing objects from all orientations and perspectives. There seem to be preferred orientations and perspectives for most objects which depend on experience with them. There are even some rather subtle recognitions (e.g., a particular person's face) which cannot be made if the absolute orientation is too different from that normally experienced. (The reader is referred to Rock's (1974) studies of orientation for further information.) How are we to resolve these apparent inconsistencies?

One type of solution is shown in Figure 1 for orientation information. Both "absolute" and relative orientations of the face and eyes are encoded. The "absolute" orientation (shown in brackets) of a structure is encoded as the orientation of its long axis (from broad to narrow) relative to the ground. The important parameters for the representation of the face schema or frame are the relationships (shown in ovals) between the "absolute" orientations of the eyes and face. If the face were rotated, the absolute orientations of each structure change, but their orientations relative to each other remain constant. Thus, the representation is not orientation specific, and a face can, in principle, be recognized as a face in any orientation.

The absolute orientations have two functions. In the expectation-driven mode, the orientation of a known (or hypothesized) structure will determine the orientation of an expected structure via the relationships between them. For example, if some set of data have been tentatively interpreted as a face at +45, then the expected orientation of the eyes are +135 and -45 as computed from the encoded relationships. The other function of absolute orientation information is to suggest possible interpretations for sensory data in the data-driven mode. It is here that the preferred orientation will affect perceptual identification. A face, for example, is more likely to be suggested as a possible interpretation for an oval at 80° orientation than at a 90° orientation. According to this account, the difficulty in recognizing forms at atypical orientations is due to a reduction in the effectiveness of the data-driven mode. For all objects that have typical absolute orientations, however, using them to generate possible interpretations is an efficient heuristic. A price is paid when the orientation is wrong, but the net effect is probably favorable.

Simply stated, I propose that parametric information be represented relative to its immediate structural context: superordinate whole and the other component parts at that level. Some problems remain. Should all types of parametric information be represented in this way, or are there some types (e.g. color) that are better represented absolutely. What constitutes a level within a given object, and what level of one object corresponds to a given level of another object when they appear together? Kosslyn's (1974) recent experiments in imagery lead me to the view that levels within a scene may be defined by parameters like the overall size of the perceptual element. If the perceptual system analyzes only at one level of resolution at a time, then those perceptual elements that fall within that level will be encoded together and relative to each other.

Conclusion

It was my original intent to clarify some issues involved in the analog/propositional debate in order to reach a deeper level of understanding about the nature of perceptual knowledge. After having made this attempt, I wonder whether anything has really been clarified. At first glance, the big, messy problem of whether perceptual representations are analog or propositional has become a series of almost-as-messy, little problems. To make matters worse, these little problems are subtly intertwined within the fabric of a complex and integrated system. Even so, I think we are on firmer ground with the more basic problems and their interrelationships than we were before. This does not mean that we should focus myopically on the smaller problems, for in the end they must fit together into a coherent whole. In analogy with the view presented earlier that there are both separable parts and integrated wholes, we must work at both levels in our attempts to understand the nature of perceptual representation.

During the course of the discussion, I have presented my own approach to an integrated system for representing perceptual knowledge. At present, it is but a sketch of what a complete theory might look like. (The system is presented in greater detail in another paper (Palmer, 1975) to which the reader is referred for further information.) In developing this formalization I have tried to build-in a number of general properties. A few of the desirable features are as follows.

1. Flexibility and generality. The proposed system blends analog, feature, and propositional representation into a single formalism, even though the representational format is itself propositional. As noted earlier, the primitive point-level representation is (second-order) isomorphic with the stimulus and can be viewed as an analog representation. At each level within the network, the structural units of the representation are associated with the

global features of that unit at its holistic level of resolution. The various units are organized into structural networks characteristic of many propositional systems. The blend is achieved naturally from a simple rule: each structural unit (or "schema" or "frame") is represented globally by its holistic parameters and locally by its structural parts.

2. Variable resolution. The proposed networks are capable of representing different level of resolution through the different structural levels at which they are described. Each scene, object, or part has many possible levels at which it can be examined. The appropriate level of analysis will vary with the stimulus itself (e.g., its size on the retina), the task at hand, and the contextual information available.

3. Functional autonomy. The inclusion of global properties for each structural unit gives them functional independence (within limits). In the data-driven mode, any unit may be activated directly by sensory data via the global parameters without first activating the lower level units. This means that analysis can begin at virtually any level within the network. In the expectation-driven mode, any unit can look for confirmation of its global parameters before requesting more finely resolved information from lower structural levels.

4. Interactive capability. The structural and parametric relationships within the network provide communication mechanisms through which the units can interact in both data-driven and expectation-driven modes. The current state of one unit can be communicated to related units such that positive results for related units serve to strengthen each other. This allows for a "bootstrap convergence" process among the units of a schema during interpretation.

5. Contextual dependence. The structure of the network and the relative representation mechanisms give contextual factors paramount importance in a simple but powerful way. Given a non-random world in which contextual regularities are frequent, the ability to use this information is desirable, possibly even necessary.

Regardless of whether my particular approach proves to be a good way to realize these objectives, I have confidence that they are important in the design of a workable model of perceptual representation and processing.

REFERENCES

Baylor, G.W., A treatise on the mind's eye: an empirical investigation of visual mental imagery. Unpublished doctoral dissertation, Carnegie-Mellon University, 1971.

Bobrow, D.G., Dimensions of representation. In D.G. Bobrow and A.M. Collins (eds.) Representation and Understanding. New York: Academic Press, 1975.

Bobrow, D.G. and Norman, D.A., Some principles of memory schemata. In D.G. Bobrow and A.M. Collins (eds.) Representation and Understanding. New York: Academic Press 1975.

Bower, G.H., Mental imagery and associative learning. In L.W. Gregg (ed.) Cognition in Learning and Memory. New York: Wiley, 1972.

Cooper, L.A., Mental transformation of random two-dimensional shapes. Cognitive Psychology, 1975.

Cooper, L.A. and Shepard, R.N., Chronometric studies of the rotation of mental images. In W.G. Chase (ed.) Visual Information Processing. New York: Academic Press, 1973.

Egeth, H.W., Parallel versus serial processes in multidimensional stimulus discrimination. Perception and Psychophysics, 1966, 1, 245-252.

Frost, N. and Wolf, J., How "visual" is visual memory? Paper presented at the 14th Annual Meeting of the Psychonomic Society, St. Louis, Missouri, November 1973.

Garner, W.R., The Processing of Information and Structure. Potomac, Maryland: Erlbaum, 1974.

Kosslyn, S.M., Constructing visual images: An exercise in neo-mentalism. Unpublished doctoral dissertation, Stanford University, 1974.

Kosslyn, S.M. and Pommerantz, J.R., Mental imagery reconsidered: An analysis of Pylyshyn's critique, in preparation.

Nickerson, R.S., Same-different reaction times with multi-attribute stimulus differences. Perception and Motor Skills, 1967, 24, 543-554.

Norman, D.A. and Bobrow, D.G., On the role of active memory processes in perception and cognition. In C.N. Cofer (ed.) The Structure of Human Memory. San Francisco: W.H. Freeman, 1975.

Norman, D.A. and Rumelhart, D.E., Memory and Knowledge. In D.A. Norman, D.E. Rumelhart, and LNR Research Group. Explorations in Cognition. San Francisco: W.H. Freeman, 1975.

Palmer, S.E., Structural aspects of perceptual organization. Unpublished doctoral dissertation, University of California, San Diego, 1974.

Palmer, S.E., Visual perception and world knowledge. In D.A. Norman, D.E. Rumelhart, and LNR Research Group, Explorations in Cognition. San

- Francisco: W.H. Freeman, 1975.
- Pylyshyn, Z.W., What the mind's eye tells the mind's brain: A critique of mental imagery. Psychological Bulletin, 1973, 80, 1-24.
- Quillian, M.R., Semantic Memory. In M. Minsky (ed.) Semantic Information Processing. Cambridge, Massachusetts: MIT Press, 1968.
- Reed, S.K., Psychological Processes in Pattern Recognition. New York: Academic Press, 1973.
- Reed, S.K., Structural descriptions and the limitation of visual images. Memory and Cognition, 1974, 2, 329-336.
- Rips, L.J., Shoben, E.J. and Smith E.E., Semantic distance and the verification of semantic relations. Journal of Verbal Learning and Verbal Behavior, 1973, 12, 1-20.
- Rock, I., Orientation and Form. New York: Academic Press, 1974.
- Rosch, E., Cognitive representations of semantic categories. Journal of Experimental Psychology, General, in press.
- Shepard, R.N. and Metzler, J., Mental rotation of three-dimensional objects. Science, 1971, 171, 701-703.
- Smith, E.E., Shoben, E.J. and Rips, L.J., Structure and processing in semantic memory; A feature model for semantic decisions. Psychological Review, 1974, 81, 214-241.
- Winston, P.H., Learning to identify toy block structures. In R.L. Solso (ed.) Contemporary Issues in Cognitive Psychology: The Loyola Symposium. Washington, D.C.: Winston, 1973.

Figure 1:
An illustration of structural and parametric differences.




Figure 2:
An example of stimuli used in the mental synthesis task and corresponding synthesis times (Palmer, 1974).

Synthesized Figure		MEAN SYNTHESIS TIME (in seconds)	
		Figure shown	All 20 figures
Good Parts		2.16	2.04
Bad Parts		5.03	5.05

Figure 3:
An example of structural organization showing points, lines, parts, and the whole figure.




Figure 4:
The structure of integral and separable parameters.




Figure 5:
An illustration of relative and absolute representation of orientation parameters for a FACE and two EYES.

