DRAFT

December 5, 1978

LABORATORY FOR COMPUTER SCIENCE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Metaphysical and Epistemological Foundations for a Linguistically Oriented Semantic Network

by

William A. Martin

> Finally going one step higher, we might consider networks whose primitive elements are language-specific. The only formalism that I know of at the current time that embodies this view is OWL, whose elements are expressions based on English. In such a formalism, one would presumably "...take seriously the Worfian hypothesis that a person's language plays a key role in determining his model of the world and thus in structuring his thought" [Martin 1977, p. 985]. In OWL, there is a basic concept-structuring scheme (see [Hawkinson 1975]) which is used to build expressions, and strictly speaking, the principles of "specialization", "attachment", and "reference" are the primitives of the language. However, these primitives are neutral enough to be considered implementational, and thus the knowledge itself can be considered to form the structure of the data base. This seems operationally reasonable when OWL is looked at in detail -- the two expressions, (HYDRANT FIRE) and (MAN FIRE), while both specialized by FIRE, can have the specializations "mean" different things based on the rest of the network structure. This linguistic level represents perhaps the most radical view of semantic nets, in that the "primitives" are language-dependent, and are expected to change in meaning as the network grows. Links in linguistic level networks stand for arbitrary relationships that exist in the world being represented.

As Brachman suggests, the significance of an OWL concept derives from the rest of the network, but this does not preclude the establishment of logical and epistemological conventions; and an orientation toward language can influence choices made at the logical and epistemological levels. In what follows some conventions of OWL at these levels will be set forth. This done, some comparison of OWL with other systems including those which Brachman has listed at the logical and epistemological levels can be ventured.

There are two ways in which language and metaphysics can interact. On the one hand we can try to arrive analytically at a set of principles which can be used to elucidate such linguistic issues as how to distinguish different senses of a word. This is working bottom up. An example of this can be had in the tract by Marcus (1960). He comes to a related but different interpretation of some of the phenomena discussed here. The second possibility is to work top down. To use intuitions about language use as a basis for assumptions at the metaphysical level. The reader will find both sorts of argument in this paper. There is a sharp contrast to, for example, Montague (1969), where the emphasis is on making the minimum number of assumptions about the entities present in the world which will enable the construction of a system of semantics. For example, the question of existence is a complex one, but certainly there are senses of "exist" for which an entity doesn't exist just because some speakers appear to refer to it as if it did. A "scientific" course for logicians and philosophers, then, is to determine the minimum assumptions which must be made, and to discover properties of the resulting systems which may be used as a norm to guide and evaluate theories of what people might actually think or do. Our goal has been one of engineering rather than science, to establish a conceptual base for semantic networks which is principled and plausible, and in which the practical problems of representation of "linguistic" knowledge will "work out". As such, it draws on, but does not follow, analytic philosophy. Hopefully this paper will give the reader further insight into the issues which one must face in doing this and the comparison of our work with other systems will help him evaluate the choices made here.

2. Notation

English words and phrases referred to in text will be quoted. Entities in our ontology of the world will be written in italics. Concepts will be written using all capital letters. For example, corresponding to the word "dog" we might have the concept, DOG, with the referent, *dog*.

Concepts will correspond to nodes in a semantic network. The links between nodes will be referred to as zone relations. Zone relations will be written in capitals like concepts, but also prefixed by #, e.g. #CHARACTERIZATION. The notation

```
[PROFESSOR #CHARACTERIZATION  FACULTY-MEMBER
           #EXEMPLAR  ASSISTANT-PROFESSOR
                      ASSOCIATE-PROFESSOR]
```

will be called a complex. The concept following the left bracket is termed the subject of the complex. The above complex states that its subject, PROFESSOR, has a #CHARACTERIZATION link to FACULTY-MEMBER and #EXEMPLAR links to ASSISTANT-PROFESSOR and ASSOCIATE-PROFESSOR.

3. Ontology

Our first step will be to establish a classification of things which we take to be present in the world. While distinctions of the type to be made here have been made since the time of Aristotle [Ackvill, 1963], philosophers have found it quite difficult to establish any one scheme by analytical means. Our goal is to devise a scheme which leads to a semantic network in which it is easy to express knowledge needed for understanding English. Our approach will be to take decisions which express our intuitive notions of how English speakers think about the world when expressing their thoughts in English. When our intuitions about English fail us, we will look for decisions which produce good language processing strategies.

3.1 Kinds and Individuals

We begin our classification by dividing the things in the world into kinds and individuals. We state that kinds exist in the world because it is our intuition that English speakers make this assumption. For example, we claim that in 1

1) The elephant plays an economic role in some countries.

the speaker is referring to elephant kind, and he is doing so as if elephant kind exists as an entity in the world. The speaker believes in individual elephants, sets of elephants, and elephant kind, which is something different than the set of all elephants which exist at any one time or have existed at all times. The existence of kinds like elephant kind is not universally accepted. For example, Lyons (1977) develops an ontology which admits the existence of the controversial entities, propositions, but not the existence of kinds.

In addition to kinds, we also think of the world as containing particular things, as Strawson (1963) says:

> "We think of the world as containing particular things some of which are independent of ourselves; we think of the world's history as made up of

> particular episodes in which we may or may not have a part; and we think of these particular things and events as included in the topics of our common discourse, as things about which we can talk tc each other. These are remarks about the way we think about the world, about our conceptual scheme. A more recognizably philosophical, though no clearer, way of expressing them would be to say that our ontology comprises objective particulars."

As examples of objective particulars Strawson lists historical occurrances, material objects, people, and their shadows.

The idea of individuals seems quite natural, but there are circumstances which show more clearly how the existence of individuals must be viewed as an assumption people make. Consider, for example, a world consisting of nothing but two spheres rotating around their common center of mass, Black (). Since the two spheres are indecernable, can it make any sense to talk about separate individuals? We are prepared to do it; we look at the problem as one of developing a test which will let the two individuals be distinguished, Kripke () has argued that it is wrong to think of names as shorthand for definite descriptions because we really believe that the name names an individual independent of whether a description can pick it out.

As another example, consider an individual place, like one's favorite spot in a forest. What grounds have we for considering this an individual. Couldn't we also have another place just slightly shifted from this, thought to include or leave out an additional tree. The forest has as many individual places as we wish to conceive of. There is nothing innate about the forest which determines which individual places we will choose to identify. We take a substructure of a structure like a forest to be an individual whenever it is an aid to thought.

Expanding upward from this example, there is a sense in which the entire world is a structure. Important substructures of this structure are thought of as individuals.

3.2 Predicates and Descriptions

Corresponding to each kind we will postulate a predicate on individuals which determines if the individual is of that kind. Montague () calls such a one place predicate a property. He defines a property in terms of a function which enumerates the individuals having the property:

> "To be more specific, let I be the set of all *possible worlds*; for each member i of I, let A_i be the set of *individuals existing in the possible world* i; and let U be the set of all *possible individuals*. Then U will include the union of all sets A_i for i in I, and may indeed coincide with it, though we need not explicitly impose such a limitation. A *property of individuals* is a function having I as its domain and subsets of U as its values. (If P is such a property and i a possible world, $P(i)$ is regarded as the set of possible individuals that partake of P in i. For example, the property of being red is the function that assigns to each possible world the set of possible individuals which in that world are red.) More generally, an *n-place predicate of individuals* is a function having I as its domain and of which the values are sets of n-place sequences of members of U. If P is a predicate of individuals and i is a possible world, we regard $P(i)$ as the *extension* of P in i; the extension of an n-place predicate will always be an n-place relation (in the extensional sense)."

Montague stipulates that two properties are identical just in case they have the same extension in all possible worlds. Such a criterion is useful for thought experiments.

However, for semantic nets what we need is a representation of the property which will allow us to do intensional reasoning. Therefore, we will associate with each property a description telling things about any individual having the property. In doing this we part ways to a certain extent with A/I researchers advocating procedural semantics. These people would represent the property with procedures which could be executed to enumerate or identify individuals having the property. The procedures must be written in terms which some processor can interpret. By descriptions we intend a more non-procedural representation. It would require a more powerful processor to interprete these descriptions and thus identify or enumerate individuals to which they apply.

We believe that there are individuals in the world and that our knowledge of the world consists of knowing

a) What individuals there are which have a given property - my wife exists but no unicorns exist.

b) For two properties A and B whether individuals which have A also have B - a dog is an animal, a dog might be a pet.

A predication like "dogs are animals" we take to mean that any individuals which have the description "dog" have the description "animal" as well. This predication is a description of a state and is true if that state exists in the world. Thus truth has to do with the correspondence between symbolic descriptions and the world.

When we describe an individual as an elephant we will refer to him as a particular elephant. When we describe as a worker, we will refer to him as a particular worker. Clearly the mapping between kinds and individuals is many to many.

Recall that above Montague introduced not just one place predicates, but n-place predicates. An n-place predicate is a function whose values are sets of n-place sequences of individuals. He gives as an example the two place predicate "lifting a stone" which has values which are individuals and stones, such as that that individual lifted that stone.

By contrast, we will use only one place predicates. Instead of n-place predicates we postulate the existance of individual conditions and events which have structure. For example, there are individual *liftings* which have individual *stones* and individual *lifters* in their structure. When we say "Bob lifted a stone" we predicate of the individual, Bob, that he may be described as a particular *lifter* of a stone.

This approach is amplified in the following discussion.

3.3 Context and Structure

Any entity, kind or individual, may have a context and a structure. That is, we may say of a second entity whether or not it is part of the context or structure of a given entity. Use of these two notions as primitives may be found in Fahlman (1978). Here, context will be taken as a derivative notion. An entity, A, will be taken to be part of the context of an entity, B, only if either B is part of the structure of A or A and B are both part of the structure of some entity C.

3.4 Classification of Kinds

Having classified entities into kinds and individuals, our next step is to classify kinds into five types:

1) Attributes / Physical-Disturbances
2) States / Processes
3) Relational Characterizations
4) Attributive Characterizations
5) Other kinds

This classification will be based on three criteria; a) that of one entity being ontologically prior to another, b) the notion of existing only for a period or moment of time, and c) a criterion regarding the requirements for the identification of individuals of a kind.

The notion of one kind being ontologically prior to another may be illustrated with the kinds *physical body* and *weight.* Since one can conceive of physical bodies without weight, but not the reverse, physical bodies are said to be ontologically prior to weight.

The issues of identification are more subtle and require a longer introduction.

0 X

\+ 0

Figure 1

Looking at Figure 1, one may say that the figure about the + is the same as the figure to the left of the X. Using another sense of same, one may say that the figure in the upper left corner is the same as the figure in the lower right corner. In the first case we have numerical identity, in the second case we have only structural identity. In LISP terms, we could say that in the first case the figures are EQ, which in the second they are EQUAL. To identify an individual we must establish numerical, or EQ identity.

Suppose that a person is blindfolded and moved about so that he loses track of his whereabouts. When the blindfold is removed, he discovers that he is in his own bedroom. He believes the room is numerically the same. He could be in a carbon copy of that room, but he firmly believes that he is not. He believes in the existence of individual lamps, chairs, etc. and the continuity of their existence in space and time. Seeing the familiar configuration of these familiar individuals he firmly believes that he has uniquely established his whereabouts.

Now if the same person were to come upon a piece of his furniture in an unfamiliar location, he would probably be reluctant to identify it as numerically the same. His memory of the description of the piece is not enough, he needs the bedroom context to be sure. Unblindfolded, he must verify the numerical identity of the room before he can numerically identify its constituents. On the other hand, since the room is composed of its constituents, one might be lead to assume that our subject could not identify the room without first so identifying its constituents. Neither the room nor the constituents come first. After all, he might be in a carbon copy. We sense, however, that this is not right. One never feels oneself to be in such a dilemma.

The obvious answer is that to verify the room is numerically the same people, in fact, only require that pieces of the same structural description are at their expected positions in the room and that there is no strong negative evidence. They consider the description of their room sufficiently unique to establish its identity. They are not concerned with whether the pieces of furniture in the room are numerically the same as their own. Once the room is verified to be numerically the same, the pieces can be also so identified if that is desired.

Underlying the previous discussion is the assumption that corresponding to the structure and context of an entity we may have structural descriptions of the entity and descriptions of the context of the entity. Further, the structural description of an entity, such as a lamp, may be used either:

1) As part of the structural description of another entity, e.g. a room. In this case the description may be referred to as a <u>role</u>. We may find that an individual fills a role without establishing the numerical identity of the role filling individual.

or 2) to establish, in concert with a numerically unique context, the numerical identity of an individual.

Using these distinctions, kinds may be classified into the five categories given above:

1) <u>Attributes / Physical-Disturbances</u>

Examples are attributes such as ***health, color,*** **and** ***length.*** **Note that one can refer to** ***health*** **as a kind, as in "health is important to happiness". He can also identify a particular health, but only by referring it to the entity for which it is a part of the structural description, as in "What is the state of Bob's health". He can't say "some health" or "a health". One can describe the color of a particular chair as red, but the description "red" does not numerically identify that individual the way describing it as the color of a particular chair does.**

Based on these considerations, attributes/physical-distrubances will be defined as kinds whose corresponding particulars are individuals which can only be numerically identified through their "dependence" on one or more other numerically identified entities. By dependence is meant either that they fill a role of the sponsoring particulars' structural description, or that they have the sponsoring particulars in their context. Attributes such as ***height, weight,*** **and** ***health*** **have corresponding particulars which fill roles. Physical disturbances like** ***weather, flashes,*** **and** ***bangs*** **have corresponding particulars identified by context.**

Some kinds in this category may have either relation to their sponsoring entity. For example, we may say either "the smell in a room" or "the smell of a room", "temperature in a room" or "the temperature of a room". As a subclass of attributes we may distinguish scales, such as ***height, weight,*** **and** ***size.*** **A scale may be modeled by a ray of a line running from zero to infinity. Scales support the notion of magnitude**

comparison. The method of making the comparison will vary with the scale. In the case of a complex attribute like *size*, for example, the method will depend on the object having that attribute. Another important subclass of attributes may be termed intervals. Examples are *degree of fullness* and *extent to which one is grown up*. Intervals may be modeled by a line segment running from zero to one.

2) States / Processes

Examples are states and processes involving physical bodies such as *hit* and *hold*. States/processes will be defined as kinds whose corresponding particulars are individuals which, as with attributes/physical-disturbances, can only be numerically identified through their dependence on one or more numerically identified sponsoring entities. However, in this case dependence means that the sponsoring entity either fills a role in the sponsored particulars structural description of serves as a context. For example, the particular event described by "Bob kissed Mary by the barn yesterday," has "Bob" specifying the role of subject, "Mary" specifying the role of object, and "the barn" and "yesterday" specifying a context. It would be difficult to argue that no condition or event can be identified by its structural description alone. Our intuition is, however, that English speakers don't do this. (Even if a speaker says "remember the time that that guy kissed that girl on the elbow" he is implying the hearer knows about the incident and is thus referring it to the hearer.)

Note that according to what has been said so far the distance between, for example, two people Tom and Fred could be considered either an attribute in the structure of the "set" of Tom and Fred or a state with Tom and Fred as individualizing particulars in its structure. The second criterion listed above, time dependence, is used to eliminate this ambiguity. Since the existence of a distance between two objects is not a function of time, distance is taken as an attribute rather than a state. In general, an attribute/physical-property does not involve the notion of time, while a state/process does.

The *rising of the sun* we take to be a process identified by having the sun in its structure and a certain *place* and *moment of time* in its context. Since we have expressed context in terms of structure, this means that there is some entity which has the *rising of the sun*, the *place*, and the *moment of time* in its structure, or else that entity is the *place* or the *moment in time*.

Montague (1969) treats this example by making *the rising of the sun* a property of the *moment in time*. While this is not really foreign to our treatment, we must consider it as insufficient. The moment in time is not itself enough to identify an event.

Since Montague is not interested in precisely how an event can be distinguished from other properties, he is not concerned with how events could be identified. Recall that he postulates a function, the intension, for each predicate which would take a possible world as argument and produce a function which could select the tuples in that possible world for which the predicate holds at a given time. But he says nothing about how such a function could be constructed. The function is a mathematical fiction which allows him to ignore this problem.

Since we are interested in representing the details of events in a semantic network, we are motivated to give events the status of entities with a structure (and attributes) of their own. This leads us to give relatively non-procedural descriptions of events. In an alternative approach, Hobbs and Rosenstein () have suggested how functions similar to Montague's intension functions could be realized in LISP. This gives them a procedural representation of events. We return to this contrast in Section 8.

States and processes may be subdivided as shown below. States and processes are described using verbs, nouns, adjectives, and adverbs and prepositions.

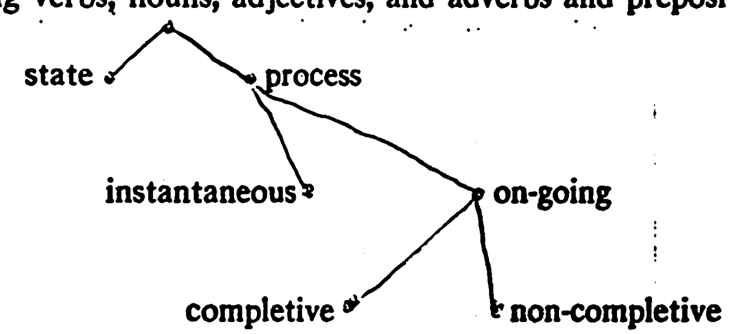

States and Processes

States may be described by what Joos (1968 p. 116) calls "status verbs". He claims there are two main groups:

"(1) psychic state, including the specific perceptions (see, hear, etc.) and the intellectual and emotional attitudes (believe, understand, hate, like, regard,

> etc.); (2) relation, such as the relations of representing, depending, excluding, and so on; it seems interesting that 'have' and the copula 'be' are status verbs of this sort, so that possession in English is a status: 'He has a farm' and 'He is a farmer' are usually synonymous."

It is frequently observed that status verbs repell the progressive, e.g., "Bob is resembling his father more each day." Another property of status verbs is that they cannot have future reference without a specific time shifter such as "will". With a process verb we can say "Don't worry: he leaves next week"; but we can't say, "don't worry: the baby resembles his father next year."

Joos points out that most status verbs can be used both as status verbs and as process verbs, but they have only one sense as a status verb.

> "Take 'hear' as a typical verb-base used both ways. As a process verb it has a number of fairly distinct meanings, as in 'The Judge is hearing a case just now; you'll have to wait;' 'I hope to be hearing from him soon', and as many others as we care to distinguish. But as a status verb it seems to have exactly one meaning."

What is the nature of a state? Joos says:

> "For relation verbs it is obvious, and for psychic-state verbs it is now not too hard to see, that the (single!) meaning of each status verb is such as to reject the time limited validity of temporary aspect. 'That makes no difference' is not a process, not an event that essentially preceeds but is now frozen for our inspection: it is instead a relation between 'that' (whatever it is) and the whole world we live in: it doesn't happen, but simply is so. Equally truly, 'I drop the tablet into this warm water, and you see it disolves quite nicely' is not a process of seeing; it doesn't mean seeing is proceeding within time so that it could be made temporary by using the temporary aspect. Instead, this perception, a kind of psychic state if you like, is a sort of relation between the beholder and the disolving tablet."

While status verbs have a sense which always describes states and not processes, other verb senses can be used to describe either states or processes. For example, "The pig ate a sandwich" describes a process, while "pigs eat corn" describes a state. "Dinosaurs ate kelp" can describe either a state or a process. The distinction between states and processes will be further clarified in Section 9.

A process can be either <u>instantaneous</u>, e.g., "hit", or <u>on-going</u>, e.g., "crying" or "building a house". "Building a house" is <u>completive</u>, but "crying" isn't. One can say "how long did it take you to build the house", but no "how long did it take you to cry" (disregarding the interpretation "to start to cry"). Also, if one has been crying then he has cried, but if one has been building a house, he has not built it. On these distinctions see, for example, (Evans).

Note that while "hit a tree" describes an instantaneous process, "hitting trees" can describe an on-going process which consists of repeated hits. Note also that while "walking" is a non-completive process, "walking to the store" is a completive process.

Adjectives can be sorted into those which normally describe states and those which normally describe non-completive processes. The former include "long" and "male", and resist the form:

I found them _____ .

which the latter, e.g. "sick", "true", accept.

Prepositional phrases, e.g. "in the woods" normally describe non-completive processes.

3) Relational Characterizations

Examples are *father*, *employee* and *neighbor*. These kinds have corresponding particulars which are individuals which can be numerically identified without any sponsoring entities. However, such a kind must have an ontologically prior state or process. For example, to be a natural, legal or ersatz father requires in each case that a certain state or process exist. A husband may be identified by saying who is the husband of, or what marriage he is involved in. I may identify someone as a husband without identifying his wife. But since he may have been married more than once, this is not numerical identification of this particular. To get this I must specify the wife. In general, the numerical identification of a particular of a kind in this category requires at least the identification of a sponsoring entity which is either an entity of its ontological prior kind or an entity filling a role in the structural description of this ontologically prior kind.

One must make a rather fine distinction at the boundary between relational characterizations and attributes. An example of the distinction we have in mind was noticed by Fillmore (1970). He shows that "hit" takes a location and "break" an object. So,

> "27) I broke the top of the table.
> 28) I hit the top of the table.
> In (27) the noun top must be referring to the top of a table as a more or less distinct object, while in (28), it can refer either to that or a portion of the surface area of the table."

The sense of "top" in (27) is a relational characterization. In (28) we have either this sense or the sense of "top" as an attribute of the table.

4) Attributive Characterizations

Examples of these kinds are *bachelor* and *teenager*. These kinds have corresponding particulars which are individuals which can be identified without any sponsoring entities. However, such a kind must have an ontologically prior attribute. For example, there can be no notion of *teenager* without a notion of *age*.

5) Other Kinds

Particulars of these kinds may in principle be identified without reference to any other entities.

4. Descriptions

We now endeavor to set down rules for the formation of descriptions of entities. The rules will be based on the ontology of entities just presented.

4.1 We Require Identifying Descriptions of Individuals to Describe Their Context

The first problem we take up in dealing with descriptions manefests itself in an issue discussed by Moore (1922)

> "How can a thing 'appear' or be 'thought of' unless it is there to appear or be thought of? To say that it appears or is thought of, and yet that there is no such thing, is plainly self-contradictory. A thing cannot have a property unless it is there to have it... When I think of a unicorn, what I am thinking of is certainly not nothing; if it were nothing then, when I think of a griffin, I should be thinking of nothing and there would be no difference

> between thinking of a griffin and thinking of a unicorn. But there certainly is a difference; and what can the difference be except that in one case I am thinking of a unicorn, and in the other a griffin? And if a unicorn is what I am thinking of then certainly there must be a unicorn, in spite of the fact that unicorns are unreal. In other words, though in one sense of the word there certainly are no unicorns - the sense, namely, which to assert that there are would be equivalent to asserting that unicorns are real - yet there must be some other sense in which there are such things; since, if there were not, we could not think of them.

One view, and the one to be adopted here, is that things like unicorns, subsist rather than exist. As Eaton (1925) says:

> Subsistence is a kind of being to which existence may be added, but to which existence is not necessary. In the realm of subsistence, the lion and the unicorn lie down together. To think of a lion or of a unicorn is to think of entities that partake equally of this impartial being, this thinner reality, which includes the possible and the imaginary as well as the real."

Eaton goes on to suggest that subsistence can be interpreted to mean that we have procedures which we can and do apply to check for existence. That we can give a procedural interpretation to the notion of subsistence.

The importance of this issue in semantic networks arises from the simplicity which would derive from equating the existence of a description in a semantic network with the existence of the entity it describes. This possibility is ruled out by the need to represent hypotheticals like "a blue yo-yo" in the non-specific reading of "John wants a blue yo-yo". The representation of this statement will contain the description corresponding to "a blue yo-yo", but it would be wrong to conclude from this statement that a blue yo-yo exists.

On the other hand, if we require that an identifying description of any individual include a description of the individual's context, then we may equate the existence of such a description with the subsistence of the individual in that context. An individual is said to exist when it subsists in the real world.

Recall that in 2.2 context was related to structure, so that to specify the context of an individual, A, means to relate it to another entity, B, whose structure contains A, or to relate it to entities which are part of the structure of some individual, C, of which A is also a part. The effect of this scheme, then, would be to make every individual part of the structure of another entity. To this we make the exception that the real world is not a part of another entity and can be described without context.

Re-examining "John wants a blue yo-yo" in this light, we would describe *a blue yo-yo* as being part of the structure of the *want* entity. Understanding the nature of the subsistance of the blue yo-yo then becomes a matter of understanding the structure of *want*. In section 7 this approach will be compared with a popular alternative of using contexts which are separate from the structures of entities they contain.

4.2 Relating Descriptions in Terms of the Individuals They May Describe

There are two ways which one entity may be related to another.

a) an individual can be described as a particular kind.

b) an entity can be part of the structure of another.

This section and the next section explain how these two realtionships are reflected in descriptions. Given a description it makes sense to talk about the set of individuals which satisfy that description. These individuals will be termed potential referents. For example, while "a dog" has a single referent, it has many potential referents. A description may or may not have a unique potential referent. It may have no potential referents, e.g., "a round square"; it may have one, e.g. "the President of the United States in 1978"; or it may have many, e.g. "a dog". Since a single individual can be a potential referent of many descriptions, descriptions can be extensionally related by a Venn diagram of their potential referent sets as shown in Figure 2.

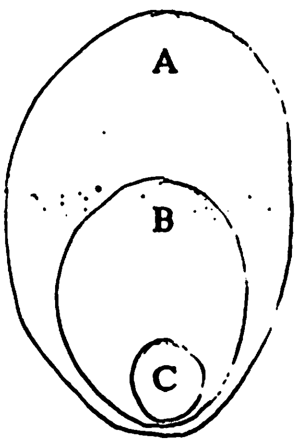

1) description A is a #CHARACTERIZATION of description B when the potential referent set of A is a superset of the potential referent set of B.

2) description C is an #EXEMPLAR of description B when the potential referent set of C is a proper subset of the potential referent set of B.

Figure 2

Venn diagram showing the relationship between the potential referent set of a description and those of its characterizations and exemplars

This is the traditional IS-A relationship defined in terms of the extensions of the concepts, except that IS-A is usually taken to be a proper subset relationship. This means that when using IS-A an = notion is required to indicate that two descriptions refer to the same individuals where here mutual characterization is used.

It is also useful to have the notion the one description, C, is the apposition of two others A and B. As shown in Figure 3 the potential referent set of the apposition of the two descriptions is the intersection of the potential referent sets of each of them.

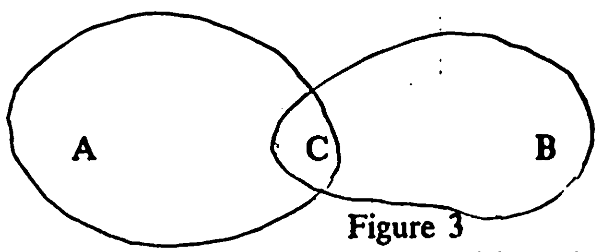

Figure 3
Apposition, C, is the apposition of A and B

The disjunction of descriptions A and B has a potential referent set which is the union of the potential referent A and B.

Finally the negation of a description A has a potential referent set which includes everything except the members of the potential referent set of A. See Eaton (1925) on this interpretation of negation.

4.3 Description of Structure

A description A can be a role of another description B. This is interpreted to mean that any referent of A subsists in the structure of the referent of B. For example, suppose we define 1978-PRESIDENT to be a role in U.-S.-GOVERNMENT.

Then any entity such as *Carter* which is a referent of 1978-PRESIDENT subsists in the structure of *U.S. Government*. Of course, *Carter* may also be described in other ways and these descriptions characterize one another as described in the last section.

One important structure is what is usually called a set with each element explicitly described. This will be called an itemization. Each element description is a role in the itemization. We will denote the itemization of A, B, and C by {A, B, C}.

4.4 Functions over Concepts

To describe an entity, *a*, meeting the description, A, we can do any or all of the following:

1. Declare it has an entity, *b*, meeting a certain description, B, in its structure; [A #ROLE B].

2. Declare it is in the structure of an entity, *b*, meeting a certain description, B; [B #ROLE A].

3. Characterize:
   a) the entity, [A #CHARACTERIZATION C]
   b) an entity in its structure, [A #ROLE B] [B #CHARACTERIZATION C]
   c) an entity whose structure the given entity is in, [B #ROLE A] [B #CHARACTERIZATION C].

4. Describe a state on:
   a) the entity
   b) an entity in its structure
   c) an entity whose structure the given entity is in.

Such a state may of course involve other entities as well.

The distinction we have in mind between characterizing an entity and describing a state on an entity is that between "is-a" and "is" predication which goes back at least to Aristotle (). In English we see the distinction in:

Snow is white.	is
White is a color.	is-a

Note that these two statements do not imply that snow is a color.

One may question whether two kinds of predication are needed or whether one can be reduced to the other. In fact, Russell popularized the use of just "is" predication in mathematics. [???] gives a discussion of the issues involved.

We have chosen to reduce "is" predication to "is-a" predication. In order to do this and to implement option 4. above it is useful to introduce the notion of functions. For example, we will take TALL to be a function which maps the description, A, into the description TALL-A. The potential referent set of TALL-A comprises those entities, *a* in the potential referent set of A whose *height* attribute has a *measure* which lies in the upper range of the distribution of the *measures* of the *height* attributes of the peer group of that potential referent, *a*. That is, the function TALL derives from a state, *s*, on the *measure* of the *height* attribute of the entity, *a*, described by the argument, A, to TALL.

Descriptions and functions will be referred to collectively as concepts. Descriptions refer; functions do not refer but map concepts into concepts. Functions may take either descriptions or other functions as arguements. For example, we might have the function, VERY, applied to the function, TALL, yielding the function, VERY-TALL; then VERY-TALL applied in turn to the description, BOY, yielding the description, VERY-TALL-BOY.

Since functions do not refer, it doesn't make sense to relate concepts which are functions by #CHARACTERIZATION and #EXEMPLAR links. It does make sense, however, for concepts which are functions to have #ROLE's.

Two functions F1 and F2 may or may not commute. We make the convention that a concept A is modified to the concept which is the result of apply commutative functions F1 and F2 to the unmodified A by writing

[A #FUNCTION F1 F2]

Functions may be classified by the way in which the mapping between their arguments and values is defined. First, we distinguish <u>adjuncts</u>, <u>conjuncts</u>, and <u>disjuncts</u>.

a) Adjuncts add to or modify a description - e.g. TALL.

b) Conjuncts produce an itemization which contains their argument as one element - e.g. ONLY. For example, ONLY might take THE-DOG into the itemization {THE-DOG, NOTHING-ELSE}.

c) Disjuncts produce an itemization which contains their argument as one element and a description of a condition on their argument as a second element, e.g. SURPRISINGLY. SURPRISINGLY might take HE-SNORES into the itemization {HE-SNORES, THAT-HE-SNORES-IS-SURPRISING}.

Adjuncts may be further subdivided into <u>hedges</u> (Lakoff) and <u>non-hedges</u>. A hedge builds its output description from its input description. For example, in

John is a regular fish.

the hedge, "regular", selects some part of the description of the function of a fish, but not the description of the form of a fish.

In "alleged communist" the hedge ALLEGED applied to COMMUNIST creats ONE-SAID-TO-BE-A-COMMUNIST. That is, ALLEGED is derived from a state and it uses its argument to construct an exemplar of a description of that state.

An example of a non-hedge adjunct is TALL, discussed above. Non-hedge adjuncts are always derived from states.

Quirk and Greenbaum (1973) point out that "old friend" can mean either "a friend who is old" or "someone with whom an old friendship" exists. From this we note that when an adjunct is applied to a relational characterization like FRIEND, it may

a) be derived from a state on the referent of FRIEND

b) be derived from a state on the state (or process), FRIENDSHIP, on which the relational characterization FRIEND is based.

5. Parts of Speech

It is not the purpose of this paper to develop a grammar of English. But since the examples used derive from English words and phrases, it may be helpful to at least relate the material in the previous section to the traditional parts of speech: noun, verb, adjective, adverb.

5.1 Assertions

Vendler (1967) notes that while one can say either

The collapse of the Germans was in 1944.
The collapse of the Germans is a fact.
The collapse of the Germans was an event.

one cannot say

* That the Germans collapsed was an event.
* That the Germans collapsed was in 1944.

but only

That the Germans collapsed is a fact.

In this way, he argues for a distinction between events and facts. He says, for example, that although some people followed the collapse of the Germans, they did not follow the fact that the Germans collapsed.

The phrase "that the Germans collapsed" contains a tensed (or finite) verb. Joos (1968) maintains that the finite verb is distinguished in that it makes an assertion, whereas an infinitive like "to put" merely names its action. The Kiparskys (1970) point out that in comparing

That the Germans collapsed is tragic.
That the Germans collapsed is true.

we see that "tragic" presupposes that its subject is a fact, while "true" does not.

From this we are lead to understand Joo's distinction to say that a finite verb describes an assertion which may or may not be presupposed true. An infinite describes a state.

The sentence

The Germans collapsed.

may be divided into a subject "the Germans" and a predicate "collapsed". The predicate is applied to the subject to form a sentence. This can be modelled by taking COLLAPSED as an assertion function from descriptions into descriptions of assertions. This is in line with Jespersens (1937) intuition that the predicate is secondary to the subject and parallels the treatment in Montague (1969) and Thomason and Stalnaker (1973).

We postulate the existence of a process, *collapse*, which has a relational-characterization *subject-of-collapse* in its structure. Thus, to meet the description COLLAPSE is to be a *collapse*, to meet the description SUBJECT-OF-COLLAPSE is to be something that collapsed. The intransitive verb "collapse" has the sense SUBJECT-OF-COLLAPSE. That is "to collapse" is to admit of the description of carrying out the *subject* role of a *collapse*.

We let assertion function PAST-TENSE have an OBJECT role which is characterized as a condition. The assertion function COLLAPSED is like PAST-TENSE except that its OBJECT role is characterized as SUBJECT-OF-COLLAPSE. COLLAPSED applied to THE-GERMANS creates the assertion that THE-GERMANS may be characterized as SUBJECT-OF-COLLAPSE.

Thomason and Stalnaker point out the difference between predicate modifiers and sentence modifiers. For example, "frequently" is a sentence modifier in

Frequently, somebody got drunk.

and a predicate modifier in

Somebody frequently got drunk.

This shows the need for maintaining the predicate and the predication (or assertion) as separate entities to be modified. The above formulation does this.

5.2 Correspondences

Table 2 gives sample correspondences between parts of speech and concepts.

	Part-of-speech	Concept-type
1)	noun	description
2)	verb	description of state of being the subject of a state or process
3)	adjective	function
4)	adverb	function

Table 2

Jespersen (1937) worked out a theory of ranks. For example, in "not very carefully worded remark" he claims the structure is ((((not very) carefully) worded) remark). "Remark" is a primary, "worded" a secondary, and the rest tertiaries. By taking descriptions as his primaries, functions of descriptions as his secondaries, and functions of functions as his tertiaries, the reader can get an idea of how our concept types will translate into a grammar.

6. Concepts

Descriptions and functions are referred to collectively as concepts. Each concept will have a corresponding node in our semantic network.

6.1 Concept Definitions

By what has been said thus far, concepts may be classified as shown in Table 3. For each relational characterization description we know there is a condition on which the relational characterization is based, e.g. *husband* on *marriage*, and for each attributive characterization description we know there is an attribute on which it is based, e.g., *teenager* on *age*. Similarly, we know there is a state which each inherent function is derived from, e.g., LONG on a state on measures of *length*.

Type	Example(s)
Descriptions	
Attribute/Physical-Disturbance	temperature
state/process	party
Relational Characterization	husband
Attributive Characterization	teenager
Other	dog
Functions	
Inherent Function	long, very
Non-inherent Functions	
Hedge	regular
Base State Function	civil (engineer)

Table 3

Such information can be used to organize the concepts in semantic memory, but it seems unlikely that it constitutes grounds for the definition of concepts. The traditional method of defining concepts is to start with some primitive set of notions and then build up definitions by combining these according to some logic. It seems unlikely that the notions we have introduced are adequate to that task.

Rather than define concepts in this traditional sense, we assume that concepts acquire their uniqueness by the unique set of structure (#ROLE) and abstraction (#CHARACTERIZATION, #EXEMPLAR) relationships they have to other concepts.

Earlier it was suggested that the world may be profitably viewed as one big structure and that people choose to regard various substructures of this structure as individuals. Similar individuals may be considered in concert to form descriptions at various levels of abstraction.

By <u>semantics</u> one understands the question of whether the descriptions in a semantic network are faithful to the real world. Since any system of knowledge representation usually provides a method of combining descriptions to form new ones, a potentially infinite number of descriptions can be generated from those currently existing in a network. There is the question of whether those generated descriptions will also be faithful to the world.

The most popular approach to this problem stems from Tarski (). His method is based on using two kinds of rules:

a) Rules to verify that some "base" descriptions correspond to the world.

b) Rules for verifying the composition of descriptions given that the individual descriptions can be verified.

This approach leads naturally to the idea that concepts like DOG should be verified in terms of rules involving sense data. The sense data will somehow provide the set of primitive terms used to build up the definition of concepts.

By contrast, we suggest that the focus should be on verifying that the entire semantic network represents some portion of the world, taken as as a single structure. There are several reasons for preferring this wholistic approach in the context of semantic networks.

First, consider the fact that the typical computer system employing a semantic network suffers more sensory deprivation than a blind, deaf person. It sends and receives only character strings. The question of concept verification in terms of sensory primitives is not currently relevant. However, suppose we changed the name of every node and link in a very detailed semantic network to a number. We postulate that it would still be possible to determine if that network was a faithful model of the portion of the world, since only one portion of the world, if any, would admit to being described by that particular network. We would do this by comparing the computer's model with our own.

The relevant process in semantic networks is the verification of the computers model of the world by comparing its structure with the structure of our own model.

There are other factors which contribute to an interest in the wholistic approach. One is a certain pessimism about ever getting a model of the world which is sufficiently accurate to permit confidence to be placed in deductions of any complexity. Certainly this can be done in specialized areas like mathematics. But in general it seems necessary to operate with very incomplete and inaccurate data. There is a question of whether the difficulties of prediction faced, to take an extreme example, by economists, can ever be overcome.

One must also recognize that some concepts like ELECTRON, BACHELOR, and SLAVERY seem inherently to rest on their relationship to other concepts rather than their definition in terms of sense primitives.

The thesis that a primary mode of thinking is the expansion of concepts into their definitions has never made much headway. For example, Schank's () moves in this direction lead to a combinatorial explosion of irrelevant inferences. Authors such as Minsky (), Bobrow and Winograd, and Fodor () have taken positions against such an expansion. If expansion of definitions is not of key importance in on-going thought, their key use if any must be primarily organizational. This paper is intended as a step towards finding organizational principals other than definitions in terms of a set of basic notions. Of course, it remains to be seen whether inferencing and concept learning can take place in the framework presented here. We can imagine a machine with a keyboard as its only sense organ learning a great deal about the structure of the world and the abstractions people use to describe it, but it is hard to predict just how far this could go.

The question of definition expansion comes up in the application of functions derived from states, e.g. long, to descriptions, e.g. snake. Rather than have a concept in memory for the function, long, why not just use the underlying state and eliminate the need for the function. That is, represent long snake by representing the state on the measure of the length of snake. Our answer lies in the belief that it is usually not the definition of long that matters in dealing with a long x, but rather the implications of an x being long. For example, a long train needs several engines and must be waited for longer at crossings while a long snake is often more dangerous to man. We want to know that an x would satisfy a description resulting from the application of a function more often than we want to compute with the description which would result from the function application.

6.2 Expectations

We recognize a sentence like
The duck was too old to eat.

as ambiguous since it could mean either that the duck could not eat or could not be eaten. We find less ambiguity in
The bread was too old to eat.
The dog was too old to eat.

This must be because we know what things are eaten and who the eaters are. Highly implausible interpretations like taking in the car in
I road down the street in the car.

to the give location of the street are ignored. We can mentally construct a car big enough to hold a street, but this mental exercise is not invoked by the above sentence. Instead, in the car is taken to apply to ride; cars being the archtypical thing to ride in.

To describe what things are eaten and who the eaters are we can make statements at different levels of abstraction. For example, we could describe what eats bamboo shoots or what eats at all.

Given that the level of abstraction is chosen, we may specify the eaters and the eaten by, for example, taking SUBJECT-OF-EAT and OBJECT-OF-EAT to be relational characterization roles in the structure of EAT. Then, for example, SUBJECT-OF-EAT can be related by #CHARACTERIZATION and #EXEMPLAR to descriptions of eaters. For example, we might have

```
[ANT-EATER #CHARACTERIZATION SUBJECT-OF-EAT-ANTS]
[SUBJECT-OF-EAT-ANTS #EXEMPLAR ANT-EATER]
[FAVORITE-ACTIVITY-OF-ANT-EATER #ROLE ANT-EATER
        #EXEMPLAR EAT-ANTS]
```

In the case of functions, we will write

```
[A #APPLICABLE-FUNCTION B C]
```

to indicate that B and C are functions which can be applied to A. For example, we might have

```
[SUBJECT-OF-OLD #EXEMPLAR WINE LIVING-THING]
[WINE #APPLICABLE-FUNCTION OLD YOUNG RED WHITE]
```

The reason for allowing links both from WINE to OLD and from OLD to WINE is two fold. First, it is sometimes the case that given either end of a link we wish to be able to find the other without an exhaustive search.

The second reason is to provide flexibility in handling rare cases. For example, by writing only

```
[SUBJECT-OF-LAYS-EGGS #EXEMPLAR BIRD REPTILE]
[PLATYPUS #CHARACTERIZATION SUBJECT-OF-LAYS-EGGS]
```

when checking if something can lay eggs we will have only BIRD and REPTILE to consider. A platypus is known to lay eggs, however, because this is directly on PLATYPUS. Of course, when asked what lays eggs we will find from SUBJECT-OF-LAY-EGGS that birds and reptiles do. When asked if we are sure nothing else does, we will have to think of all the strange animals we know and check if any of them do.

7. Specialization

Specialization is our particular method for constructing unique names for concepts. Since it is easy to conceive of semantic networks containing tens of thousands, if not millions, of concepts, an orderly naming scheme is very important.

McDermott () has suggested that concepts should be given names in some invented language, lest we confuse our understanding of the name with an understanding of the concept. It seems unrealistic, however, to ask someone to learn several thousand names in order to use or evaluate a system. McDermott's suggestion seems practically limited to small or one person systems. For others, it will be necessary to name concepts using terms the user is familiar with, either directly or by supplying a translation. Since a natural language, such as English, is the only widely known form of expression for non-mathematical concepts, it is the only real choice for concept names. Most all semantic networks do, in fact, use English in creating their node and link names. The issue is to what extent this can be done in a principled way.

In a semantic network concepts can be named and they can also be individuated (or defined) by their relationships to other concepts in the network. When a person is studying one concept in the network it is extremely helpful if he can get at least a general idea of what the concepts are that it is related to in the network just from their names. Otherwise mapping the semantic networks onto his knowledge of the world becomes a combinatorial puzzle. Fortunately, this can be made possible by capitalizing on a natural skill of every English speaker, the ability to extend a sense of a word for use in a particular context.

7.1 Word Sense Extension

Rosch () suggests that people identify entities by comparing them to prototypes. For example, the robin and the sparrow might be prototypical birds. For this reason, people feel comfortable with (Lakoff)

A robin is a bird par excellance.

but not

A chicken is a bird par excellence.

Suppose people identify several prototypes and examples with a particular sense of a word. That word may be used for anything which is a good match for one of the prototypes. When a new thing comes along which doesn't match any of the prototypes too well, we may extend the sense of an existing word.

Wittgenstein () has pointed out that senses of a word have a family resemblance

- while any two have much in common, little is common to them all. A given sense can be extended in different directions and each of these extensions can be further extended.

In thinking about specialization it is helpful to have examples of how different senses of a word are related. Examples of word sense change may be found in Bloomfield ().

Narrowing:

Old English *mete* 'food' > *meat* 'edible flesh'

Old English *deor* 'beast > *deer* 'wild ruminant of a particular species'

Old English *hund* 'dog' > *hound* 'hunting-dog of a particular breed'

Widening:

Middle English *bridde* 'young birdling' > *bird*

Middle English *dogge* 'dog of a particular (ancient) breed' > *dog*

Latin *virtus* 'quality of a man (*vir*), manliness' > French *vertu* (> English *virtue*) 'good quality'

Metaphor:

Primitive Germanic *['bitraz] 'biting' (derivative of *['bi:to:] 'I bite') > *bitter* 'harsh of taste'

Metonymy - the meanings are near each other in space or time:

Old English *ceace 'jaw'* > *cheek*

Old French *joue* 'cheek' > *jaw*

synecdoche - the meanings are related as whole and part:

*Primitive Germanic *['tu:naz] 'fence' (so still German Zaun)* > *town*

pre-English *['stobo:] 'heated room' (compart German *Stube*, formerly 'heated room', now 'living room') > *stove*

. *Hyperbole* - from stronger to weaker meaning:

pre-French **ex-tonare* 'to strike with thunder' > French *etonner* 'to astonish' (from Old French, English borrowed *astound, astonish*)

Litotes - from weaker to stronger meaning:

pre-English *['kwalljan] 'to torment' (so still German *qualen*) > Old English *cwellan* 'to kill'

Degeneration:

Old English *cnafa* 'boy, servant' > *knave*

Elevation:

Old English *cniht* 'boy, servant' (compare German *Knecht* 'servant') > *knight*.

If we accept the hypothesis that word sense change occurs by extension to a new sense and then extinction of the old, then we see that there will be periods where old and new senses of the types above exist simultaneously. Thus two senses of a word may be related in all the ways shown above.

7.2 The Specializer

Let us identify word senses with concepts. Suppose that the prototypical HOUSE is one for humans. It is misleading to refer to a house for dogs as a HOUSE, because it differs from the HOUSE prototype in too many ways. What English lets us do is to indicate by DOG HOUSE that the HOUSE concept is to be modified in some way so as to make it appropriate to DOG. As Marchand () points out, this process of word formation is not limited to nouns, e.g. "color blind", and "hog tie". We postulate that when someone hears a term like "dog house" he can use some mix of the following three strategies in understanding it.

a) He may have a concept for which he knows this to be an unanalyzed name, e.g. "bull dog" or "skid row". How many people know that a bull dog is a dog for fighting bulls or that skid row was a row of shantys along a log skid in early day Seattle Washington.

b) He may have a concept for which this would be an appropriate name. For example, although people don't have a name for the little plastic cylinder on the end of a shoestring (Woods 1978) they know of it and they can understand the sentence "my shoestring end came off". This presumably because END could be extended in the context of SHOESTRING to mean this cylinder, and this can come off.

c) He may create a concept by principled analogy. For example, by analogy with "dog house" and "bird house" one could understand "grasshopper house". We contend that this process of analogy is always at work. We understand "crackage" by analogy with "breakage", "pilferage", "damage". We understand "inchage" by analogy with "mileage" and "yardage".

7.3 Construction of Names

The problem with using expressions like "dog house" as the names of concepts is that they can be ambiguous. For example, "snake poison" can mean either poison from snakes or poison for snakes. A "steel drill" can be either made from steel or for drilling steel. A "river bank" could be a financial institution.

The ambiguity is bad for two reasons. First the user does not know which concept is meant. While this is bad, the user can always study the semantic network in order to see how the concept is used and thus what is meant. A worse problem is that the naming scheme does not generate enough distinct names; it proposes the same name for two different concepts. We wish a scheme which will allow us to generate a distinct name for each concept.

The ambiguity in "river bank" arises from the ambiguity in "bank". It seems reasonable to resolve this by saying that names will be constructed from senses of words instead of words. Thus, RIVER BANK and RIVER FINANCIAL-BANK. The ambiguity in "snake poison", however, does not seem to arise from any ambiguity in "snake" or "poison". Instead it arises from two different relationships between these terms. We have already described in section 7.2 how a person uses his knowledge of the world in order to determine what relationship is meant. It has long been recognized that there are certain regularities to this process. For example, corresponding to any substance such as steel, a "steel x" can mean an x composed of steel.

Rhyne (1976) discusses attempts by various authors to capture these regularities in the case of compound nouns. These may be divided into attempts to classify the semantic relationship between the two constituents of the compound and methods for deriving the compound from a relative clause or other expression.

The difficulty with expressing the relation between the elements of a compound by a linguistic expression may be seen by considering "widget store". Everyone knows what stores are for and so one meaning of "widget store" is clear. But how should this be expressed by a relative clause. Is it a store where widgets are sold, a store which carries widgets, a place to go to buy widgets, etc.

We postulate that somewhere in the semantic network there is a concept which corresponds to the store's wares. A "widget store" is a store where this wares concept may be characterized as widgets. There may be many different relative clauses which can express the relationship between the store concept and its wares concept, since in general they will be linked by a network of different relationships.

Loosely speaking, our solution to this problem will be to take names as triples. The second element of a triple will indicate how the first and third are related. In an example like "widget store" the second element will indicate how WIDGET and STORE are related by providing access to the STORE's WARES concept mentioned above.

It is necessary to introduce some notation to describe exactly how this is done.

7.4 Notation for Specialization

We let semantic memory be made up of concepts and symbols. Symbols are written as character strings between double quote marks, e.g., "ENGINE". As shown in Figure 4, the most general concept is SUMMUM-GENUS. Symbols are taken to be atomic in the sense that they cannot be decomposed in any way. Concepts are non-atomic. They are constructed from SUMMUM-GENUS and symbols by using the binary operation, <u>specialization</u>. Specialization is written:

(<u>genus</u> <u>specializer</u>) where <u>genus</u> is a concept and <u>specializer</u> is a concept or symbol.

We say that a concept is a <u>specialization</u> of the concept in its genus position.

SUMMUM-GENUS

(SUMMUM-GENUS "ENGINE") (SUMMUM-GENUS "FIRE")

((SUMMUN-GENUS "ENGINE") (SUMMUM-GENUS "FIRE"))

Figure 4

For example in Figure 4 we have constructed two specializations of SUMMUM-GENUS: (SUMMUM-GENUS "ENGINE") and (SUMMUM-GENUS "FIRE"). We have then specialized (SUMMUM-GENUS "ENGINE") by (SUMMUM-GENUS "FIRE").

Moving up in the genus direction, it is clear that concepts are the nodes of a tree with SUMMUM-GENUS at the root. SUMMUM-GENUS is taken as a specialization of itself;

SUMMUM-GENUS = (SUMMUM-GENUS "SUMMUM-GENUS")

We say that any concept, C, forms a <u>class</u> which contains in the sub-tree whose root is C, including C itself. It is convenient to refer to all of these concepts as specializations of C.

If specialization is carried to very many levels, the expression for a concept quickly becomes unwieldy. We avoid this through the familiar mechanism of labeling. The expression

<u>label</u> = <u>concept</u> where label is any string of letters digits, hyphens, and periods.

assigns <u>label</u> to <u>concept</u>. A label is just a notational abbreviation for the parenthesized expression that exhibits the genus and specializer of a concept; it has no semantic significance in and of itself. Using labels we might rewrite Figure 4 as Figure 5

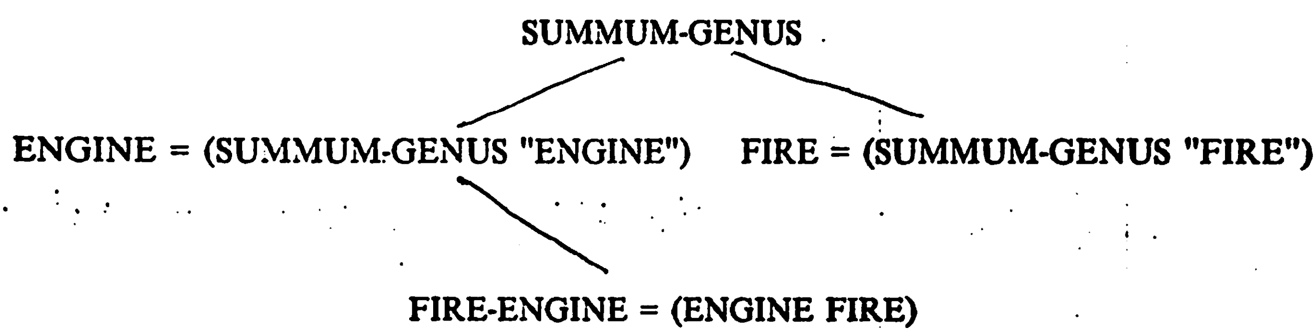

Figure 5

7.5 <u>Meta-Specializers</u>

Specialization as defined is a binary process, combining a genus and specializer. To implement triples, A B C, we will write ((A B) C) as (A*B C) and refer to A as the genus, B as the meta-specializer, and C as the specializer. Meta-specializers may be classified into seven types. These seven types may be further classified into three categories as shown in Figure 6

category	meta-specializer	abbrev.	use	English phrase
idiomatic	#SPECIES	S	(DOG*S BULL)	bull dog
"	#STEREOTYPE	T	(MAN*T FAT)	(circus) fat man
"	#INDIVIDUAL	I	(DOG*I FIDO)	Fido
structural	#ROLE	R	(OBJECT.*R HIT)	object of hit
"	#APPOSITIVE	A	(DOG*A PET)	pet dog
"	#FUNCTION	F	(DOG*F FAT)	fat dog
property	any-concept except the other 6 types of meta-specializers	none	(HIT*OBJECT. BALL)	hit ball
	#PARTITIVE	P	(THAT*P SOME)	some of that

Figure 6
Types of Meta-specializers

The first three meta-specializers, #SPECIES, #STEREOTYPE, and #INDIVIDUAL all tell the user that the concept is idiomatic. For example, when the user sees (DOG*S BULL) he knows that either he already knows what a "bull dog" is or he will have to look at how the concept is used in the semantic network to figure this out. His knowledge of the concept BULL will not be of any help. Any concept with the meta-specializer #INDIVIDUAL must represent an individual or a kind. If it represents a kind, the specializer must be KIND. #SPECIES is useful in setting up a Linnean classification system. If (A*S B) and (A*S C) are any to any two concepts with the same genus and #SPECIES as meta-specializer, then their potential referent sets are mutually exclusive. #STEREOTYPE is used for the remaining idiomatic specializations.

The structural meta-specializers #ROLE, #APPOSTIVE, #FUNCTION, and #PARTITIVE indicate that the genus and specializer are related by meta-level constructs. Specifically, (A*R B) is a role in B. (B*A C) has a potential referent set which is the intersection of the potential referent sets of B and C. (A*F B) has the same potential referent set as the concept which results from the application of function B to concept A. (A*P B) must be defined in terms of inheritance, which is described in the next section. Roughly speaking, (A*P B) inherits both A and B, with B dominating in case of conflict.

By way of example, it is interesting to contrast (MAN*F FAT) and (MAN*T FAT). Notice that <u>fat man</u> is ambiguous between the sense of a man who is fat, and the idiomatic sense of a man who might work in a circus. The first sense may be pronounced with a slight pause between fat and man. The distinction can be seen in <u>very fat man</u> vs. <u>circus fat man</u>.

This distinction may be described as a choice between an adverbial and a lexical combination of the words. By definition, the adverbial relies primarily on the meaning of its constituent parts, while a lexical combination relies on the lexicon or memory for its interpretation. These two readings will be written (MAN*F FAT) and (MAN*T FAT).

Now suppose we have a concept, (A*B C) where B is not one of the above six meta-specializers. In this case we stipulate (B*R A) must be a role node in A, and

(A*B C) is a specialization of A whose role node corresponding to (B*R A) is characterized by C. For example (DRILL*SUBSTANCE STEEL) is a specialization of DRILL which implies that (SUBSTANCE*R DRILL) exists and that we have [(SUBSTANCE*R (DRILL*SUBSTANCE STEEL)) #CHARACTERIZATION STEEL].

Recall the example of "widget store" used in Section 7.3. This would become, for example, (STORE*WARES WIDGET) where we create the role (WARES*R STORE) and characterize it by the concept (or concepts) in the network referring to the wares in STORE. This is done so that the reader can understand precisely what is meant by (WARES*R STORE).

7.6 Inheritance

Recall that we proposed that the concept (HOUSE*S DOG) is an extension of the concept HOUSE. That is, the concept HOUSE forms a starting point for the construction of the concept (HOUSE*S DOG). We model this by stipulating that a concept inherits all the #ROLE's and #FUNCTION's of its genus, unless the concept itself has the same role or a contradictory function explicitly given. Some concepts will inherit a great deal and others will inherit very little. No idiomatic specialization can be completely understood by inheritance and interpretation of the specializer.

To create unambiguous names for the two senses of "mine" as a pronoun and a place, we could write

PRONOUN-MINE = (PRONOUN*S "MINE")

```
PLACE-MINE = (PLACE*S "MINE")
```

PRONOUN-MINE will then be a concept with the properties of pronouns while PLACE-MINE will be a concept with the properties of places. (These can be recognized as senses of "mine" if we make the convention that concepts specialized by a symbol are senses of the word spelled like that symbol.)

As an example of widening of a word sense, consider "animal". Living beings may be divided into plants and animals. Animals may in turn be divided into human beings and animals. Animals may in turn be divided into animals, fish, insects, birds, etc. Suppose that we consider the narrowest sense of "animal" to be the most central. Then we might state

```
 ANIMAL = (LIVING-THING*T "ANIMAL")
[NON-HUMAN-ANIMAL = (ANIMAL*T "ANIMAL")
        #EXEMPLAR ANIMAL FISH INSECT BIRD]
[SCIENTIFIC-ANIMAL = (NON-HUMAN-ANIMAL*T "ANIMAL")
        #EXEMPLAR NON-HUMAN-ANIMAL HUMAN-BEING]
```

This way of looking at things says that even though we know that all ANIMAL's may be described as NON-HUMAN-ANIMAL's, in fact, the concept of NON-HUMAN-ANIMAL is formed by modifying the concept of ANIMAL.

This representation makes it necessary to override most of the properties which NON-HUMAN-ANIMAL inherits from ANIMAL. Within the conventions presented here one could also choose to state

```
[SCIENTIFIC-ANIMAL = (LIVING-THING*S "ANIMAL")
        #EXEMPLAR NON-HUMAN-ANIMAL HUMAN-BEING]
[NON-HUMAN-ANIMAL = (SCIENTIFIC-ANIMAL*S "ANIMAL")
        #EXEMPLAR ANIMAL FISH INSECT BIRD]
 ANIMAL = (NON-HUMAN-ANIMAL "ANIMAL")
```

In this case, very little is available to be inherited and we can also indicate the Linnean classification. The choice between these alternatives is not dictated by our scheme.

Note that
[DOG #CHARACTERIZATION ANIMAL]

means that any individual which meets the description DOG also meets the description ANIMAL. It would be reasonable to state then that DOG's inherit the properties of ANIMAL's. The reader may then wonder how this inheritance differs from that implied by, for example,

DOG = (ANIMAL*S "DOG").

The answer is that in the case of specialization we are not guaranteed that any individual which is a DOG is also an ANIMAL. Specialization does not imply that a criterial description is inherited. In fact, just the opposite may be the case. For example, consider a concept formed with a hedge like "regular". Individuals meeting the description (FISH*F REGULAR) do not meet the description FISH.

7.7 Descriptions of Kinds and Individuals

English distinguishes between unquantifiable entities, like *health*, and quantifiable entities, like *dog* and *water*. Quantifiable entities may be further broken down into countable entities, like *dog*, and mass entities, like *water*.

There are three types of generic sentence involving countable entities.

The motor car is a popular means of transport.
Motor cars are a popular means of transport.
A motor car is a popular means of transport.

That these are not equivalent may be seen by comparing

The motor car was a major influence in this century.
A motor car was a major influence in this century.

Thus, in the case of countable entities, it is useful to distinguish the concept corresponding to "the motor car" from that corresponding to "a motor car", that is, to distinguish a description of a kind from a description of a corresponding individual. The same problem does not arise with mass or unquantifiable entities.

We are lead to the following conventions:

The description of an individual dog will be notated, for example,
DOG = (DOMESTIC-MAMMAL*S "DOG")

and that for dog kind
DOG-KIND = (DOG*I KIND)

A practical application may have little reference to concepts such as DOG-KIND and thus they may often not be formed.

In the case of mass or unquantifiable entities we will just let, for example,
WATER = (LIQUID*S "WATER")

represent either the kind or an indeterminate amount.

7.8 The Adjectives, Long, White, and Thin

Some final insight into specialization may be gained by comparing the phrases "long worm", "long train", "long movie", "white snow", "white person", and "white wine". Precisely the same sense of "long" may be applied to "worm", "train", and "movie". This sense of "long" is a function which relates the length of its argument to those of the argument's peer group. A long worm is much shorter than a long train because worms run shorter than trains. A movie is long in time because the length of a movie is its length in time.

By contrast, different senses of white are required for snow and wine, e.g. (WHITE*SUBJECT SNOW). Comparing "thin", "slim", and "skinny" and "fat", "plump", and "pudgy", we see that an adjective can also be specialized to reflect the speakers attitude towards the condition on which it is based, e.g. "slim" is desirable "thin", "skinny" is undesirable "thin".

8. Selectional Restrictions

Selectional restrictions have been used to semantically constrain, for example, the subject of senses of "bark" so that

The block barked.

is semantically unacceptable, while

The dog barked.
Fido barked.
The animal barked.
The father barked.
The seal barked.

all are acceptable, and

The Sargent barked an order.

must have a sense of "bark" different from the first.

In the representation system proposed here we would have

```
[(SUBJECT*R ANIMAL-BARK)
    #CHARACTERIZATION C1  C2  ...
    #EXEMPLAR         E1  E2  ...
    #FUNCTION         F1  F2  ...
```

#APPLICABLE-FUNCTION AF1 AF2...]

and in addition (SUBJECT*R ANIMAL-BARK) might have any number of roles similarly described. The question is how these descriptions could be used in determining the acceptability of various candidate subjects of ANIMAL-BARK.

Almost anything can be the subject of a verb in some context, e.g.

> A funny thing happened at work this morning. A fellow came into my office carrying a dog and a shoebox. The shoebox barked. I wonder if the dog was a ventriloquist.

So the question we ask here is not whether a certain candidate concept could be used to fill a role, but only whether the candidate is at odds with what we know about the traditional fillers of that role.

To begin, recall that #CHARACTERIZATION and #FUNCTION apply to every individual while #EXEMPLAR and #APPLICABLE-FUNCTION apply only to some individuals. Thus a candidate could be ruled out by showing that it could not have a given #CHARACTERIZATION or #FUNCTION. It could be ruled in by showing that it could be characterized by any #EXEMPLAR. Showing that it could be described by an #APPLICABLE-FUNCTION neither rules it in or out, but is perhaps evidence for keeping it in.

The usual practice in existing computer programs is to rule candidates in through class inclusion, e.g. to ANIMAL-BARK you must be a dog or a seal. Ruling in by class inclusion is perhaps used in existing programs because they are typically restricted to those of the argument's peer group. A long worm is much shorter than a long train because worms run shorter than trains. A movie is long in time because the length of a movie is its length in time.

By contrast, different senses of white are required for snow and wine, e.g. (WHITE*SUBJECT SNOW). Comparing "thin", "slim", and "skinny" and "fat", "plump", and "pudgy", we see that an adjective can also be specialized to reflect the speakers attitude towards the condition on which it is based, e.g. "slim" is desirable "thin", "skinny" is undesirable "thin".

8. Selectional Restrictions

Selectional restrictions have been used to semantically constrain, for example, the subject of senses of "bark" so that

> The block barked.

is semantically unacceptable, while

> The dog barked.
> Fido barked.
> The animal barked.
> The father barked.
> The seal barked.

all are acceptable, and

> The Sargent barked an order.

must have a sense of "bark" different from the first.

In the representation system proposed here we would have

```
[(SUBJECT*R ANIMAL-BARK)
     #CHARACTERIZATION C1  C2  ...
     #EXEMPLAR         E1  E2  ...
     #FUNCTION         F1  F2  ...
```

the genus and all characterizations of any concept reached during running up except SUMMUM-GENUS. We run up from both the candidate and the exemplar and we note all concepts at which a path from the candidate intersects a path from the exemplar. Since all paths meet at SUMMUM-GENUS, only three situations can obtain.

a) The intersection is the exemplar, e.g. "Fido barked" intersects at DOG. Accept.

b) The intersection is the candidate, e.g. "The animal barked" intersects at ANIMAL. Accept.

c) The intersection is at some other concept, e.g. "The block barked" intersects at, say, PHYSICAL-OBJECT. "The father barked" intersects at, say ANIMAL. In this case we accept unless the two paths come in through mutually exclusive SPECIES, e.g. (PHYSICAL-OBJECT*S "BLOCK") and (PHYSICAL-OBJECT*S "ANIMAL") or unless the two paths pass through mutually exclusive FUNCTION's, e.g. "hen crowed" would be ruled out because crowed requires a rooster which is male and a hen is female. More explicitly,
ROOSTER = ((CHICKEN*F ADULT)*F MALE)
HEN = ((CHICKEN*F ADULT)*F FEMALE)
MALE = (SEX-FUNCTION*S "MALE")
FEMALE = (SEX-FUNCTION*S "FEMALE")
The paths from ROOSTER and HEN intersect at (CHICKEN*F ADULT) and pass through MALE and FEMALE respectively. These later are explicitly known to be mutually exclusive.

If <u>any</u> intersection is rejected then the candidate is rejected, since this means there is a way of showing the potential referent sets of the candidate and exemplar to be mutually exclusive.

In the psychological literature one finds the quick comparison of a candidate to a prototype being done by comparison of sets of features of each. Tversky () suggests that the first step is to determine the salient dimensions of the comparison. For example, Cuba is similar both to Russia and to Puerto Rico but Russia is not similar to Puerto Rico because the dimensions of their comparison to Cuba are different. One abstracts the appropriate features from each concept. The candidate and exemplar are then compared by the formula

$$W = A\Sigma a_i + B\Sigma b_j + C\Sigma c_k$$

where a_i is the saliency of a feature had only by the candidate, b_j is the saliency of a feature had both by the candidate and the exemplar and c_k is the saliency of a feature had only by the exemplar. Saliency is some function of factors such as the strength of the feature in a concept, its uniqueness to the concept, its well formedness, and its familiarity. Such an approach could be implemented in the framework we provide, but it would appear to require too much computation for current hardware.

Comparison to prototype could be useful in forming new characterizations of concepts, but it alone does not solve the problem of selectional restrictions. For example, what features of "father" and "goat" would be needed so that "the father barked" is more acceptable than "the goat barked". Goats are very similar to dogs but we know they do not bark. Fathers are a diverse set, not having much to do with dogs in particular but we know a dog can be a father.

It appears best to resort to comparison by match of features or structure only when deduction based on explicit knowledge is not possible.

9. Quantifiers

In Section 7.7 we chose to interpret the quantifiable concepts

DOG = (DOMESTIC-MAMMAL*S "DOG") countable individual
WATER = (LIQUID*S "WATER") mass kind or individual

as countable concepts refering to single individuals rather than to an unspecified number of individuals. This can be implemented by giving the most abstract countable concept a quantity of one.

[(QUANTITY*R COUNTABLE-ENTITY)
#CHARACTERIZATION SINGULAR-NUMBER]

All countable entities like DOG will then inherit this unless they have contradicting information.

"Dogs" will be written
DOGS = (DOG*QUANTITY PLURAL-NUMBER)

and "any 5 dogs" will be written
ANY-5-DOGS = ((DOG*QUANTITY 5)*QUANTITY ANY)

Then, if numbers are organized as shown in Figure 7, DOGS will not inherit the SINGULAR-NUMBER characterization of QUANTITY from DOG because the PLURAL-NUMBER characterization of the QUANTITY of DOGS is mutually exclusive with and will thus override SINGULAR-NUMBER. At the same time, ANY isnot mutually exclusive with 5 so the QUANTITY of ANY-5-DOGS will be characterized by both ANY and 5.

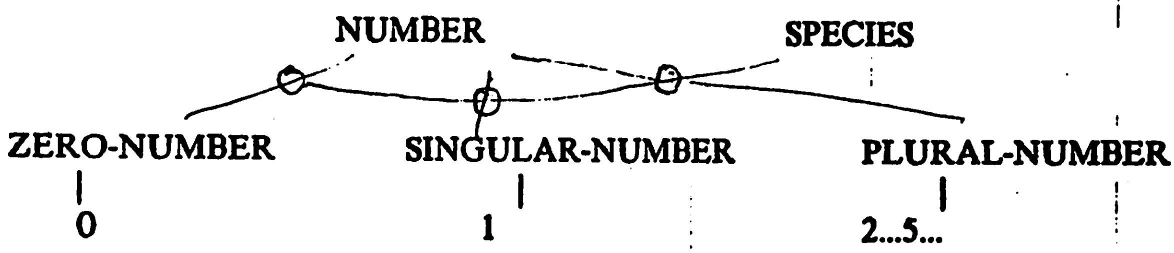

Figure 7
Organization of Numbers

9.1 Groups and Quantities Distinguished

When two conjoined verbs share the same object noun phrase, one expects that they will both take it in the same sense. This accounts for the oddity of a sentence like:

She baked a cake for and called her son John.

So when a sentence like

I opened and ate a can of beans.

seems perfectly normal, one assumes that both these verbs take the same sense of a can of beans in spite of the fact that a person opens the can but eats the beans. A can of beans is an entity like a coke, in

I opened and drank a coke.

a container with something in it, which has properties different than those of an empty container or those of a quantity of the stuff contained. For example, compare

I broke a glass.
? I broke a glass of wine.
I broke a glass containing wine.

Perhaps with less force, the same argument holds for count groups like pack of dogs. A pack of dogs is more than a quantity of dogs. We know quite a bit about dog packs as an entity in their own right.

This behavior of groups is in contrast to the behavior of quantities. In expressions like "5 of the dogs" we know nothing about the quantity except its measure. For this reason we will distinguish between groups and quantities.

Note that the sentences
I crushed 3 cans of bottles.
This container has 3 cups of coffee in it.

are ambiguous. They can mean either that the cans or cups are physically involved, or that they are only measures of the amount. A measure we take to be a form of quantity. Thus one reading of "cup of coffee" is a group, the other is a quantity.

9.2 Representing Groups and Quantities

The three senses of "cup" displayed in the previous section can be represented by the hierarchy:

```
COUNT-QUANTITY = (COUNTABLE-ENTITY*S "QUANTITY")
  MEASURE = (COUNT-QUANTITY*S "MEASURE")
    CUP-MEASURE = (MEASURE*S "CUP")
GROUP = (COUNTABLE-ENTITY*S "GROUP")
  CUP-GROUP = (GROUP*S "CUP")
CUP = (COUNTABLE-ENTITY*S "CUP")
```

The two senses of "3 cups of coffee" will be represented as:

```
((CUP-GROUP*DOMAIN COFFEE)*QUANTITY 3)

((COFFEE*P CUP-MEASURE)*QUANTITY 3)
```

For the partitive reading COFFEE is taken as the genus since a CUP-

MEASURE has all the properties of coffee - except that CUP-MEASURE's are countable and COFFEE is not. The #PARTITIVE meta-specializer is used to denote that the QUANTITY attribute is inherited from the specializer rather than the genus. Recall that the #PARTITIVE meta-specializer stipulates that the concept inherits from the specializer as well as the genus, and the specializer overrides the genus.

9.3 The Partitive Applied to "Kind", "Type", and "Sort"

According to Zandvoort (),

> **231.** In groups with *kind of, sort of* (*this kind of tool, that sort of speech*), *kind of* and *sort of* are often felt to be subordinated in meaning to the following noun. In that case the plural ending is added, at least in familiar English, not to *kind* or *sort*, but only to the last word of the group: *these kind of tools, those sort of speeches.*
>
> What kind of trees are those?
>
> In literary English we also find *all manner of* (*benefits*, etc.).
>
> *Kind* and *sort* take the plural ending after *all*: *all kinds of men, all sorts of people.* The same applies in non-colloquial English to *kind* and *sort* preceded by other qualifiers: *these kinds of tools, those sorts of speeches.*
>
> A genitive-ending would, as usual, be added to the last word of the group: *that kind of student's idea of work, those kind of people's notions of honesty.*
>
> **232.** When *kind* or *sort* are of equal importance with the following noun, they take the plural ending, the following noun being usually in the singular.
>
> What kinds of cherry flourish best in this region?

Constructions with words like "kind" we take to be a special kind of partitive construction. While the partitive "cup of coffee" inherits all but its quantity attribute from "coffee" and the quantity attribute from "cup", "flavor of ice-cream", for example, inherits all but this and its most specific and relevant characterization from "ice-cream."

We see in the examples

I can't identify that flavor (kind, type) of ice-cream.
I can't eat that flavor (kind, type) of ice-cream.

which can be represented

(ICE-CREAM*P (FLAVOR*A THAT))

that verbs like "identify" rely on words like "flavor", "kind", or "type" to specify how they should be identified. As a more complex example, the sentence

I ate some of each of the flavors of ice-cream,

can be represented

```
((((((ICE-CREAM*P ((FLAVOR*QUANTITY PLURAL-NUMBER)*A THE))*P
      COUNT-QUANTITY)*QUANTITY EACH)*P
            MASS-QUANTITY)*QUANTITY SOME)
```

9.4 Interpretation of Quantified Expressions

In a recent thesis Van Lehn (1978) reports that when people are given a sentence like

A quick test confirmed that every drug was psychoactive.

they claim they understand it, but are then unable to state whether there was one test per drug or only one test for all. This ambiguity can be expressed in predicate calculus using the universal and existential quantifiers: Using the standard rule that an existentially quantified variable may only depend on the universally quantified variables to its left, we may write the two interpretations:

For all drugs$_i$ there exists a test$_j$ such that that test$_j$ of a drug$_i$ was positive.

There exists a test such that for all drug$_i$ that test of drug$_i$ was positive

The difficulty with this notation is that it lacks a form which leaves the ambiguity unresolved. If people were forced to resolve the ambiguity in order to come to some understanding of the sentence they would not respond to VanLehn as they do. From this we may surmise that people have some representation of sentences which proceeds what Chomsky calls translation into logical form. In this representation, predicates are applied to expressions like every drug, without it being determined whether or not there is a separate instance of the predicated conditions for each member. For example, all the boys in

I gave all the boys a cookie.

might be treated either collectively, using a single cookie, or distributively, using a cookie for each boy. We postulate the existence of a form of description where this is not resolved.

Even saying that a predicate applies either distributively or collectively once it is interpreted seems to be an over simplification. As Bierwisch (1971) notes, a sentence like

The boys hit the girls.

represents a statement whose truth would be supported by hit conditions existing between individual boys and individual girls. Indeed (Nunberg, 1975)

Have the boys ever hit the girls?

reguires the reporting of even a single instance. However, since we have no information about what boy/girl pairs are involved, there seems to be no sensible way to interpret this statement in terms of conditions on individuals without making arbitrary choices. Were all the boys and girls involved? Was any girl hit by more than one boy? We don't know. What is more, it seems likely that, in most contexts, not only would these details not be available, but they would not be needed in order to make whatever inferences are desired. If the details for interpretation are not available and a specific interpretation is not needed for making inferences, it seems the uninterpreted form would normally be the best representation to use for this sort of statement. Other example of this sort are:

The white oppress the negroes.
The Romans destroyed Carthage.
The Chinese of the seventh century know porcelain.

The question of what level of interpretation of such statements is required for inferencing would seem to be the same as that faced with a statement like

John broke the vase with the tongs.

On reading this we know that John controlled the tongs in causing the vase to break. But how, exactly, did he do it? Did he smash the vase with the tongs, did he squeeze too hard, did he pick the vase up with the tongs and smash it against the wall? We don't know. Inferencing may or may not require us to select one of these interpretations. We envision that the interpretation would involve the construction of a procedural specification on the lines of that suggested by Woods (1977). For example, to interpret

A requirement for the course is the carving of a block of wood into each of the 12 designs.

we must construct a procedure. Our reasoning might be "Well, Let's see. We take the wood and carve the first design. Oh! Oh! Now how do we carve a second design, the block is used up. "Well, we could have 12 blocks and carve one after the other, or we could cut up the block into 12 pieces and carve them, or maybe we could fit the 12 designs on one block."