Computational Models of Natural Language Processing
B.G. Bara and G. Guida (eds.)

SEMANTIC PRIMITIVES OR MEANING POSTULATES: MENTAL MODELS OR PROPOSITIONAL REPRESENTATIONS?

Philip N. Johnson-Laird

MRC Applied Psychology Unit
Cambridge, UK

The controversy over whether images are equivalent to sets of propositions, and the controversy over whether semantic primitives are equivalent to sets of meaning postulates, are currently conducted independently. This article brings them together in relation to the problematical nature of the transitivity of spatial relations. Both controversies can be resolved by assuming that spatial relations are represented as mental models constructed by a procedural semantics that decomposes the meanings of words into more primitive notions.

1. INTRODUCTION

This paper confronts two controversies with one phenomenon. The first controversy concerns the nature of mental representations and whether there is any essential difference between an image and a set of propositions. The second controversy concerns the meanings of words and whether they are represented by semantic primitives, or by meaning postulates that specify the semantic relations that hold between them. Both controversies are currently in progress and both are primarily matters of theory rather than fact. They have not previously been related to one another; yet, if the fundamental argument of this paper is sound, they can both be resolved by confronting them with one and the same phenomenon -- a seemingly trivial aspect of everyday inference: the transitivity of certain spatial relations. A necessary preliminary, however, is to lay out the two controversies in more detail.

2. THE FIRST CONTROVERSY: IMAGES VERSUS PROPOSITIONS

Many human beings claim to be able to form and to manipulate mental images in the absence of corresponding visual stimuli. The phenomenon has been studied ever since Galton's questionnaire on his correspondent's ability to imagine their breakfast tables (Galton, 1928, originally published in 1880). More recent investigations have examined a variety of aspects of images, incluiding their use as mnemonics (Bower, 1972; Paivio, 1971), their mental rotation and transformation (Cooper, 1975; Shepard, 1975), their suppression by other tasks (Brooks, 1967; Byrne, 1974), and their use in retrieving information about objects (Hayes, 1973; Holyoak, 1977; Kosslyn, 1975, 1980; Moyer, 1973; Paivio, 1975). No one seriously doubts the existence of the subjective experience of imagery. What is problematical, however, is the explanation of the experience and the ultimate nature of images as mental representations. It seems unlikely that they are simple pictures in the head, because this conjecture leads to a number of undesirable consequences including the need for an homunculus to perceive the pictures, and accordingly to the imminent danger of an infinite regress (Dennett, 1969). There remain two schools of thought.

On the one hand, there are those who argue that an image is distinct from a propositional representation (Bugelski, 1970; Kosslyn and Pomerantz, 1977; Paivio, 1971, 1977; Shepard, 1975, 1978; Sloman, 1971). These authors attribute a variety of properties to images. The most important are as follows:

1. The brain processes underlying an image are similar to those underlying the perception of an object or picture.

2. An image is a coherent and integrated representation in which each element of a represented object occurs only once.

3. An image is amenable to apparently continuous mental transformations, such as rotations or expansions, in which intermediate states correspond to intermediate states or views of an actual object undergoing the same transformation. Hence, a small change in the image corresponds to a small change in the object or its appearance.

4. Images represent objects. They are analogical in that the structural relations between their parts correspond to those between the parts of the objects represented. There may indeed be an isomorphism between an image and an object, though this claim only makes sense with respect to an object viewed as decomposed into parts with particular relations between them.

On the other hand, there are theorists who argue that the subjective experience of an image is epiphenomenal and that its underlying representation is propositional in form (Anderson and Bower, 1973; Baylor, 1971; Kieras, 1978; Moran, 1973; Palmer, 1975; Pylyshyn, 1973). These authors attribute a variety of properties to propositional representations. The most important are as follows:

1. The brain processes underlying a propositional representation are similar to those underlying the perception of an object or picture; as Norman and Rumelhart (1975) remark, "the first stage of conscious awareness is already in the form of propositions."

2. The same element or part of an object may be referred to by many of the different propositions that constitute the description of the object. However, when propositions are represented in the form of a semantic network, then the representation can be thought of as coherent and integrated, with each element of the represented object occurring only once with all its relations to other elements readily accessible.

3. A propositional representation is discrete and digital rather than continuous. However, it can represent continuous processes by small successive increments of the relevant variable(s), such as the angle of an object's major axis to a frame of reference. Hence, a small change in the representation can correspond to a small change in the object.

4. Propositions are true or false of objects. Their representations are abstract in that they do not resemble either words or pictures, though they may be needed to provide an interlingual[1] between words and pictures (Chase and Clark, 1972). Their content carries the significant information about what they represent.

The critics of imagery often allow that an image can be constructed from its propositional description, but such an image does not introduce any new information, it merely makes the stored description more accessible and easier to manipulate. Gelernter's (1963) program for proving geometric theorems, and Funt's (1977) program for making inferences about the stability of arrangements of blocks, are both considerably enhanced by the use of procedures that operate on diagrammatic representations. However, Pylyshyn (1973) argues that picture-like representations are not necessary for such purposes: the same function can be served by propositional descriptions. This view has been pushed still further by Palmer (1975):

> The arguments in favor of analogical representations tend to emphasize the relative ease with which certain operations can be performed on them compared to the difficulty in performing the same operations on propositional representations. These arguments, however, generally overlook the fact that propositions can encode quantitative as well as qualitative information. In addition, it is not often recognized that propositions are capable of encoding an analog image.

Palmer then goes on to establish both a way in which a shape such as a triangle can be encoded propositionally and a method for rotating such representations once they have been decomposed into their propositional constituents.

Evidently, the two sorts of representation share a number of properties; they differ mainly on the fourth of the characteristics listed above, i.e. whether they represent objects or describe them. This apparent similarity and the view that they are readily transformed into one another has indeed led some commentators to conclude that the controversy is neither fundamental (Norman and Rumelhart, 1975) nor resolvable (Anderson, 1976, 1978). In particular, Anderson (1978) argues that "any claim for a particular representation is impossible to evaluate unless one specifies the processes that will operate on this representation." He shows that given certain assumptions a theory based on images can be mimicked by one based on propositions. As we shall see, however, there are some important exceptions to the principle.

3. THE SECOND CONTROVERSY: SEMANTIC DECOMPOSITION VERSUS MEANING POSTULATES

Semantic theorists are generally agreed that a major burden for the meanings of words is to account for the relation between such assertions as "Polly is a parrot" and "Polly is a bird". If the first assertion is true, then plainly, so is the second. There is a major disagreement between theorists about the nature of the semantic machinery needed to explain such relations. One school of thought, whose recent ancestry can be traced back to the work of Katz and Fodor (1963) -- though it has a much longer history reaching back into antiquity -- holds that the meaning of a word such as "parrot" is represented in the mental lexicon as a set of semantic elements that includes, amongst others, those corresponding to "bird". The relation between the two sentences is accordingly captured by the semantic primitives in the entries of the mental lexicon. A wide variety of psychological theories of meaning are committed to some sort of semantic primitives (Clark and Clark, 1977; Collins and Quillian, 1972; Miller and Johnson-Laird, 1976; Norman and Rumelhart, 1975; Schank, 1975; Smith, Shoben and Rips, 1974).

An alternative view is that there are no semantic primitives (Fodor, 1976; Fodor, 1977; Fodor, Fodor and Garrett, 1975; Kintsch, 1974; Lyons, 1977). Entailments that depend upon the meanings of words are, according to these theorists, captured by meaning postulates. A meaning postulate as formulated by Carnap (1956) is an assertion that stipulates the semantic relations between words of a language, generally a formal one, e.g. _for any x, if x is a parrot then x is a bird_. Such rules are introduced into a model-theoretic semantics of a language in order to render some models inadmissible, namely, those for which the meaning postulates are not true. Latterly, the idea has been cut loose from formal semantics and imported into psychological theory. Fodor _et al_ (1975) assume that sentences in a natural language are translated into a corresponding mental language, and that meaning postulates couched in the mental vocabulary are used to make inferences from the translated sentences.

Although there have been attempts to resolve the controversy experimentally, the results so far are equivocal. There is evidence against the idea that comprehension depends on a process of semantic decomposition (Kintsch, 1974; Fodor _et al_, 1975); there is evidence against the idea that it depends on meaning postulates (Clark and Clark, 1977; Johnson-Laird, Gibbs, and de Mowbray, 1978). But, as yet, there are no results to resolve the issue of whether or not the mental representations of meaning depend on semantic primitives. Indeed, there has been a tendency to accept the view of Katz and Nagel (1974) that there is no fundamental distinction between semantic primitives and meaning postulates. There are, in fact, several arguments that could be made to establish a difference in their psychological plausibility. I shall present only one: it is the argument that will also resolve the controversy about images and propositions.

4. TRANSITIVITY IN EVERYDAY INFERENCE

My central thesis concerns simple inferences based on premises in ordinary language. Consider the following argument about a pencil, a box, and an envelope:

The pencil is in the box.

The box is in the envelope.

∴ The pencil is in the envelope.

Is this inference valid? Before we can answer this question, it would be as well to remind ourselves that an inference is valid if it is impossible for its premises to be true and its conclusion false. Obviously, the present inference _is_ valid since given that the premises are true, the conclusion must be true, too.

Meaning postulates provide an initially plausible basis for a psychological mechanism that makes such an inference. The premises are translated into a propositional representation, which according to Kintsch (1974) might take the following sort of form:

(IN, PENCIL, BOX)

(IN, BOX, ENVELOPE)

and then the meaning postulate that captures the transitivity of "in":

For any x, y, z (If (IN, x, y) & (IN, y, z)) then (IN, x, z)

is applied to the representation of the premises to yield the conclusion:

(IN, PENCIL, ENVELOPE).

And this propositional representation can, if necessary, be translated back into natural language.

Although the details of the various processes of translation have not been formulated explicitly by any theorist, they are not at issue as far as the present paper is concerned. Its argument applies to any processes that lead parsimoniously to propositional representations and to the application of meaning postulates to them. There is nothing privileged about meaning postulates here, they may be replaced by any rules of inference that apply to propositional representations.

The heart of my argument depends on the following sort of inference:

Luke is on Mark's right.

Mark is on Matthew's right.

∴ Luke in on Matthew's right.

It is not immediately clear whether this inference is valid. If the three individuals are sitting in a straight line on one side of a table, then the relation denoted by "on X's right" is transitive, and the inference is valid. But if they are sitting at equal intervals round a small circular table, then the relation denoted by "on X's right" is not transitive, and the inference is invalid.

A natural way to try to accommodate this phenomenon within the framework of a propositional theory is to propose two different meanings for "on the right" and its cognates, one to which a meaning postulate expressing transitivity applies, and one to which a meaning postulate expressing intransitivity applies. Indeed, only this manoeuvre will save the theory if it is to account for the way that ordinary inferences are made and evaluated. However, if a number of people are seated round a large circular table, then the previous inference could be valid, but one might have doubts about the following one:

John is on Luke's right.

Luke in on Mark's right.

Mark is on Matthew's right.

∴ John is on Matthew's right.

As more and more individuals are added round the table, there will inevitably come a point where transitivity breaks down. (As a matter of fact, there is likely to be a region of uncertainty, but this possibility merely exacerbates the problems of a meaning postulate theory). In general, "on X's right" may denote an intransitive relation, or relations with a transitivity that varies over any number of items from three to an arbitrarily large number. Each of these extents requires its own separate meaning postulate with the number of premises in its antecedent directly correlated with the number of items over which transitivity holds -- two premises for transitivity over three items, three premises for transitivity over four items, and so on <u>ad infinitum</u>. Because there is no limit to the number of items at which transitivity ceases to hold, there is no limit to the number of separate meaning postulates that are required to cope with the semantics of this single relation. This conclusion is psychologically unacceptable on the reasonable criterion, decisive in other contexts (e.g. Miller and Chomsky, 1963), that human beings do not have an unlimited capacity for learning or for storing information.

It should be emphasized that these difficulties are not peculiar to "right" and "left". English vocabulary is plagued by the same sorts of problem, and it is hard to find any simple spatial term that has an unequivocal logic.

A proponent of meaning postulates might argue that once the transitivity of "on the right" ranges over some large number of items, say, 100, then it can be taken to have an unlimited extent. This <u>ad hoc</u> proposal has at least the virtue of limiting the required meaning postulates to a finite number. Yet, it does not solve the problem: no matter how large the radius of a circle and how densely the individuals are packed around it, it <u>is</u> a circle and transitivity must break down. Another way to save the meaning postulate theory is to assume that there is some machinery for <u>generating</u> meaning postulates. However, this proposal highlights another difficulty. It is clear that any feasible system will depend on some mechanism for determining the nature of the situation referred to explicitly or implicitly by the premises. In the case of our examples, it will depend on information about the table and the seating arrangements, which in turn will be used to select the appropriate meaning postulate.

Once the need to deal with reference situations is admitted, a still more obvious, though deliberate, gap in the meaning postulate theory of semantics becomes apparent. It is again best illustrated by a simple example. Given the following arrangement of letters:

B A

any competent speaker of English knows that it is true to say of them, "A is on the right of B" and false to say of them, "A is on the left of B". This distinction reflects the difference in meaning between "right" and "left"; yet, there is no way to capture it using meaning postulates. One can, of course, establish that there is a difference in meaning between the two terms, e.g. for any x and y, x is on the right of y if and only if y is on the left of x, and for any x and y, if x is on the right of y then x is not on the left of y. Such postulates establish that a difference exists, but they do not specify its nature. For that, it is necessary to make explicit what it is that underlies our knowledge that A is indeed on the right of B in the example above.

The idea lying behind the psychological exploitation of meaning postulates, and indeed most psychological theories of meaning, is that it is feasible to specify the semantic relations between words without considering how they relate to the world (see Johnson-Laird and Herrmann, 1983, for the implications of this assumption for semantic network theory). Intensions can be profitably pursued independently from extensions. The principle seems plausible for meaning postulates in their original context of formal semantics, where the real world is replaced by a model, or a set of models, in which the extensions of terms are assigned directly. But the precedent is misleading for natural language where, as we shall see, the only way to account for the proper relations between words, and for inferences based upon them, is by giving a complete specification of their meanings, including their relations to the world. What is deliberately missing in the meaning postulate account is a definition of how "right" and "left" relate to the world. The reason for this omission is obvious: the relations are so basic that there is no way to define them in ordinary English. It is for this reason that a complete theory of meaning must rely upon some more primitive notions.

Is it possible to save a propositional theory by sacrificing meaning postulates? The answer depends, of course, on what processes are used to make inferences in their stead. Any system that relies on rules that manipulate propositions will have to introduce some machinery to handle transitive relations, and hence it will be in imminent danger of falling into precisely the same difficulties. The only escape route will be a method for handling the facts of transitivity without relying on rules, postulates, or productions, for transitivity itself. This prescription may seem to be impossible to fulfil; fortunately, there is at least one way in which it can be met.

5. PROCEDURAL SEMANTICS AND A PROGRAM FOR SPATIAL INFERENCE

The definitions of spatial terms and the uncertainties of their transitivity can be accommodated within a theory based on a "procedural semantics" (see Davies and Isard, 1972; Johnson-Laird, 1977, 1983; Miller and Johnson-Laird, 1976; Woods, 1981). The theory can be illustrated by considering a computer program (written in the list-processing language, POP-10) that I have devised in order to model a theory of spatial inference. The purpose of the program is to evaluate premises about the spatial relations between objects. It works by building up a model in the form of a two-dimensional array that satisfies the premises given to it, and indicates whether a premise is implied by, or is inconsistent with, what it has already been told. It accordingly contains a number of general procedures for constructing, recursively manipulating, and interrogating, sets of arrays. One such procedure, for instance, given the location in an array of one item mentioned in a premise, inserts the other item into the array at a place that is appropriate according to the meaning of the premise. Another general procedure is used to verify whether the relation specified to hold between two items, say, A and B, obtains within an array. It works by locating B and then by scanning along a line from B in order to determine whether or not A is somewhere on that line. If A is found to lie on the line then the premise is true, otherwise it is false. The verification procedure contains two parameters, DI and DJ, whose values specify the direction in which to scan: they give the respective increments on the two axes of the array that define the locations to be examined. This use of parameters to specify directions is common to all the general procedures used by the program, including those for inserting new items into an array. This uniformity makes it possible to define the meanings of relational terms as procedures that work in a way that is utterly remote from meaning postulates and conventional decompositional theories (e.g. Norman and Rumelhart, 1975; Schank, 1975; Smith, Shoben and Rips 1974).

The meaning of "on the right of" consists of a single procedure: FUNCT(% 0, 1 %). This takes whatever general procedure is about to be executed, and which has been assigned as the value of the variable, FUNCT, and "freezes in" the value of 0 to its DI parameter and the value of 1 to its DJ parameter. The decorated parentheses are a standard device in POP-10 for freezing in the values of parameters, with the effect of converting a general procedure into a new more specific one. The effect of FUNCT(% 0, 1 %) on the verification procedure is accordingly to produce a more specific procedure that scans a sequence of locations lying in a particular orientation: the value of one axis is held constant (i.e. incremented by 0) while the value of the other axis is progressively incremented by 1. In other words, if you imagine the array laid out on a table in front of you, the procedure examines a sequence of locations lying progressively further to the right of B. It looks to see whether A is on the right of B. The same process of freezing in the values of parameters is used to convert the program's other general procedures into specific ones that depend on the relation specified in a premise.

The program's lexical entries define how words relate to its world; but they stipulate nothing about transitivity or intransitivity. However, a relation such as "on the right of" has the emergent property of transitivity, that is to say, whenever A is on the right of B and B is on the right of C, then as a matter of fact A will be on the right of C, whether the program is building, manipulating, or interpreting an array. The program can accordingly make transitive inferences even though it contains no rules, postulates, or productions, for transitivity itself. This facility depends on the use of spatial models and procedural definitions that relate directly to them. The definitions express meanings in terms of the primitive components of specific coordinate values that are only interpretable with respect to the spatial arrays. The meaning of a word is accordingly not a procedure that can do anything by itself; it is a procedure that applies to other procedures.

Several theorists have claimed that the apparatus of spatial arrays smuggles in the principle of transitivity by the back door. (Both Jerry Fodor and Zenon Pylyshyn have independently made this suggestion in conversation, though it would be wrong to saddle them with a strong commitment to it). Its plausibility rests on its confusability with a very different claim, namely, that it is possible to define transitive relations over spatial models. Obviously, this condition must hold if a transitive relation is to be represented in a spatial array. Spatial arrays, however, do not have transitivity built into them: if they did, they would be unable to represent intransitive relations, or relations that are transitive to varying degrees. Indeed, if the locus of the individuals in a reference situation is circular rather than rectilinear, then the same lexical entries used in relation to the frame of reference that each individual defines give rise to local transitivity, but sooner or later it breaks down as the individuals depart further and further from the required sequence of locations passing through the initial one in the series. Hence, with the introduction of procedures for representing an individual's frame of reference and the spatial arrangement, it will be possible to deal with the vagaries of transitivity without having to postulate an indefinitely large number of different meanings for a relational term. A single entry in the lexicon will suffice.

The use of meaning postulates to make spatial inferences places an unacceptable demand on the storage capacity of human memory: it becomes necessary to store a vast number of postulates just to cope with the meaning of a single relational term. Semantic primitives, however, do the job parsimoniously. The controversy over semantic primitives and meaning postulates is accordingly resolved: meaning postulates must be replaced by semantic primitives that relate words to mental models of the world. There is no need to suppose, however, that comprehension calls for a process of decomposition into semantic primitives: they play a direct role in the tacit construction of mental models.

6. MENTAL MODELS

The notion of a mental model, which was introduced in the previous section, is to be distinguished from the current conception of an image. It shares the crucial property that its structure represents information, but this structure need not necessarily possess any immediately 'pictorial' features. Of course, a mental model might combine both pictorial and non-pictorial information, such as a spatial array that represents the positions of chess pieces and, by a further convention, the tactical relations between them (cf. Simon and Barenfeld, 1969). But a mental model can also be entirely schematic in that only its topological properties are relevant to the processes that interpret it. This idea can be difficult to grasp, and it is best explained in terms of a recent theory of syllogistic inference and the interpretation of quantified sentences (Johnson-Laird, 1983; Johnson-Laird and Bara, 1983).

According to the theory, a statement about, say, a set of people in a room, such as:

Some of the scientists are drivers

is mentally represented by a model in which an arbitrary number of elements stand for the relevant set of scientists, and a subset of them are linked to other elements standing for drivers:

s = d

s = d

(s) (d)

Here, each <u>s</u> represents a scientist, each <u>d</u> represents a driver, and the arbitrary number of parenthesized elements indicate that there may be scientists who are not drivers and <u>vice versa</u>. The identity signs denote an identity between the entities that they link.

In order to make an inference, as opposed to evaluating a given conclusion, it is necessary to possess a heuristic, because logical principles alone can never determine which of the potentially infinite number of valid, though mainly trivial, conclusions should be drawn from a set of premises. Unfortunately, this point has often been overlooked in psychological theories of reasoning (e.g. Inhelder and Piaget, 1958, 1964) and in theories based on meaning postulates. The theory of mental models lends itself naturally to heuristics for drawing informative conclusions (see Johnson-Laird, 1983). In the case of relational inferences, which include syllogisms, these heuristics reduce to a simple principle: try to form connections that link up end items by way of middle items. Thus, given the premises:

Some of the scientists are drivers

All the musicians are drivers

the heuristic yields the following combined model:

```
s = d = m
s = d = m
(s) (d)
```

where each m represents a musician. It will be easy to trace through the links to derive the conclusion:

Some of the scientists are musicians

or its converse:

Some of the musicians are scientists

The heuristic has, of course, led one into error, though it is an error that many subjects make (Johnson-Laird and Steedman, 1978). In order to guarantee that a conclusion is valid, an attempt should be made to find an alternative model that falsifies it: only if there is no such model is the conclusion valid. In the case of the example, there is an alternative model that falsifies the conclusions:

```
 s = d
 s = d
(s)  d = m
     d = m
```

Mental models of syllogisms are analogical in a crucial way: sets of entities are represented by sets of mental tokens, and relations between entities are represented by relations between the tokens in the model. Hence, a model unlike an image may contain no 'pictorial' information, though many models may take the form of images. The theory does indeed account for performance in syllogistic reasoning (see Johnson-Laird and Steedman, 1978; Johnson-Laird and Bara, 1983), and it does so, like the treatment of transitivity above, without postulating rules of inference. Hence, the theory answers two hitherto embarrassing questions: how children acquire rules of quantified inference, and why adults have no introspective access to them. The answers are simply that no rules of inference are learnt in childhood, and that there is no mental logic. These claims may seem paradoxical, but the paradox disappears once one realizes that the process of testing a conclusion by searching for counterexamples may well be carried out in a haphazard manner -- the process is radically different to the quasi-syntactic manipulation of propositional representations by rules of inference. Nevertheless, if the search for alternative models is carried out accurately, then any conclusions that survive unscathed are logically valid. The logicians who first formulated rules of inference may have relied in part on reflections about the invariant properties of their own mental manipulations of models.

7. MIMICKING THEORIES

Procedural definitions couched for the construction of mental models solve the problem of transitivity. But, how does this solution stand in relation to Anderson's (1976, 1978) proof that one representational theory can be mimicked by another provided that there is a one-to-one mapping between their respective representations? Could one not resolve the controversy about images and propositions by showing that mental models can be translated into a propositional theory? The two questions are intimately related, and in order to answer them we must consider Anderson's theorem in detail.

His argument is intended to establish that given a theory which embodies assumptions about mental representations and processes, it is possible, in principle, to construct other theories with different sorts of representations that nevertheless mimic the original theory by behaving in an equivalent manner. In fact, 'mimicry' is not the right word to describe Anderson's manoeuvre: rather the second theory invades the first and takes it over like a virus taking over an organism's machinery for producing DNA. Suppose, for instance, that one wishes to show that with suitable mental operations, a propositional theory can mimic an imaginal theory. The trick is to embed the whole of the imaginal theory within the operations carried out on the propositional representations. The imaginal theory assumes, say, that a stimulus is encoded as an image, which can be mentally rotated in order to determine whether it coincides with another stimulus. The propositional theory assumes only that a stimulus is encoded as a set of propositions. The imaginal theory is embedded within the propositional theory by the following sequence of operations applied to the propositional representation of the stimulus:

1. Apply the inverse of the propositional encoding to the set of propositions in order to recover the original 'stimulus' (i.e. sensory image).

2. Apply the imaginal encoding to the stimulus in order to obtain the corresponding image.

3. Rotate the image.

4. Apply the inverse of the imaginal encoding to the rotated image in order to obtain the corresponding stimulus.

5. Apply the propositional encoding to the stimulus in order to obtain the set of propositions corresponding to the rotated image.

The decision about whether these propositions match the second stimulus can again, if necessary, rely on the imaginal theory:

6. Apply the inverse of the propositional encoding in order to obtain the stimulus corresponding to the rotated image.
(This stimulus is, of course, identical to the one obtained in step 4.).

7. Apply the imaginal encoding to the stimulus to obtain the corresponding image. (This image is identical to the one obtained from step 3.).

8. Compare the image to the one obtained from the second stimulus, and make the appropriate response.

Although this chain of operations can be postulated, its feasibility depends on two crucial conditions. First, the various functions must be computable. Second, it must be possible to apply the inverse of the propositional encoding to obtain the original stimulus, or, more plausibly, a sensory representation isomorphic to the original stimulus. However, since perception is likely to involve a many-one mapping, the inverse will not be a function since there will be many stimuli that it could produce. It is for this reason that Anderson imposes the condition that there must be a one-to-one mapping between the respective representations of the two theories. Granted this condition, the inverse of the propositional encoding can yield any of the stimuli that could have given rise to the original set of propositions, and it will not matter which stimulus is selected, because they will all be equivalent for the imaginal theory, too.

That a propositional theory can mimic an imaginal theory by importing wholesale the apparatus of images is plainly a trivial result. What is of interest is the possibility of a more direct method of mimicry that does not depend upon embedding one theory within another. Unfortunately, there is no guarantee that a direct method can always be found for two alternative representational theories. Anderson makes only the modest claim: "...it seems we can usually construct [the required operation] more simply than its formally guaranteed specification".

Moreover, if one theory encodes stimuli into classes that do not correspond one-to-one with the encodings of the other theory, then the whole system of mappings breaks down and the theorem ceases to hold.

Considerable care needs to be exercized in drawing conclusions on the basis of Anderson's theorem. He himself (Anderson, 1976, p.74) makes the following claim:

> Any behavior that can be computed from inspecting semantic primitives can be computed with the aid of 'meaning postulates' that interpret more complex semantic units. This follows from the theorem ... that any representation can mimic the behavior of any other, provided they impose the same equivalence class on their inputs.

The first assertion has, of course, proved to be false: meaning postulates cannot handle the vagaries of transitivity, but lexical entries based on procedural primitives can accommodate them. On the assumption that Anderson's theorem is sound, it follows that the two sorts of theory do not impose the same equivalence classes on their inputs. And this conclusion is clinched by considering sentences of the form: "A is in front of B, which is behind C". The sentence is unambiguous[2] and should accordingly receive a single propositional representation, but it is referentially indeterminate -- the relation between A and C is unspecified -- and can accordingly be represented by at least two different mental models. This distinction drives a wedge between sets of propositions and mental models that is not easily removed.

It might be supposed that the propositional representation could mimic the model representation, and yield two alternatives: one in which A is in front of C, and one in which C is in front of A. But, before such alternatives could be specified, it would be necessary to detect the indeterminacy in the first place. In general, a scheme for detection would have to be able to infer that the relation between certain items in a propositional representation was indeterminate. Unfortunately, this requirement leads straight back to the problems of transitivity: whether the relation between certain items is determinate or indeterminate may depend entirely on whether a transitive inference is valid or invalid. Since no finite system of rules based on a propositional representation can handle this problem, it follows that no such system can detect indeterminacies, or *a fortiori* set up alternative representations when they occur. Hence, a theory of propositional representation does not yield the same equivalence class of representations as the class yielded by the theory of mental models. The wedge remains securely in place.

Strangely enough, a major critic of imagery theories has made a related point. Pylyshyn (1973) wrote:

> It would be quite permissible ... to have a [propositional] mental representation of two objects with a relation between them such as 'besides'. Such a representation need not contain a more specific spatial relation such as 'to the left of' or 'to the right of'. It would seem an unreasonable use of the word 'image' to speak of an image of two objects side by side, without the relation between them being either 'to the left of' or 'to the right of'.

This point has empirical consquences. Mani and Johnson-Laird (1982) showed experimentally that spatial descriptions can be mentally represented either by models or by sets of propositions. When subjects formed a model of a description, they often reported that they had used imagery. Hence, the controversy is resolved: images and models can be distinguished both theoretically and empirically from propositional representations.

8. LEVELS OF DESCRIPTION

Is it really true that images are not equivalent to sets of propositions? That was the conclusion of the previous section, but doubtless it will be resisted by propositional theorists. There is one way in which they can sustain their objection, but only at the cost of trivializing the whole controversy. It all depends upon a source of much confusion in theoretical discussions, the level at which a particular theory is described. The issues can be illustrated by considering the problem of how to characterize the computer program that embodies the theory of spatial inference.

One approach is that since the program must ultimately be translated into the machine language of a computer before it can be run, we should concern ourselves with what the machine language instructions cause to happen in the machine -- the shifting of bits from one location in store to another, and so on. But this approach is misguided: the details of a specific implementation should not concern us. We should not worry about the particular computer and its machine code, since the program could be executed on some very different machines, and we do not want to make a separate characterization of the program for all these different sorts of computer. An alternative approach is provided by Scott and Strachey (1971), the pioneers of formal semantics for computing languages:

> Compilers of high-level languages are generally constructed to give the complete translation of the programs into machine language. As machines merely juggle bit patterns, the concepts of the original language may be lost or at least obscured during this passage. The purpose of mathematical semantics is to give a correct and meaningful correspondence between programs and mathematical entities in a way that is entirely independent of an implementation.

There is a very important analogy for psychologists here: psychology (the study of the algorithms) can be pursued independently from neurophysiology (the study of the machine and the machine code). The argument also provides a useful antidote to the excessive scepticism that can be induced by theorems demonstrating how one sort of representational theory can be mimicked by another. In order to try to substantiate this claim, and to clear up the confusion over levels of description, let us continue the characterization of the spatial inference program.

"It works by building up a two-dimensional array that satisfies the premises given to it." This description of the program is informal but at a high level, the level of psychological discourse. You may wonder how exactly an array is represented by the programming language. It is, in fact, a data structure of one or more dimensions in which the elements can be accessed and updated by giving appropriate coordinates. (An array is also a function in POP-10, which permits it to be represented by a rule rather than an explicit table). A programmer needs to know no more, since procedures for manipulating arrays can be written simply by thinking of them as n-dimensional spaces where each location is specified by an n-tuple of integers. A student of the 'psychology' of computers, however, may be curious about the invisible machinery that makes such an array possible. Its representation in the computer does not involve an actual physical array of locations in memory. That is quite unnecessary. Indeed, the physical embodiment of an array is irrelevant. What matters is that it should function as an array, that is, it has a set of addresses that are functionally equivalent to an array, its elements can be accessed as in an array, and its contents displayed or printed out in the form of an array. A psychological description should accordingly be a functional one.

Consider a program for spatial inference in which an assertion such as, "A is on the right of B" is represented by the following formulae: AT(A,1,6), AT(B,1,2) where the first integer in each pair gives the y-coordinate and the second integer gives the x-coordinate of the relevant item. The general procedure for verification works by looking for sequences of ordered pairs of integers as parts of such formulae. In order to verify the assertion, "C is on the right of B", it starts with B and its associated pair (1,2), and then looks for formulae corresponding to the sequence: (1,3), (1,4), (1,5) ... up to some arbitrary number. If the program finds C associated with a pair of integers in the series, then the assertion is true; otherwise, it is false. The series is defined by the procedure representing "on the right of", which freezes in the appropriate values for the incremental parameters of the verification procedure.

It should be clear that the whole of the original theory of spatial inference can be reconstructed in this way, even to the extent of coping with the problems of transitivity. Indeed, many adherents of propositional theories may wish to claim that a propositional theory of spatial inference has here been constructed that counters all the earlier criticisms. They would be wrong, but in a way that is most instructive. The construction of the new propositional theory of spatial inference is in reality simply a reconstruction of the original mental model theory at a lower level of description. The whole of the propositional apparatus, the ordered pairs of integers, the definition of "on the right of" in terms of incremental values of parameters, is parasitic upon the unacknowledged presence of a spatial array. Perhaps it is easiest to grasp this point by asking oneself how such a system could have been set up in the first place, how it could have been learnt, and where the definition of "on the right of" could have come from. The program functions as though it uses an array, and one seen from a particular viewpoint, too.

In general, a model is only a model at a certain level of description: that level at which it functions as one. A listing of the original spatial inference program in machine code is a level of description that obscures the program's use of models. The new 'propositional' theory is similarly a redescription of the old theory at a level that obscures its reliance on models; it is a description that could well pass as a slightly more detailed account of how to set up and manipulate arrays in a certain programming language.

There is, of course, nothing inconsistent about calling such a redescription a propositional theory. It would certainly resolve the controversy in favour of the propositional theorists, but it would do so in an entirely trivial way, for it would entail that any plausible theory of any psychological phenomenon is propositional. The conclusion follows from the fact that any theory which meets the reasonable criterion of being an 'effective procedure' can in turn be described by a set of propositions that determine, from moment to moment, precisely how the system should behave. The notion of proposition is here somewhat vague, and can be replaced by the much more precise notion of a Turing machine (see Minsky, 1967, p.106 et seq). To characterize a theory as propositional would accordingly be to say nothing of any empirical consequence: it would be equivalent to saying that a theory could be instantiated as a Turing machine. It is unclear whether those who advocate a propositional representation for images intend to make so trivial a point; certainly, they have not made a direct appeal to the notion of an effective procedure. What is noteworthy, however, is that they have freely introduced propositions expressing polar coordinates, vectors, and other spatial matters. Such concepts can obviously be expressed in the scientific meta-language, but there is no corresponding terminology for them in the object language -- the language of simple shapes that is under analysis. If the term "propositional" is to have any empirical content, then the following conclusion can be drawn: what a theorist constructs in such cases is, not a propositional theory, but a reconstruction of an imaginal theory at a lower level of description.

The term "propositional" has a perfectly good sense that we have already encountered, and one that makes a sensible contrast with the term, "model". This sense concerns theories of semantics. A characteristic feature of many semantic theories is that the terms of the vocabulary in which they encode propositions are closely related to the vocabulary of natural language. This feature, which may be fortuitous, has been elevated into a theoretical principle by Fodor et al (1975). They propose that "to each morpheme of the surface vocabulary of a natural language there corresponds a primitive expression in the vocabulary of the representational system", and that "meaning postulates mediate whatever entailment relations between sentences turn upon their lexical content". Representations of this sort are truly propositional, and they are distinct from representations based on models, because the vocabulary of a propositional representation does not now allow the use of procedures that depend upon semantic primitives for which there are no readily available terms in the language under analysis. The contrast between the two sorts of representation is important, and the present resolution of the controversy about them establishes the need for both mental models and propositional representations. Human beings have recourse to representations of both sorts.

9. THE FORMS OF MENTAL REPRESENTATION

Three different sorts of representation have been discussed in this paper: images, propositions, and models. Since the lines of demarcation between them may seem to be vague, I want finally to try to draw some finer distinctions.

What is a propositional representation? The answer obviously depends on what a proposition is. One view, which has much to recommend it, is a generalization of the commonplace notion that to understand a proposition is to know what the world would have to be like for it to be true. If we consider all the different ways in which the world might be, as well as the way it actually is, that is, the set of all 'possible worlds', then a proposition is, in principle, either true or else false of each member of the set. Hence, we can treat a proposition as a function from the set of possible worlds onto the set of truth values.[3] A logician might, in turn, treat this function as a set of ordered pairs, each comprising a possible world and a truth value (of the proposition in that world), but this conception is highly abstract since the set of possible worlds is plainly infinite.

If a proposition is a function, then its representation is the representation of a function. The way to represent a function is to express it in a language, and, as Fodor et al (1975) have argued, it is useful to think of a propositional representation as an expression in a mental language. Although we may never delineate the details of the mental language, we do know that it must have both a syntax and a semantics (Fodor, 1976). It must be capable, for example, of representing conjunction, and its mental syntax could take a variety of forms, e.g. "$(\alpha K \beta)$", "$K(\alpha,\beta)$" or "$(\alpha,\beta)K$", where the Greek letters range over representations of propositions, and "K" stands for some mental token representing conjunction. The syntax is free to take any form provided it is associated with the appropriate semantics.

The propositional description of a complicated object may well consist of a large number of propositions. The question arises as to the nature of the structural relations between them. In fact, one paradigm case of a propositional representation is simply an unordered set of expressions in some symbolic language such as the predicate calculus. There are uniform proof procedures that will evaluate inferences made in such a formalism; they rely in part on procedures that will search the set for any particular atomic proposition, examining complex propositions to check whether it is a constituent of them. One might (just) imagine that such a system could pass as a possible psychological theory with the addition of a heuristic for making inferences. However, advocates of propositional theories have almost invariably relied on some sort of semantic network (see Anderson, 1976, 1978; Anderson and Bower, 1973; Baylor, 1971; Kintsch, 1974; Moran, 1973; Norman and Rumelhart, 1975; Palmer, 1975). In a network, propositions about the same entity are gathered together and attached to the single node for that entity. The structure facilitates the processes that encode or retrieve information: it does not have an essential analogical role, since its effect could as well be achieved in unstructured formulae by referential indices that symbolize co-reference.

Mental models have a completely analogical structure. The way in which they contrast with propositional representations is clearest in relation to quantified assertions. Quantifiers, in fact, present a major problem for propositional representations. Even in the guise of semantic networks, propositional representations rely on the methods of formal logic (see Woods, 1975; Hendrix, 1979). One snag with all of them is that they do not readily account for the systematic biases and errors that occur in everyday reasoning. A further, more technical, difficulty is the possibility -- strongly urged by Hintikka (1974) -- that the semantics of natural language quantifiers is very much richer than is often supposed, and outstrips the strictly linear representations made available by standard quantification theory. Mental models solve both of these difficulties. They accommodate numerical and quasi-numerical quantifiers such as "most", "many", and "few". Likewise, they can represent discourse in general and account for the use and interpretation of anaphoric expressions (see Johnson-Laird, 1983).

The reason why mental models are so readily generalizable should be obvious: they assume, not an abstract propositional representation of statements, but one that mirrors the relevant aspects of a real state of affairs that would satisfy the statements. Thus, they use elements to stand for individuals, and links to stand for identities between them. But, they possess one other feature that distinguishes them from propositional representations. They represent a set of entities by introducing an arbitrary number of elements that denote exemplary members of the set. Propositional representations of the sort proposed by Fodor et al (1975) do not contain arbitrary features, whereas models ordinarily do so. Images share this property, too, which has often drawn comment from philosophers. You cannot form an image of a triangle in general, but only of a specific triangle. Hence, if you reason on the basis of a model or image, you must take pains to ensure that your conclusion goes beyond the specific instance you considered (see e.g. Hume 1896, vol. 1). The heuristic advantage of a model is balanced by the need for procedures that test the conclusions that can be derived from it -- a point that is borne out by the way in which the models for quantified assertions and spatial relations have to be manipulated in order to ensure validity. There must be recursive procedures that can revise models in order to correct assumptions that turn out to have been mistaken (see Johnson-Laird, 1983).

Of course models can have a richer analogical structure than those required for quantifiers. They may be two- or three-dimensional; they may be dynamic; they may take on an even higher number of dimensions in the case of certain gifted individuals. One advantage of their dimensional structure is that they can be scanned in any direction, regular or irregular, since the dimensional variables controlling the search can be determined from moment to moment by any mentally computable function. In the case of a propositional representation, as Simon (1972) points out, direct scanning can be performed only in those directions that have been encoded in the representation.

Models and propositions are interesting to compare on the criterion of economy. On the one hand, if a series of assertions are highly indeterminate, and one is not required to draw any profound inferences from them, it may be more economical to remember the propositions that were asserted rather than to interpret them in the form of a model. There is certainly a limit to the extent that human beings can manipulate models in order to ensure validity, and even certain syllogisms are taxing for this reason. The most plausible theory of comprehension accordingly calls for two stages: firstly, the utterance is translated into a propositional representation, which takes the form of a superficial linguistic representation that encodes verbatim information; secondly, and optionally, the propositional representation is used as the input to a procedural semantics which uses it, together with inferences from general knowledge, to construct a mental model. At the heart of the theory lies the following idea: mental models represent the extensions of assertions, i.e. the situations they describe, and the superficial linguistic representations, together with the recursive machinery for constructing and revising models, represent the intensions of assertions, i.e. the sets of all possible situations that they could describe. This theory and the evidence that supports it is described in detail in Johnson-Laird (1983).

If the possession of arbitrary components and an analogical structure distinguishes models from propositions, then what is it that distinguishes models from images? It seems likely that models are the basis of images, which simply correspond to those features of them that are directly perceptible in the equivalent real-world objects. An image may simply be a projection of a three-dimensional mental model onto an internalized two-dimensional surface, thus giving rise to a 'view' of the model from a particular standpoint. Conversely, models may underlie thought processes without necessarily emerging into consciousness in the form of images. Models are also likely to underlie the perception of objects by providing prototypical information about them (see Roberts, 1965) in a form that can be directly used in the interpretation of what Marr (1976) has referred to as 'the primal sketch', the output of lower level visual processes.

10. CONCLUSIONS

To return to the two controversies with which I began, if the concept of a proposition has any empirical force, then there are indeed distinctions to be drawn between a set of propositions and a model or its perceptible counterpart, an image:

1. A model represents an object or a set of assertions analogically: its structure is constrained and is a crucial part of the representation. A propositional representation describes an object and corresponds to the linguistic structure of a set of assertions.

2. A model based on a linguistic description embodies a number of arbitrary assumptions since language is inherently vague but models must be determinate. A propositional representation does not contain any arbitrary elements. The two sorts of representation accordingly do not have the same equivalence classes, and hence there is no guarantee that a theory based on one sort of representation can be made to mimic one based the other sort.

3. A model represented in a dimensional space can be directly constructed, manipulated, or scanned, in any way that can be controlled by dimensional variables. A propositional representation lacks this flexibility and can be directly scanned only in those directions that have been laid down between the elements of the representation.

There are likewise distinctions to be drawn between semantic primitives and meaning postulates:

1. Semantic primitives are required by the semantic system that establishes a word's relation to the world, or strictly speaking to models of the world. Meaning postulates are neither intended to perform this function nor contain any machinery for defining truth conditions.

2. The logical properties of a term need not be specified within the definitions of its truth conditions, rather they are emergent properties of that definition. Meaning postulates, however, are rules that explicitly specify the logical properties of terms, and the logical relations between them.

3. Semantic primitives are ineffable in that there are no corresponding lexical items for them in the language under analysis. Meaning postulates in their psychological guise require a mental language that can be mapped virtually one-to-one onto the natural language.

What I have argued in this paper is that comprehension at its deepest level calls for an initial propositional representation and the construction of a mental model, which may take the form of an image of the state of affairs described in the discourse. The process of constructing a mental model depends on a representation of meaning that makes use of semantic primitives. Such a theory, unlike one based solely on propositional representations and meaning postulates, can cope with the logical vagaries of natural language. It can also account for performance in a variety of inferential tasks. It does not follow that all interpretations consist of mental models: no case has been presented against the use of other forms of representation or rules of inference in other contexts. However, if human beings habitually used only propositional representations and meaning postulates, vagaries in the logical properties of language could never have arisen in the first place.

REFERENCES

Anderson, J.R., Language, memory and thought. (Hillsdale, N.J., Erlbaum, 1976).
Anderson, J.R., Arguments concerning representations for mental imagery. Psychological Review, 85, (1978) 249-277.
Anderson, J.R., and Bower, G.H., Human associative memory. (New York, V.H. Winston & Sons, 1973).
Baylor, G.W., Programs and protocol analysis on a mental imagery task. First International Joint Conference on Artificial Intelligence, 1971.
Bower, G.H., Mental imagery and associative learning. In L. Gregg (ed.), Cognition in learning and memory. (New York, Wiley, 1972).
Brooks, L., The suppression of visualization by reading. Quarterly Journal of Experimental Psychology, 19, (1967) 280-299.
Bugelski, B.R., Words and things and images. American Psychologist, 25, (1970), 1002-1012.
Byrne, B., Item concreteness vs. spatial organization. Memory and Cognition, 2, (1974), 53-59.
Carnap, R., Meaning and necessity: a study in semantics and modal logic. Second edition. (Chicago, University of Chicago Press, 1956).
Chase, W.G. and Clark, H.H., Mental operations in the comparison of sentences and pictures. In L.W.Gregg (ed.), Cognition in learning and memory. (New York, Wiley, 1972).
Clark, H.H. and Clark, E.V., Psychology and language: an introduction to psycholinguistics. (New York, Harcourt Brace Jovanovich, 1977).
Collins, A.M. and Quillian, M.R., Experiments on semantic memory and language comprehension. In L.W. Gregg (ed.), Cognition in learning and memory. (New York, Wiley, 1972).
Cooper, L.A., Mental rotation of random two-dimensional shapes. Cognitive Psychology, 7, (1975), 20-43.
Davies, D.J.M. and Isard, S.D., Utterances as programs. In D. Michie (ed.), Machine Intelligence, 7. (Edinburgh, Edinburgh University Press, 1972).
Dennett, D.C., Content and consciousness. (New York, Humanities Press, 1969).
Fodor, J.A., The language of thought. (Hassocks, Sussex, Harvester Press, 1976).
Fodor, J.D., Semantics: theories of meaning in generative grammar. (Hassocks, Sussex, Harvester Press, 1977).
Fodor, J.D., Fodor, J.A., and Garrett, M.F., The psychological unreality of semantic representations. Linguistic Inquiry, 4, (1975), 515-531.
Funt, B.V., WHISPER: a problem-solving system utilizing diagrams. Fifth International Joint Conference on Artificial Intelligence, (1977), 459-464.
Galton, F., Inquiries into human faculty and its development. (London, Dent, 1928, originally published 1880).
Gelernter, H., Realization of a geometry-theorem proving machine. In E.A. Feigenbaum and J. Feldman (ed.), Computers and thought. (New York, McGraw-Hill, 1963).
Hayes, J.R. On the function of visual imagery in elementary mathematics. In W.G. Chase (ed.), Visual information processing. New York: Academic Press, 1973.

Hendrix, G.G., Encoding knowledge in partitioned networks. In N.V. Findler (ed.) Associative networks: Representation and use of knowledge by computers. (New York, Holt, Rinehart, and Winston 1979).
Hintikka, J., Quantifiers vs. quantification theory. Linguistic Inquiry, 5, (1974), 153-177.
Holyoak, K.J., The form of analog size information in memory. Cognitive Psychology, 9, (1979), 31-51.
Hume, D., A treatise of human nature. Vol. I. Edited by L.A. Selby-Bigge. (Oxford, Clarendon, 1896).
Inhelder, B. and Piaget, J., The growth of logical thinking from childhood to adolescence: an essay on the construction of formal operational structure. (London: Routledge and Kegan Paul 1958).
Inhelder, B. and Piaget, J., The early growth of logic in the child: classification and seriation. (London, Routledge and Kegan Paul, 1964).
Johnson-Laird, P.N., Procedural semantics. Cognition, 5, (1977), 189-214.
Johnson-Laird, P.N., Mental Models: Towards a Cognitive Science of Language, Inference and Consciousness. (Cambridge, Cambridge University Press. Cambridge, Mass., Harvard University Press, 1983).
Johnson-Laird, P.N. and Bara, B., Syllogistic inference. Mimeo, MRC Applied Psychology Unit, 1983.
Johnson-Laird, P.N. and Steedman, M.J., The psychology of syllogisms. Cognitive Psychology, 10, (1978), 64-99.
Johnson-Laird, P.N., Gibbs, G. and de Mowbray, J. Meaning, amount of processing, and memory for words. Memory and Cognition, 6, (1978), 372-375.
Johnson-Laird, P.N., and Herrmann, D.J. Only connections: a critique of semantic networks. Mimeo, MRC Applied Psychology Unit, 1983.
Katz, J.J. and Fodor, J.A., The structure of a semantic theory. Language, 39, (1963), 170-210.
Katz, J.J. and Nagel, R., Meaning postulates and semantic theory. Foundations of Language, 2, (1974) 311-340.
Kieras, D., Beyond pictures and words: alternative information-processing models for imagery effects in verbal memory. Psychological Bulletin, 85, (1978), 532-554.
Kintsch, W., The representation of meaning in memory. (Hillsdale, N.J., Erlbaum, 1974).
Kosslyn, S.M., Information representation in visual images. Cognitive Psychology, 7, (1975), 341-370.
Kosslyn, S.M., Images and Mind. (Cambridge, Mass., Harvard University Press, 1980).
Kosslyn, S.M. and Pomerantz, J.R., Imagery, propositions, and the form of internal representations. Cognitive Psychology, 9, (1977), 52-76.
Lewis, D., General semantics. In D. Davidson, and G. Harman (eds.), Semantics of natural language. (Dordrecht, Reidel, 1972).
Lyons, J. Semantics, Vols. 1 and 2. (Cambridge, Cambridge University Press, 1977).
Mani, K., and Johnson-Laird, The mental representation of spatial descriptions. Memory and Cognition, 10, (1982), 181-187.
Marr, D. Early processing of visual information. Philosophical Transactions of the Royal Society of London, Series B, 275, (1976), 483-519.
Miller, G.A. and Chomsky, N., Finitary models of language users. In Luce, R.D., Bush, R.R. and Galanter, E. (eds.), Handbook of mathematical psychology, vol. II., (New York, Wiley, 1963).
Miller, G.A. and Johnson-Laird, P.N., Language and Perception. (Cambridge, Mass., Harvard University Press; Cambrige, Cambridge University Press, 1976).
Minsky, M. Computation: finite and infinite machines. (Englewood Cliffs, N.J., Prentice-Hall, 1967).
Moran, T.P., The symbolic nature of visual imagery. Third International Joint Conference on Artificial Intelligence, (1973), 472-477.
Moyer, R.S., Comparing objects in memory: evidence suggesting an internal psychophysics. Perception and Psychophysics, 13, (1973), 180-184.

Norman, D.A. and Rumelhart, D.E., Memory and knowledge. In D.A. Norman, D.E. Rumelhart, and the LNR Research Group, Explorations in cognition. (San Francisco, Freeman, 1975).
Paivio, A., Imagery and verbal processes. (New York, Holt, Rinehart and Winston, 1971).
Paivio, A., Perceptual comparisons through the mind's eye. Memory and Cognition, 3, (1975), 635-647.
Paivio, A., Images, propositions, and knowledge. In J.M. Nicholas (ed.), Images, perception and knowledge. (Dordrecht, Holland, Reidel, 1977).
Palmer, S.E., Visual perception and world knowledge: notes on a model of sensory-cognitive interaction. In D.A. Norman, D.E. Rumelhart, and the LNR Research Group, Explorations in cognition. (San Francisco, Freeman, 1975).
Pylyshyn, Z.W., What the mind's eye tells the mind's brain: a critique of mental imagery. Psychological Bulletin, 80, (1973), 1-24.
Roberts, L.G., Machine perception of three-dimensional solids. In I.J.T. Tippett et al (eds.), Optical and electro-optical information processing. (Cambridge, Mass., M.I.T. Press, 1965.
Schank, R.C., Conceptual information processing. (Amsterdam, North-Holland, 1975).
Scott, D. and Strachey, C., Toward a mathematical semantics for computer languages. Proceedings of the Symposium on Computers and Automata, Polytechnic Institute of Brooklyn, April, 1971.
Shepard, R.N., Form, formation, and transformation of internal representations. In R. Solso (ed.), Information processing and cognition, The Loyola Symposium. Hillsdale, N.J., Erlbaum, 1975.
Shepard, R.N., The mental image. American Psychologist, 33, (1978), 125-137.
Simon, H.A., What is visual imagery? An information processing interpretation. In L.W. Gregg (ed.), Cognition in learning and memory. (New York, Wiley, 1972).
Simon, H.A. and Barenfeld, M., Information-processing analysis of perceptual processes in problem solving. Psychological Review, 76, (1969), 473-483.
Sloman, A., Interactions between philosophy and artificial intelligence: the role of intuition and non-logical reasoning in intelligence. Artificial Intelligence, 2, (1971), 209-225.
Smith, E.E., Shoben, E.J. and Rips, L.J., Structure and process in semantic memory: a featural model for semantic decisions. Psychological Review, 81, (1974), 214-241.
Woods, W.A., What's in a link: foundations for semantic networks. In D.G. Bobrow and A. Collins (eds.), Representation and understanding: Studies in cognitive science. (New York, Academic Press, 1975).
Woods, W.A., Procedural semantics. In A.K. Joshi, I. Sag and B.L. Webber (eds.), Elements of Discourse Understanding. (Cambridge, Cambridge University Press, 1981).

FOOTNOTES

1. There is a danger of an infinite regress here. If an interlingua is needed to mediate between words and pictures, then perhaps a language is needed to mediate between words and the interlingua, or between the interlingua and pictures, and so on and on (see Anderson, 1978).

2. Expressions such as "in front of", in fact, have two distinct spatial senses, a deictic sense that depends on the speaker's point of view, e.g. "Stand in front of the rock", and another sense that depends on the intrinsic parts of certain sorts of object, e.g. "The river was in front of the house" (see Fillmore, 1971; Miller and Johnson-Laird, 1976, Sec. 6.1.3). This complication is not relevant to the present argument, and I have otherwise ignored it.

3. Sentences in isolation rarely express single propositions. It is the sentence and its context that convey a proposition (see e.g. Lewis, 1972).

ACKNOWLEDGEMENTS

This paper is an extensively revised version of one that originally appeared in Italian in Ricerche di Psicologia in a special issue on language, 1983, edited by Bruno Bara. I am grateful to Bruno Bara, Kate Ehrlich, Alan Garnham, Steve Isard, Christopher Longuet-Higgins, Stuart Sutherland, and Eric Wanner for helpful criticisms of the earlier paper.