American Journal of Computational Linguistics *Microfiche 71*

SPATIAL REFERENCE AND SEMANTIC NETS

Norman K. Sondheimer

Department of Computer and Information Science
The Ohio State University
Columbus, Ohio 43210

SUMMARY

This paper presents an analysis in a semantic net formalism of the semantic structure of English sentences containing references to spatial location. Spatial reference, hereafter SR, provides either static location or motional information

John is at home.
Fred ran across the street to the store.

The task for the semantic analysis of sentences with SR's is to make clear what is being positioned. This has been difficult to do. Previous proposals have left unanalyzed many phenomena including important motional references. This paper's main conclusion is that a much improved analysis can be obtained by representing the SR's as positioning abstract events and states of affairs.

The analysis in semantic nets has the location of an event or state of affairs represented as a node which is linked to the node showing the event or state by arcs indicating its status as the spatial attribute. A few SR's are shown as naming these locational entities, which we call place object. These SR's involve examples with "where", "here", and "there" However, most SR's are represented as relating place objects to the position of objects in the manner of prepositional phrases. This primacy of prepositions is argued for in the paper. Motional references are allowed for by functions represented in the nets which produce parts of place objects which are then positioned by prepositional forms. The necessary ordering that comes with motional references is allowed for by associating temporal elements with the functions.

While the positioned elements are simple, the overall semantic structure of the sentences containing SR's is often complicated by the involvement of more than one event or state of affairs. The paper includes a survey of the sentential semantic structures necessary to deal with SR's. A similar complexity is necessary to deal with the information on the location of objects which is gained from sentences with SR's. The paper suggests the use of inference rules to allow for this.

The most surprising of the paper's conclusions is that a strong tie exists between references to space and temporal information. In fact, the locations of all events and states of affairs placed by SR's are argued to be locations in both space and time. The effect of this conclusion is most clearly seen in a formalized definition of the primitives of the semantic structures, which is also presented in semantic nets. There, as one possible interpretation of the place object, it is shown as a set of pairs of volumes in space and points in time.

TABLE OF CONTENTS

I. Introduction

This paper presents an analysis in a semantic net formalism of the semantic structure of English sentences containing references to spatial location. Spatial reference, hereafter SR, provides either static location or motional information:

1.1 John is at home
1.2 Fred ran across the street to the store.

The task for the semantic analysis of sentences with SR's is to make clear what is being positioned. This has been difficult to do. Previous proposals have left unanalyzed many phenomena including important motional references. This paper's main conclusion is that a much improved analysis can be obtained by representing the SR's as positioning abstract events and states of affairs.

The analysis in semantic nets has the location of an event or state of affairs represented as a node which is linked to the node showing the event or state by arcs indicating its status as the spatial attribute. A few SR's are shown as naming these locational entities, which we call place object. These SR's involve examples with "where", "here", and "there". However, most SR's are represented as relating place objects to the position of objects in the manner of prepositional phrases. This primacy of prepositions is argued for in the paper. Motional references are allowed for by functions represented in the nets which produce parts of place objects which are then positioned by prepositional forms. The necessary ordering that comes with motional references is allowed for by associating temporal elements with the functions.

While the positioned elements are simple, the overall semantic structure of the sentences containing SR's is often complicated by the involvement of more than one event or state of affairs. The paper includes a survey of the sentential

semantic structures necessary to deal with SR's. A similar complexity is necessary to deal with the information on the location of objects which is gained from sentences with SR's. The paper suggests the use of inference rules to allow for this.

The most surprising of the paper's conclusions is that a strong tie exists between references to space and temporal information. In fact, the locations of all events and states of affairs placed by SR's are argued to be locations in both space and time. The effect of this conclusion is most clearly seen in a formalized definition of the primitives of the semantic structures, which is also presented in semantic nets. There, as one possible interpretation of the place object, it is shown as a set of pairs of volumes in space and points in time.

The paper has nine sections following this one. In the first, the limitations of previous analyses of the semantic function of SR's is considered. Then in one section, the semantic net formalism and, in the next, the syntactic distinctions used in the study are introduced. The next four sections present ever more complex situations. The first section shows simple direct analyses involving one event or state. The next section presents complex sentential structures with non-movement SR's. Motional references are analyzed in the next. The connection between time and SR's is discussed in the fourth section. Following these analyses of sentential semantic structures, a section is given over to the formalization of the definition of the structures used. The paper ends with a discussion of the limitations of the proposal and possible extensions to it.

There is available a discussion in greater detail of a preliminary analysis to the one given here (Sondheimer, 1975). There is also available for comparison an analysis by this author of the same meaning phenomena, in the competing paradigm of model-theoretic semantics (Sondheimer, 1978). The current

paper is distinguishable by its better developed semantic net formalism and its emphasis on producing computationally justified structures.

II. Previous Efforts

The past has seen many studies of SR phenomena. There has been interest in connecting language and scenes, e.g., Coles (1968), Kochen (1969), Winograd (1972), Badler (1975), and Tsotsos (1976). The use of language to capture the spatial structure of the physical world has been studied, e.g., Hobbs (1975) and Kuiper (197'). The conceptual structure of the terms used in SR and the pragmatics of evaluating them has been studied, e.g., Cooper (1968), Bennett (1975), and Denofsky (1976). Finally a number of studies have considered our topic: the position of a SR within the semantic structure of a sentence.

Studies of our sort tend to be distinguishable by the type of entities SR's are claimed to locate. In some cases, the SR's apply to only physical objects. In others, they apply to only abstract forms identifying events and states o affairs. A broad third type of analysis shows different sorts of entities being modified. Each has its limitations.

II.1 Analyses Using Physical Objects

The paradigmatic phenomenon for the analyses that claim physical objects as the referents of SR's is the noun phrase modifier:

2.1 The man in the car left.

The SR in the above is the phrase "in the car". The proposals of Norman and Rumelhart (1975), Abrahamson (1975), Geis (1975a, b, anc c) and Schubert (1976), among others, would try to show the relation of the SR to "the man" directly. Figure 2.1 illustrates the typical structure in the style of Schubert (1976). This figure shows "the man" being located (LOC) at a time, indicated by the T link, and at a location which was in "the car".

This style of analysis seems simple and direct. It appeals to the intuition that only physical objects take up space. It promises to be easy to apply,

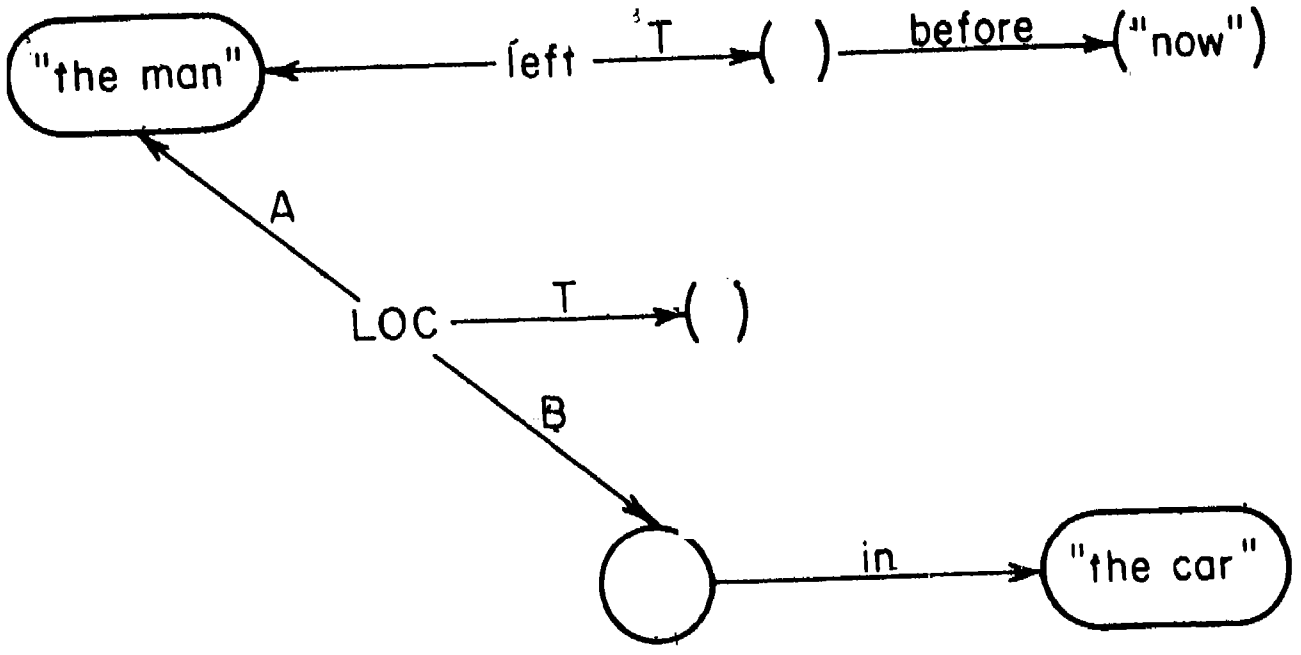

FIGURE 2.1 "The man in the car left" in the style of Schubert (1976).

since all that is required is to associate SR's with the sentential elements which are modified which reference physical objects. Unfortunately, there are problems.

It can be difficult to find all or any objects with which to associate an SR. Often there is more to an event than just its participants' locations:

2.2 John is playing solitaire in the basement.

If we hear example 2.2 then more than John is known to be in the basement. His cards are, for example. Further, the location of the action is more than the instantaneous position of John and his cards. For example, space where the cards may potentially be placed must be included. Similarly, the following does not indicate that John is next to the school:

2.3 John is playing baseball next to the school.

He might be playing outfield 300 feet from it. It can be difficult to find any objects to associate with an SR:

2.4 In France, literary criticism is a high art form.
2.5 The explosion was in the garage.

In both the above examples, only complex analyses showing many understood and potential participants can allow for object-reference.

Allowing for motional sentences is a very serious problem for object-reference analyses. The typical proposal is to show motion as change from one static location to another:

2.6 The man walked from New York to Chicago.

Example 2.6 would be shown as a man's walking causing a change of location from New York to Chicago. Some sentences show intermediate points:

2.7 The man walked from New York to Chicago via Pittsburgh.

Here, successive changes seem to be appropriate. However, one class of references to motion seems to defeat this entire approach:

2.8 The man walked across the puddle.
2.9 The man walked around the puddle.
2.10 The man walked through the puddle.

Examples like the above involve duration in a key way and can not be shown with reference to one position. For example, at no time was the man "across" the puddle like Raleigh's cloak was across it. Similarly, two points showing the man's change of position are inadequate since the same initial and final positions are acceptable in all three cases. Finally, adding an intermediate point will not be adequate, since the man might reach that point while on a path that otherwise holds a different relation to the puddle. As shall be seen, the lesson to be learned from these examples is that in allowing for motion, it is the entire path that must be considered and not selected positions of objects.

II.2 Analyses Using Events and States of Affairs

A second uniform type of analysis postulates events and states of affairs as the subject of SR's (see for example, Davidson, 1967, Lakoff, 1970, and Harman,

1972). Events and states of affairs are said to be the two types of situations that utterances describe. Taking them as the subjects of SR's claims that it is not the participants but the overall situation that is being referenced. This can be seen in Figure 2.2, which shows one of Davidson's analyses in a semantic net notation. The diagram shows that there is a strolling by John which has a particular time and space coordinate. The benefits of this analysis include the independence of event and state existence from discussion of spatial location, the ability to handle location of vaguely bounded events and states, and the simplicity of application. However, again the simple direct methods that have been proposed are unsatisfactory.

It is often difficult to simply associate SR's with a central event or state since SR's in some utterances must modify different entities:

2.11 John held the ice bag to his head in the car.

In 2.11, only the ice bag is to John's head but John and the ice bag are in the car. Motion is still a problem:

2.12 John walked from his car across the yard to the house.

How the event of 2.12 can be "from", "across", and "to" simultaneously and also have these aspects temporally ordered is nowhere explained in these analyses. Finally, even if SR's are associated with events and states of affairs, the fact that something is often learned about participants' location must be explained. For example, from the sentence of Figure 2.2, the fact that John was in the streets of Bologna is clear, but from the semantic structure only the location of the strolling is clear. No coherent way has been presented to allow for this kind of relationship.

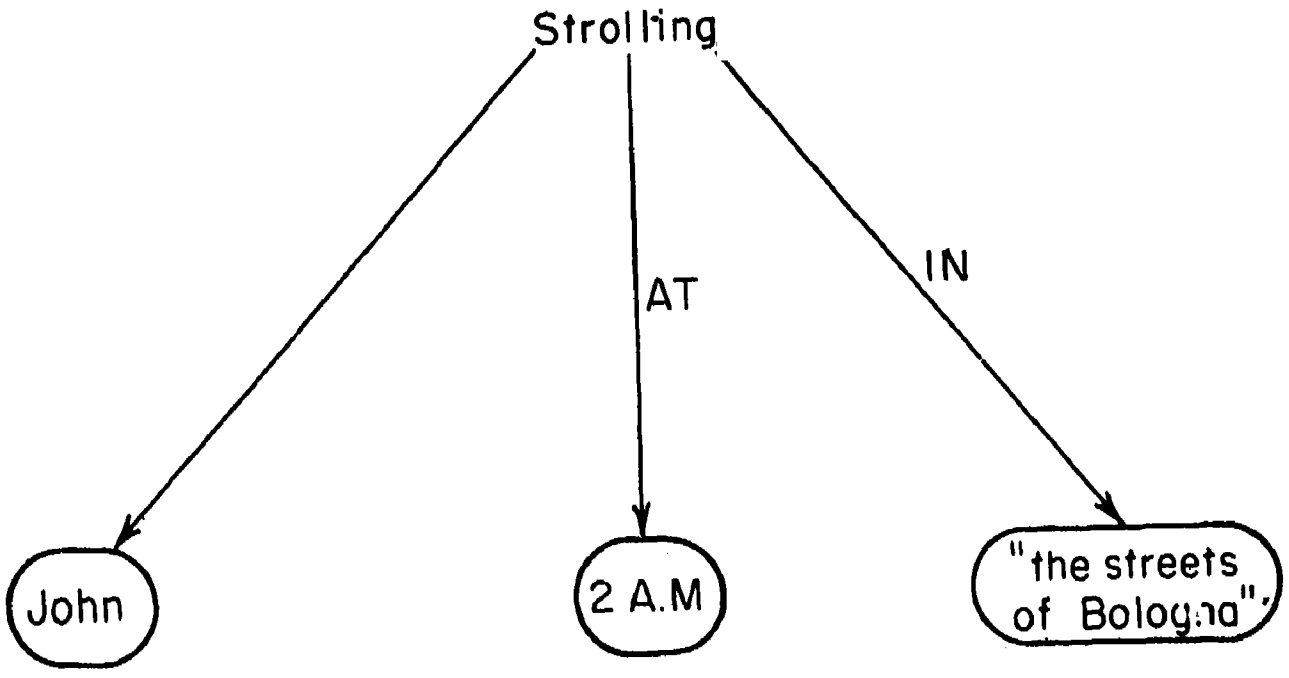

FIGURE 2.2 "John strolled through the streets of Bologna at 2 a.m." im the style of Davidson (1967).

IV.3 Nonuniform Analyses

The third style of SR analysis is nonuniform in nature. These either mix the two uniform analyses or elaborate on the simple event or state analysis. Mixed analyses claim that some SR's locate concrete objects while some locate events or states of affairs (see for example, Winograd, 1972, and Schank, 1973) By sacrificing the simplicity that comes from uniformity, these analyses avoid the uniform analyses' complementary problems. However, the mutual problems, especially motion, are left unsolved.

The nonuniform analyses that elaborate on the nature of events and states of affairs are best represented by Case analyses, see Bruce (1975). They claim that either the overall location or specific aspects of events and states are located. Taking Fillmore (1971), as opposed to the better known but earlier Fillmore (1968), as the model, four spatial cases can be seen. An SR can either reference a static location (the Location case), place of origin (the Source case), place of termination (the Goal case), or location of intermediate motion (the Path case). In terms of events and states of affairs, the first case can either be

used for overall event or state location or it may be used to locate an aspect of the event. The final three cases all relate to different aspects of a motional event. This allows for examples like 2.12, with inherent temporal ordering among the cases allowing for the ordering of the SR's.

The Case analyses still has problems. The two uses of the static case conflict in sentences with two static locations such as the one where the ice bag is held to the man's head while the man and ice bag are said to be in the car. Two instances of the Location case seem to be required, but if both appear, there is no way to identify their differing function. Also, motion is still troublesome. As Fillmore (1971) points out, instances such as the underlined phrases in the following seem to indicate a need for an unbounded number of instances of the Path case:

2.13 He walked down the hill across the bridge through the pasture to the chapel.

The underlined phrases refer to motion ordered in time, e.g., he walked the hill before the bridge. However, Case analysis gives no way to order instances of the same case. Gruber (1965) points out the same problem with the Goal case:

2.14 I walked to New York to my mother's.

Finally, the Case proposal must be given some physical interpretation. Any representation of meaning must at some point be related to a model of the world. In this instance the idea of a source, goal, and path must be somehow related to models of motion.

This paper presents a proposal for an analysis that is nonuniform in the same way the Case analysis is. A uniform source for locations modified by SR's is given, but the predication of these spaces by SR's is shown to be much more complex than previously thought. Further, sentences are not seen as being as

simple with respect to SR's as previously supposed. Before presenting the analysis, two sections will be devoted to preliminary topics: our semantic net formalism and the syntactic status of the phenomena considered.

III. Semantic Nets

The results of this analysis of SR phenomenaare formalized in semantic nets or networks. This formalism is currently a popular choice for semantic analyses. It allows clear, expressive graphic presentations and possesses many positive computational properties. Because of its popularity, it also allows wide dissemination of ideas. Working against this last claim is the proliferation of versions of the representation, for example, Hendrix (1976), Norman and Rumelhart (1975), Shapiro (1971), Simmons (1973), and Woods (1975). This section clarifies what is meant here by the formalism, which can be seen to most closely resemble that of Brachman (1977).

Semantic nets have been used for representing many aspects of intelligence. Often they are used to represent factual information concerning objects, actions, and states. They have separately been used to show the semantic structure of utterances It is this use that mainly interests us. However, there is a connection between the two uses. All semantic structures must be related to structures that represent facutal information and each use of a type of object, event, or state of affair must be related to a concept that explains it. This can be thought of as paralleling the relationship between a semantic structure shown in the predrcate calculus and a model in which that structure has a truth value.

In a complete net, the above translates into the necessity of nodes for concepts representing types of events, states of affairs, and objects and nodes for instances of tokens of these concepts. The "token" nodes must link to "type" nodes that define them. These definitions must include specification of attributes of an instance in terms of restrictions on values, functional role of the attribute, and other things. The instance nodes must be connected to instantiations of the attributes. Concept nodes must also be related to other concepts, have overall structural conditions, locate inference rules that may apply, etc. All this information is essential to any artificially intelligent entity, just as

the model is essential to any analysis in the predicate calculus. However, for showing the semantic relations in which we are mainly interested, an abbreviation is sufficient just as only the formulas are sufficient in most studies using symbolic logic. Hence a special abbreviation will be used in all sections except IX where the definitional level will be discussed.

Central to our abbreviation will be nodes that collapse types and tokens. These will identify the verbal concepts that characterize the events and states of affairs. We will call them "event/state" nodes. They will be circled and capital letters will be used for abstract types, such as CAUSING. Nonabstract forms will be shown with names that suggest the interpretation, e.g., Sleeping will suggest the sleeping state. When a node represents a physical object, identifying information will be included in quotes, e.g., "the bus". Names placed on arcs will abbreviate and suggest the functional roles of attributes. For example, ANTE for antecedent and CONS for consequence will be used with CAUSING. Case names will be used with many event and state of affairs types. These will include:

T for "Time" showing the time an event occurred or state held.
A for "Agent" showing the instigator of an event or state.
O for "Object", the neutral case (as Fillmore (1971) explains it "the wastebasket").

Restrictions on types of entities which will be necessary will be shown by non-oval shapes for nodes. For example, time instances will be shown in parentheses and time intervals in square brackets. Finally, because it is not essential for our purposes, specification of time will often be left out of most semantic structures. Similarly, we will consider only declarative statements. Figure 3.1 shows a typical structure.

Some concepts that act as functions will also be used. Each of these will look like a relation associating parameters with a value. The value will be identified by a VALUE arc. Inference rules will be presented in the form of

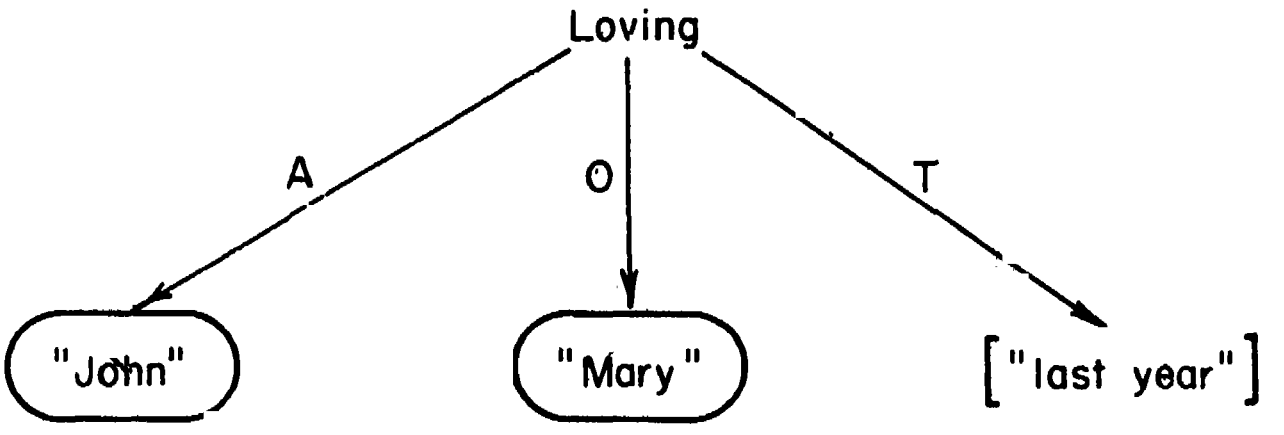

FIGURE 3.1 "John loved Mary all last year."

"subnet$_1$" $\Rightarrow$ "subnet$_2$", where on seeing subnet$_1$, subnet$_2$ is to be added to the semantic net. These rules will include variables within nodes, where the variables are to be bound on matching and referenced on inferencing. These variables will be in the form of capital letters, e.g., X.

To summarize, our semantic net formalism uses concept names, descriptions of objects, mnemonic arc names, and mnemonic shapes for nodes to abbreviate the two levels in a semantic net. Also used are functions and inference rules. This will be enough to represent the semantic relations involving reference to space that are being considered. Unfortunately, it is one more unique formalism. However, it adds no new structures, only abbreviating others. We leave as an unproven claim that it will fit in with any formalism which shows identifiable event and state of affairs nodes such as Norman and Rumelhart (1975) and Schank (1973).

IV. The Syntactic Structure of Spatial References and the Primacy of Prepositions

Semantic structure is the topic of our paper, but the syntactic structure of sentences with SR s is also important. Its consideration clarifies the range of phenomena being studied. With SR, the basic syntactic structures involve prepositional phrases. All other SR are analyzable in terms of these structures. In this section, the syntax function of prepositional phrases will be considered and arguments for their primacy will be presented.

Our main interest in syntax is in structuring our discussion of semantics. However, the problems of parsing and generation make the syntax of SR's independently important. These are not our topics here. However, in an earlier issue of this journal we presented a parsing scheme that produces semantic from syntactic structure and applied the scheme to current class of phenomena (Sondheimer and Perry, 1975).

IV.1 Locative Prepositions

Prepositional phrases that express SR's can be called locative. They appear contiguously, as in example 4.1, or discontiguously, as in example 4.2:

4.1 I put it on table.
4.2 The table I put it on is broken.

The discontiguous example can be taken as derivable from (reducible to) the contiguous forms in generation (interpretation). Hence only contiguous examples will be considered. These are primarily employed in four syntactic roles: complement, qualifier, adjunct, and locative object. There is also one special dependent usage that will be described at the end of this section.

The complement usage of locative prepositions arises only when they are the "complement" of the verb "be":

4.3 He is in the kitchen.

Quirk et al., (1972) distinguishes them from predicate adjective and nominal

usages. Locative objects and adjuncts with copulative sentences can be distinguished from complements by the presence of these adjectives and noun phrases:

4.4 There are lions in Africa.
4.5 He was important in Chicago.

The qualifier usage of locative prepositions is part of noun phrases and shows the location of the reference of the noun phrase:

4.6 The man in the car left.

The strings in some sentences may make it appear that locative prepositions are part of noun phrases when they are not:

4.7 I put the knife on the table.
4.8 She took care of John in Chicago.

In these cases, the passive test and cleft-sentence test (Jacobs and Rosenbaum, 1968, p. 38) can be applied:

4.9 *The knife on the table was put by me.
4.10 *What I put was he knife on the table.
4.11 *John in Chicago was taken care of by her.
4.12 *What she took care of was John in Chicago.

The asterisk "*" here and throughout marks ungrammatical sentences. The ungrammaticality of the above examples indicate that the strings in question are not noun phrases. Hence the prepositional phrases cannot be qualifiers.

Adjunct usages are prepositional phrases that are external to the clause of a sentence:

4.13 I met John on the train.

Locative object usages are objects of verbs and internal to clauses:

4.14 I put the lamp in the corner.
4.15 He yelled at John.
4.16 He saw her in the park.

There is some controversy on the distinction between these two types. We can present two syntactic and one semantic classification procedures. First, adjuncts are never required for grammaticality, while locative objects can be:

4.17 I met John.
4.18 *I put the lamp.

Second, adjuncts always allow shifting to presubject position without loss of grammaticality or shift in meaning:

4.19 On the train I met John.
4.20 *At John, he yelled.
4.21 In the park, he saw her.

Note that in 4.21, the man is definitely placed while in one interpretation of 4.16, the locative object one, he is not.

Semantically, we claim that adjuncts locate the entirety of events and states discussed, while locative objects can locate only part of what is described. For example, in 4.16, the locative object reading shows only the woman's position in the park, not the location of the "seeing" as a whole. The following is also informative:

4.22 He dropped it behind the door.
4.23 Behind the door, he dropped it.

Both examples are similarly ambiguous with respect to the SR. One sense, the most likely to be identified in 4.22, is that the end result of the dropping was that the object came to be behind the door. The second sense, the most likely for 4.23, is that the dropping took place behind the door. The first sense shows partial predication and a locative object usage. The second shows overall predication and an adjunct usage.

Some forms that seem to be adjuncts do not at first glance appear to make overall predication:

4.24 On the train, he commented on the Empire State Building.
4.25 In Chicago, John wrote to his mother.

The Empire State Building's and John's mother's position are independent of the train and Chicago. However, we can claim there is still overall predication since the commenting and the writing were done on the train and in Chicago, respectively.

Durational adjuncts also complicate the semantic test:

4.26 He cried through the tunnel.
4.27 He sat still from New York to Chicago.

These prepositional forms show duration of the crying and sitting and should be taken as adjuncts. The first gives overall predication. The second example shows two phrases that individually give partial predication. However, together they give overall predication. Further, they cannot be used individually:

4.28 *He sat still from New York.
4.29 *He sat still to Chicago.

These are the four primary uses of locative prepositions. We claim that the semantic structure of other SR's can be represented through these forms. We will now show this. In general, this will be done by observing the SR's structure or by paraphrase arguments.

IV.2 Other Spatial References as Locative Prepositions

Some spatial terms can have syntactic and semantic functions similar to prepositions in that they directly serve to relate two forms:

4.30 San Francisco is north of Los Angeles.
4.31 The car is to the left of the building.

These examples can immediately be given prepositional-like semantic structures. In other sentences, these terms appear as nouns and adjectives:

4.32 The North is desolate.
4.33 He hit my left leg.

Here, the forms can be paraphrased in the prepositional-like form which can be taken as their underlying semantic form:

4.34 The part of the country to the north of the rest is desolate.
4.35 He hit one of my legs that is to the left of the other.

Another category, the locative prepositional adverbs, although lacking syntactic objects have assumed semantic objects. This is shown by our ability to question the missing object, which is a means for distinguishing this category from verb

particles (Quirk et al., 1972, p. 103):

4.36 He went up.
4.37 Up what did he go?
4.38 He picked it up.
4.39 Up what did he pick it?

A diverse variety of non-prepositional locative adverbs can be handled with prepositional forms. Assumed objects can also be seen in cases of paired prepositional-adverbs and prepositions. These are suggested in parentheses below:

4.40 He walked across (a walkable space) to the blackboard.
4.41 He jumped from (a jumpable place which was on) the table.

Some adverbs can be straightforwardly treated as the equivalent of prepositional phrases. These appear as the concatenation of a preposition and a noun and refer to the spatial relation referenced by the preposition with respect to the type of object referenced by the noun:

4.42 He ran uphill.
4.43 He is overseas.

The suffix "-ward" following a preposition or preposition-like term produces an adverb that can be treated as having a destinational- or orientational-like meaning as shown by the following paraphrases:

4.44 He moved leftward.
4.45 He moved to the left.
4.46 It pointed upward.
4.47 It pointed up the space.

Other adverbs can be treated as having a neutral prepositional sense like "at" or "to" in their semantic representation:

4.48 He is home.
4.49 He is at home.

Finally, many noun phrases that indicate position can be seen as having prepositions subsumed by the verbs they appear with and hence can be represented as containing prepositional phrases, see Gruber (1965) for elaboration:

4.50 He gave Susan the ball.
4.51 He gave the ball to Susan.
4.52 He jumped the fence.
4.53 He jumped over the fence.

There are a few forms in SR's that I can not always claim to be represented by prepositional forms. These are "where", "here", "there", and measures of distance. These will be dealt with separately. In general, we will deal with prepositional phrases with the assumption that all SR phenomena are covered.

Beyond the examples already given, it is hard to say what should be considered an SR. Adjectives such as "long" and "fat" involve the abstract properties of objects more than their properties as objects momentarily situated at a point in space. Many examples appear to be metaphors of SR or make oblique reference to space:

4.54 I stood trial.
4.55 I go to Ohio State.

All of these will be ignored. Doubtlessly, there are unarguable cases of SR that are being overlooked. For this, I can only apologize.

IV.3 Semantic Structure of Prepositional Phrases

Since prepositional forms are the basic method of making SR, their representation is central to this analysis. They will be given a semantic representation as concepts relating what is referenced by the SR to their own complements, see Figure 4.1. The referenced entity will be identified by the F link for "figure" and the complement by the G link for "ground" (Talmy, 1975)*. Each prepositional concept will be defined as comparing the figure's space to the location of the ground's object at the time associated with the figure (Section

*There would have to be a second ground link for "between":
I left it between the window and the door.

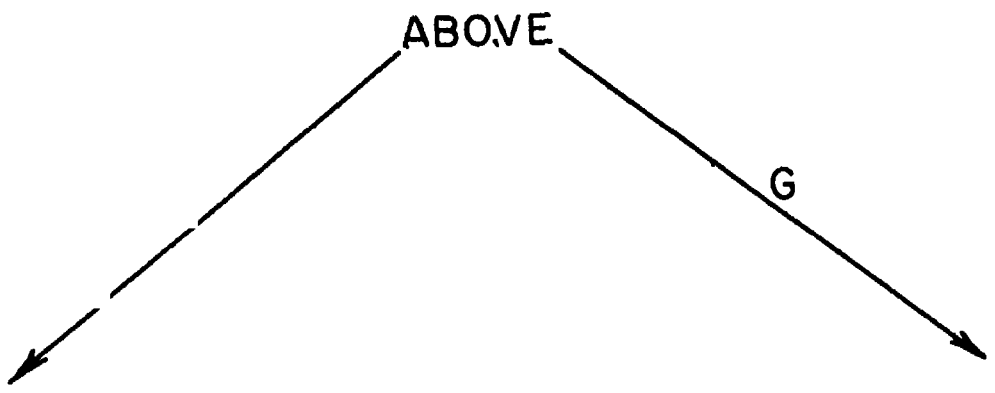

FIGURE 4.1 A prepositional semantic structure.

VIII contains more discussion on this point). Prepositional concepts will all be considered abstract and written in capital letters. The reason abstract forms are used will become clear in the following sections.

One particular dependent use of the preposition "from" falls outside the simple pattern shown in Figure 4.1 as well as outside of the four classes of prepositional usages:

4.56 John is far from home.
4.57 John is across the street from home.

In both of the above, John is distant from home. But in 4.57, John is not across the street in the usual sense of "across" stretching the width of the street. As Bennett (1975) points out, the "across" and "from" phrases combine in such a way that we understand that it is the way that must be travelled in starting from home and going to John that is "across the street".* This can be allowed for in semantic nets with a function, WAY, producing a path through space joining two points identified by INIT for Initial and FIN for Final links, see Figure 4.2. For example 4.57, G would identify the street, INIT the home, and FIN where John

*The same meaning also arises in sentences such as the following where there is an understood "from point" that must be represented:

He died across the river.

is. How John's location is to be shown is explained in the next section, where the basic and simpler SR's are analyzed.

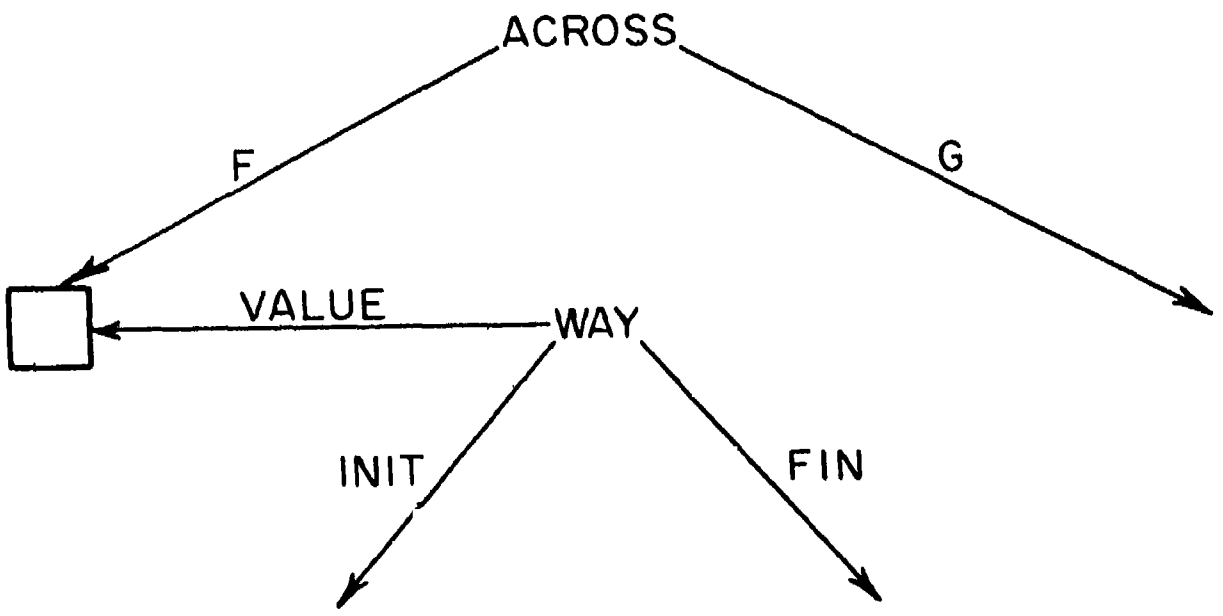

FIGURE 4.2 A prepositional semantic structure for the special "from" usage.

V. Static Adjunct, Complement and Qualifier Usages

Section I describes our basic claim: the source of the locations being referenced by SR's can be represented as being the locations of events and states of affairs. In this section, this claim is associated with the semantic net model and applied to those types of SR's for which it works immediately. These are the static adjunct, complement and qualifier usages.

V.1 Basic Structures

In our semantic net model, the locations of events and states of affairs will be shown as attributes of event/state nodes through arcs leading from the nodes to locational entities. For each event/state node involved with an SR there will be only one such arc and locational entity. These arcs will be labelled P to suggest a spatial attribute or "Place" case. The locational entities will be referred to as place objects. They are the basis of our analysis. These place objects can be taken for the time being as volumes in space. The sort of volume they are will be elaborated upon. Place objects will be identified by boxes. Figure 5.1 gives a typical diagram.

It must be asked whether place objects are required in semantic representations or simply ad hoc creations. The answer is that they are required since speakers treat them as existing by referring directly to them with some uses of "where":

5.1 Where is John living?
5.2 I found it where John was sleeping.

Place objects can not be outlined strictly in space like a solid can. This is not important, because there is no way in language to directly and completely locate any object. In the last section, it was argued that except for "where", "here", and "there", every SR is like a preposition. Hence they all give relative position. With those that do not, "where" can be shown

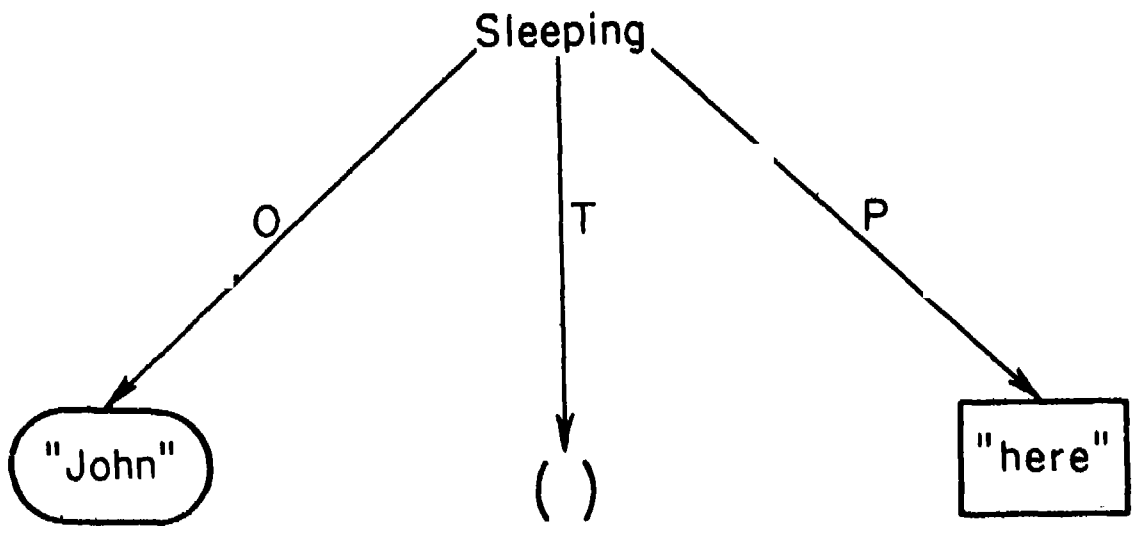

FIGURE 5.1 "John is sleeping here."

as referencing place objects not definite locations. "Here" and "there" both predicate spatial qualities of place objects not specific locations:

5.3 John was born here.

In 5.3, the location of the doing is simply associated with "here". Hence a semantic analysis that associates SR with abstract locations can work if the means of predicating these locations and of fitting them into semantic structures can be found.

V.2 Applying the Analysis

With the place object, there is a large class of phenomena that can be represented directly. These include static adjuncts (5.4) as opposed to durational ones (5.5):

5.4 At the table, John sat without moving.
5.5 From Dallas to Houston, John sat without moving.

Similarly, the static complement senses (5.6) as opposed to resultive complement senses (5.7) can be directly represented:

5.6 Chicago is far away from New York.
5.7 We are finally far away from New York.

Finally, direct analysis can be given to qualifier usages which either apply with a static sense to nouns describing physical objects (5.8) or act like static

adjuncts with respect to verbal nouns (5.9):

5.8 The man in the car left.
5.9 Swimming in the lake is fun.

These qualifier usages can be contrasted with those that show motion (5.10), act like locative objects to verbal nouns (5.11) or show extent (5.12):

5.10 The bus to Chicago left.
5.11 Swimming into a cave is fun.
5.12 The bridge from Ohio to West Virginia is old.

Applying the place object analysis to static adjuncts is easily defendable. One test for adjuncts in the last section was to see if it located the entirety of the event or state discussed. The static adjuncts are identifiable in this way. Since the place object shows the location of that entirety, static adjuncts can therefore be directly applied to them. Figure 5.2 gives a typical analysis.

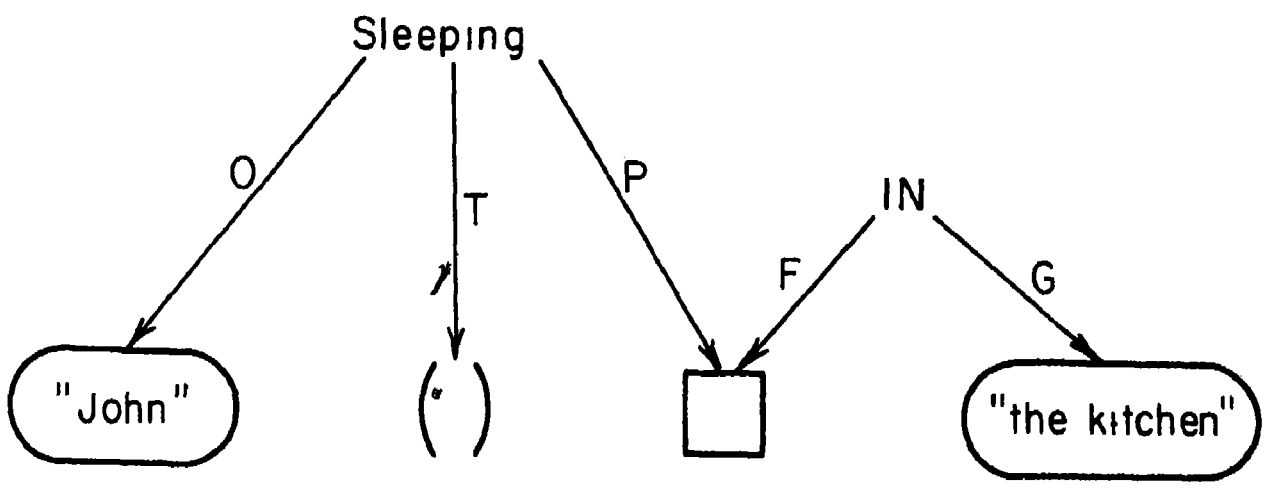

FIGURE 5.2 "John is sleeping in the kitchen."

This basic treatment extends to static complement usages. These relate an object to some location in space and time. To show this an abstract predicate, BEING-AT, can be postulated whose object case shows an entity whose spatio-temporal location is specified by Place and Time cases, see Figure 5.3.

Proposing a state of affairs to show an object's existence in space and time may at first seem artificial. But in fact, it provides representations isomorphic to the usual "direct" representation of object location. For example, Schubert (1976) uses a concept LOC which by a link A identifies an object, a link B the object s location, and a link T its time frame (see Figure 2.1). These match our BEING-AT, O, P, and T cases, respectively. Schubert sometimes abbreviates SR's when the preposition "at" is used. However, this is simply an abbreviation and his underlying form remains equivalent to ours.

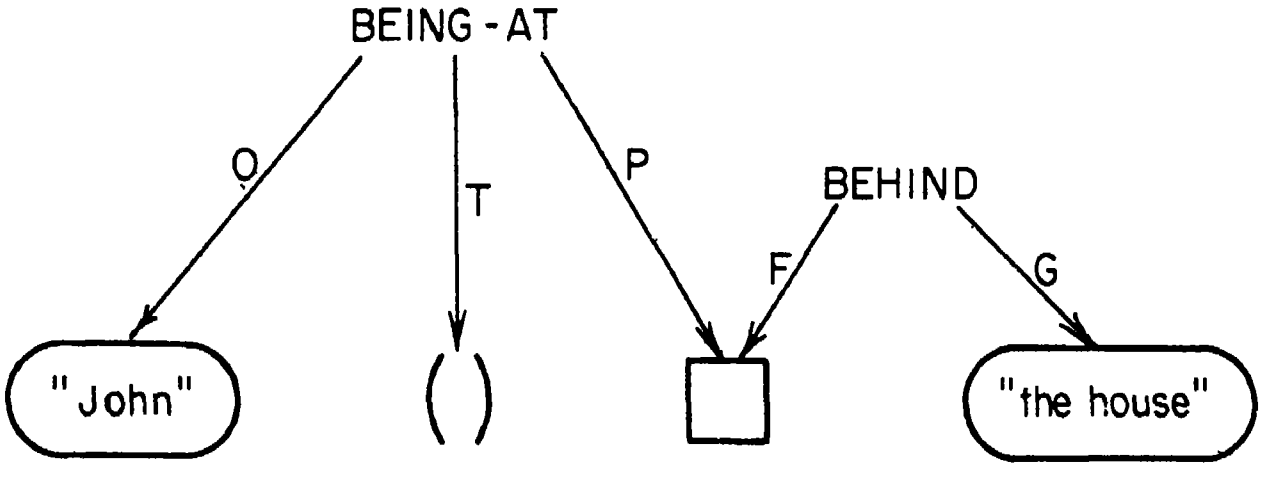

FIGURE 5.3 "John is behind the house."

Static qualifiers parallel either the adjunct or complement analysis. As an adjunct to verbal noun, we can claim that an event/state corresponding to the event or state described by the nouns can be located by the SR in the same way as an actual adjunct. With qualifiers applying to concrete nouns there can be a BEING-AT event/state showing that the existence in space and time of the object is being discussed. The qualifier can then modify its place object. This would then show thefollowing equivalently:

5.13 The man in the car yesterday left.
5.14 The man who was in the car yesterday left.

The possibility of time modification as in 5.13 is good evidence for the treatment of qualifiers as having underlying complement structure (Winograd, 1972).

V.3 Allowing for Object Position

Now that both adjunct and complement usages have been considered, our method of allowing for the positioning of objects while representing SR's as positioning event/states can be explained. As was discussed in Section II, an event/state analysis must explain how an artificially intelligent entity can discover that John was somewhere from the representation of an event or state involving John being located there. This can be taken as being something like discovering the appropriateness of the complement form (5.16) from the truth of the adjunct form (5.15):

5.15 John slept in the kitchen.
5.16 John was in the kitchen.

Within the computational paradigm, the discovery of 5.16 from 5.15 is made easy by inference rules. Whenever the semantic analyses of a sentence like 5.15 is presented to a system, rules associated with the type of event/state node involved can produce inferable information. This process allows for the human process of the deduction of specific information about participants in an event

or state of affairs from knowledge of the type of event or state of affairs. This is actually what is happening with SR's. From our knowledge of sleeping, we know that someone is where he is sleeping. From our knowledge of "working for", we know that Bill but not necessarily John is at the store in the following:

5.17 Bill is working for John at the store.

From our knowledge of contact cases such as in 5.18, we know that the location of the intersections of the objects is learnable:

5.18 The ball hit Mary on the ear.

In semantic nets, these facts can be shown by inference rules associated with the appropriate concepts. In Figure 5.4, the "subnet$_1$" ⇒ "subnet$_2$" form described in Section III is used to allow for the sleeping case. Other rules will, of course, be needed for other concepts.* The predication of place objects, which are the locations of events and states of affairs, therefore stands as the core of our analysis. How it directly applies to represent certain SR's has been shown in this section. In the next, more indirect analyses are considered.

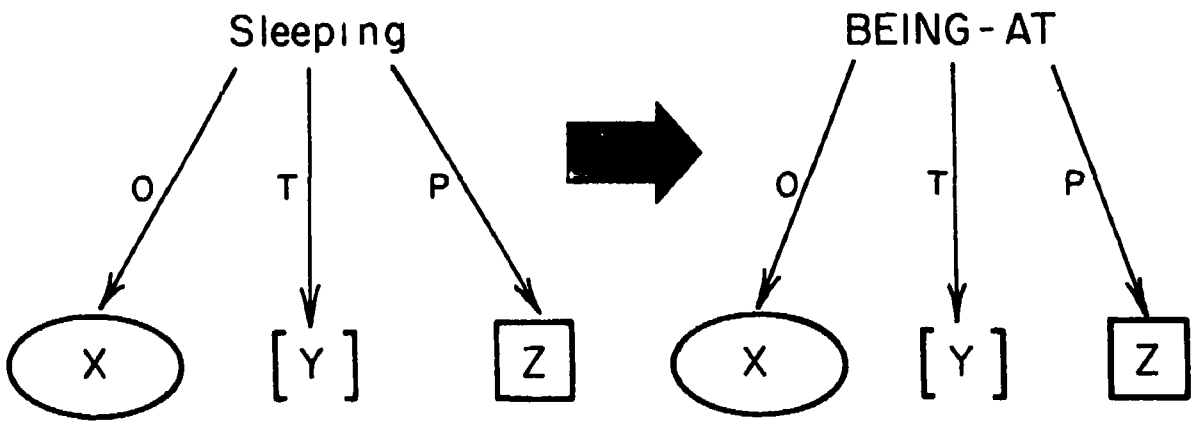

FIGURE 5.4 The encoding of the rule "If you know where something is sleeping, then you know where it is".

*A potential critic may argue that the extra processing involved with inference rules should be avoided if at all possible. However, no other analysis of SR successfully avoids its use (Sondheimer, 1975).

VI. Non-Movement Locative Object Usages

The analysis of locative object usages is not as simple as that of other forms. Looking back to Section IV, most of the problems with earlier studies arise from this class. The solution to these problems is found in an elaboration on the basic form of our event/state analysis. This elaboration proceeds in two directions. First, the semantic structure of the sentences containing the troublesome SR's is seen to be more complex than otherwise thought. Second, the nature of the SR is seen as more complex. The first case is best seen with non-movement and the second with movement SR's. This section covers the non-movement type of locative object. Section VII covers the movement type.

We can review the problems with the use of event and state location in the non-movement cases, briefly. There is a need to differentiate referents which can be seen in the following:

6.1 John held the ice bag to his head in the car.

The ice bag is to John's head, but both the ice bag and John are in the car. The first SR involves a locative object, the second an adjunct. With a simple approach to event/state location, they would not be differentiated. There is a similar problem in some adjunct references to the location of only part of an event or state of affairs. For instance, in 6.2, only the boy is placed which the hawk is definitely physically present:

6.2 In an open field, a boy watched a hawk.

In the latter case, although not the former, the use of inference rules might be suggested. However, a better answer can be found.

VI.1 Continuous Position and Perception Verbs

The semantic structure of simple sentences have often been analyzed as involving multiple events and states of affairs, see for example, Schank (1973)

and Norman and Rumelhart (1975). If we can see problematic sentences in this light, then perhaps we could assign the various SR's to different event/states. Indeed, we can do both.

Instances of causative relations between events and states of affairs are found in many problem sentences. Change-of-state events applying to separate states of affairs are seen in others. Simple instances of embedded events and states of affairs are seen in yet others. "Hold" belongs to a class of verbs that involve continuous position. Others in the class include "adhere", "cling", and "keep". With locative objects, these can all be seen as causations. Each has an action which causes some entity to remain somewhere. In our example 6.1, John's holding-type action causes the ice bag to remain somewhere. Realizing this allows us to analyze the SR's as locating events and states. The overall SR, "in the car", can be seen, as adjuncts were explained in the last section, as locating the highest event/state within the causation. The "to his head" can be seen as locating the resultant state. This is shown in Figure 6.1. The TO in the diagram represents a static sense of "to".

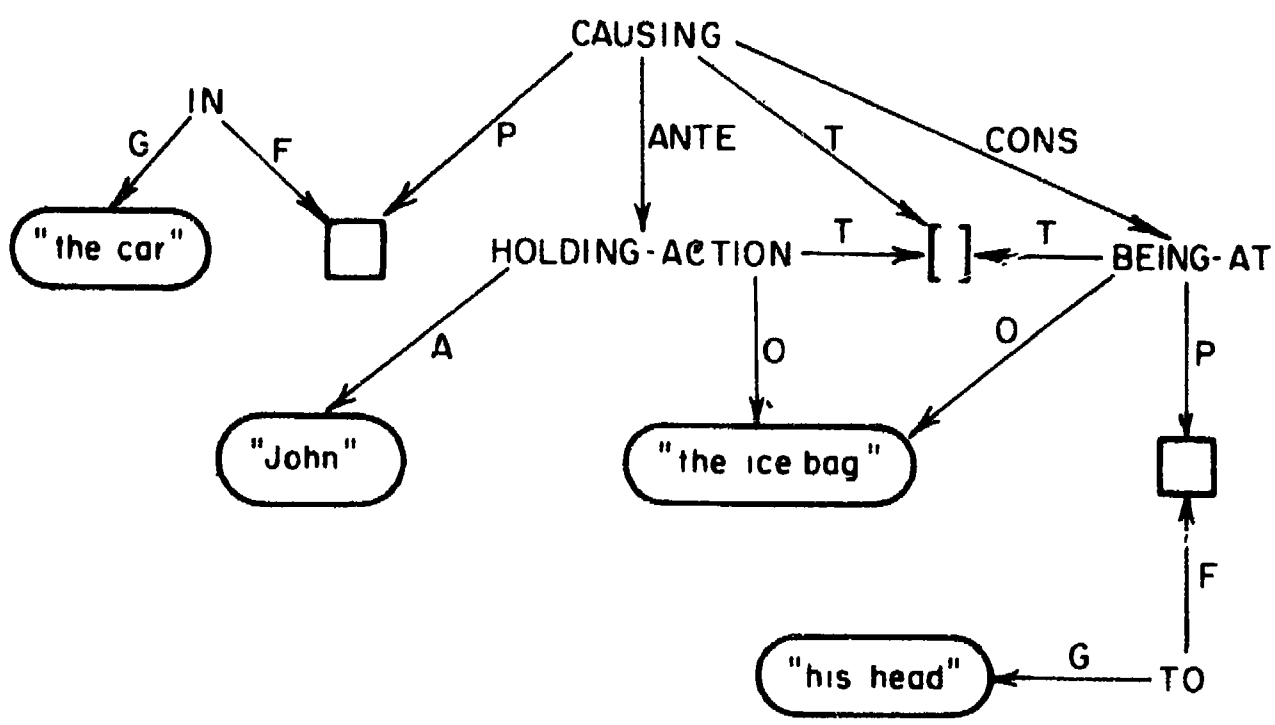

FIGURE 6.1 "John held the ice bag to his head in the car."

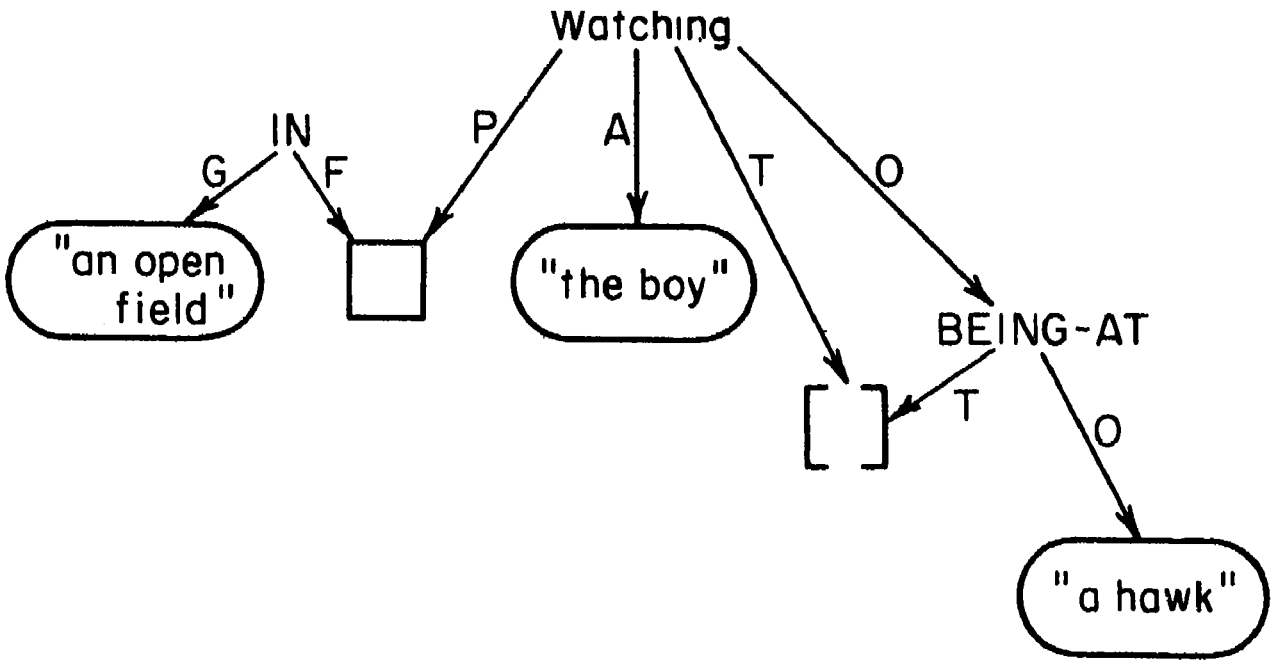

FIGURE 6.2 "In an open field, the boy watched a hawk."

"Watch" belongs to a class of verbs that includes "hear", "see", and "taste". These can all be seen as involving the perception of another event or state of affairs.* In our example 6.2, it is the being somewhere, the existing, of a hawk that is watched. How this allows for the SR to be associated with the correct event/state is evident from Figure 6.2. This analysis may seem somewhat forced here, but other exmaples show more overt event or state forms:

6.3 I watched the mating of the doves
6.4 I saw the delivery of the baby.
6.5 I heard the cooing of the doves.

*These verbs occasionally appear without an object:

I heard through the door.

On these occasions, an assumed entity can be added to the semantic structure:

I heard (something) through the door.

Inference rules play an important part in these analyses. For example, the positioning of John and the ice bag must be derivable from the structure of Figure 6.1. An inference rule must associate the position of the HOLDING-ACTION with their positions. Another rule must relate a place object for the HOLDING-ACTION as inside that of the CAUSING. Conversely, there should be no inference rule applying to the structure of Figure 6.2 to show the place object of the BEING-AT as being contained in that of the Watching.

VI.2 Other Verb Classes

There are a number of other classes of verbs that take static locative objects, see Table 1. We will survey their analysis in the remainder of the section and close with a comment on several related forms.

TABLE 1 - A CATEGORIZATION OF SOME VERBS THAT ACCEPT NON-MOVEMENT LOCATIVE OBJECTS

1. Continuous Position: adhere, cling, hide, hold, keep.

2. Perception: hear, see, taste.

3. Attachment, Containment, Posture, and Creation: build, close, crouch, draw, erect, glue, hang, lay, lean, lock, nail, paint, sew, shut, sit, stand, write.

4. Contact: grab, hit, kick, kiss, kneel, punch, slap, slug, touch.

5. Change of State: break, chop, cook, cut, fry, shatter, spill, split.

6. Discovery and Thought: dream, find, imagine, lose, recognize, remember, spot, think.

7. Copula-like: happen, gave, occur, remain, stay, take place.

8. Portability: bring, carry, send, take, wear.

The next class of verbs adds another abstract predicate to the set of forms we have considered:

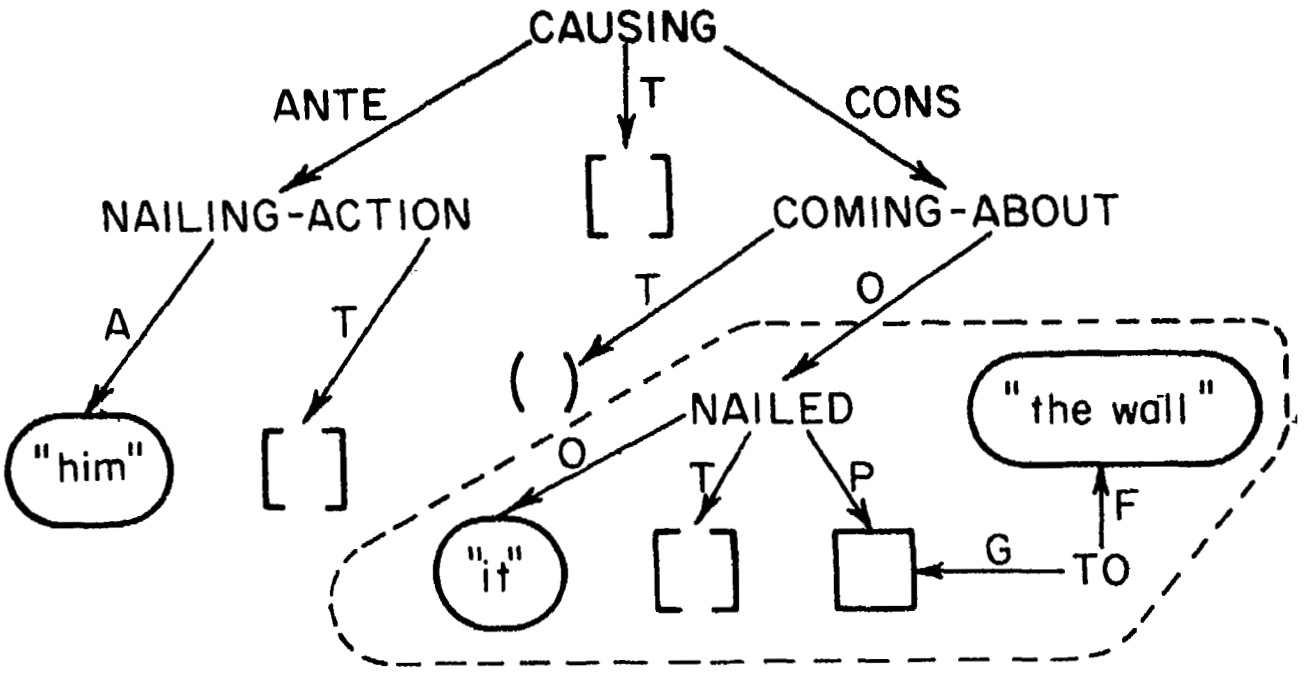

FIGURE 6.3 "He nailed it to the wall."

6.6 He nailed it to the wall.
6.7 He shut it in the room.
6.8 He sat it on the table.
6.9 He drew it on a napkin.

The above sentences involve attachment, containment, posture, and creation. Each has an element of coming-into-being that must be represented. The standard form for these sentences shows the action of an agent causing the bringing about of a state of affairs. The locative object is shown locating this state of affairs, see Figure 6.3. The coming-into-being concept in this structure is labelled COMING-ABOUT. The segment of the structure inside the dotted line is there to show the analysis these verbs take in the second type of usage they allow:

6.10 It is nailed to the wall.
6.11 It is shut in the room.
6.12 It sits on the table.
6.13 It is drawn on a napkin.

These examples lack agents and any sort of causation. The forms within the dotted line in Figure 6.3 show exactly this structure.

The prepositional form, TO, in Figure 6.3, is to be understood in the static sense just as with Figure 6.1. In fact, this is the case with all prepositional forms used here. It is an important advantage of this analysis that it uses only static senses in semantic structures. On the surface, it is often said that the locative objects of the current set of verbs have dynamic senses. However, with

a separate inchoative event/state, this is unnecessary. This allows the representation of presuppositions like "to" or "on to" either through "at" or "on" as Gruber (1965) does, or through their own static sense as in example 6.10. This is one in a series of reductions. It was shown in Section IV that some double prepositional phrase structures involved "from" can be reduced to a simpler form. It will be seen elsewhere that other simplications can be made. That underlying senses of the prepositions are being used explains why our prepositional concepts have been capitalized.

Another class of verbs shows contact. They take the two types of analyses just discussed. They also show a coming-into-being sense when no agent is present but a state is achieved. All three cases are shown in the following:

6.14 I touched her on her face with my hand.
6.15 The tree touches the window near the top.
6.16 The ball touched my leg near the knee.

The semantic structures for each of these can contain an event/state showing contact between the two objects to which the SR's can be applied.

Another class of verbs which show change of state have all three types of structures with locative object usages:

6.17 I broke it on the rim.
6.18 The cup broke on the rim.
6.19 The cup is broken on the rim.

6.17 is causative/coming-into-being, 6.19 is coming-into-being, and 6.19 is static only. The static form in each can again take the SR. In noncausative examples with these change of state verbs, SR's generally appear to act as locative objects, as a test from Section IV shows:

6.20 The cup broke on my knee.
6.21 *On my knee, the cup broke.

The noncausative examples includes SR's which reference objects not inherently possessed by the changed object, such as 6.20, but which place the entire event.

In these cases, the SR's should be treated similarly to adjuncts and shown applying to the COMING-ABOUT event/state.

Some verbs which take locative objects are like perception verbs in not requiring causative analyses to explain locative object usages. These include discovery and thought verbs, such as "spot" and "thought". They can be shown with embedded event/states. With locative object readings, 6.22 and 6.23 involve only locating of the direct objects:

6.22 I spotted her behind the dresser.
6.23 I thought of Mary at the seashore.

These entities can be shown in an event/state claiming they existed in a certain time and space with the SR predicating that event/state. This treatment would parallel the structure of overt examples of embedded events as in the following:

6.24 I spotted you stealing some bananas.
6.25 I thought of you dancing.

Finally, as Lyons (1968) notes, some verbs, which we treat here as having locative objects, seem to relate to SR's in the same way as the complement usages:

6.26 It occurred in Chicago.
6.27 It remained in New Orleans.

These can be analyzed with one event/state showing both adjunct and locative objects identifying the same entity.

So we have seen that the complexity associated with many SR's comes from their semantic environment, not themselves. With the exception of a class of verbs covered in the next section this covers the range of verbs that take non-movement locative objects. Also covered but only indirectly are a few senses left from the last chapter. We can now see how qualifiers of verbal nouns that are acting as non-movement locative objects can be analyzed. We can also see that resultive senses of the copula can be shown with a COMING-ABOUT, see Figure 6.4.

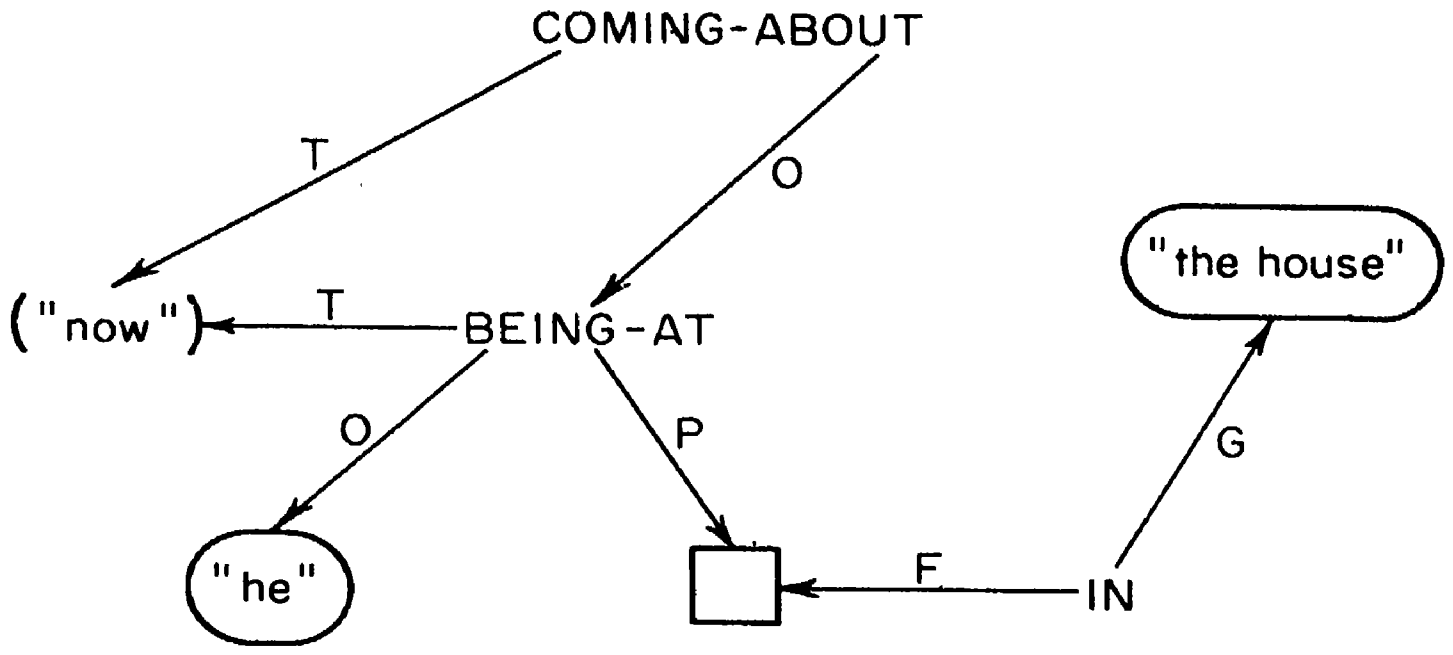

FIGURE 6.4 He is just now in the house.

VII. Motion

Problems with motion arise in every analysis of SR considered in Section II. In this section, what is, as far as we know, an entirely unique approach to the semantics of motion is presented. Our analysis centers on movement locative objects. As has been mentioned, this involves complex modification of the location of motional events. The section first presents a brief discussion of the structure of movement sentences, then motivates our view of motion predication, and finally presents the details of the representation.

VII.1 The Structure of Movemement Sentences

The verbs that take movement locative objects are numerous, see Miller (1972). They include "come", "go", "bring", "take", "climb", "drive", "bit", "punt", "set", etc. The structure of the sentences with movement locative objects resembles that of sentences with non-movement locative objects in being complex. Nearly all examples show causative structure with an action in one event/state causing motion in another (Fillmore, 1971). The appropriate analyses approximately pairs the following:

7.1 John threw the ball through the door.
7.2 John's throwing caused the ball to go through the door.
7.3 Mary walked out of the house.
7.4 Mary's walking caused her to go out of the house.

The only sentences which take simple, causative-like analyses are those with "go" and "come".

An important aspect of the analysis of movement SR's is the concept to be used in the motional event/state. The semantic equivalents of "go" or "come" will not do. These verbs have special deictic conditions on them (Fillmore 1966):

7.5 Go there.
7.6 *Go here.
7.7 *Come there.
7.8 Come here.

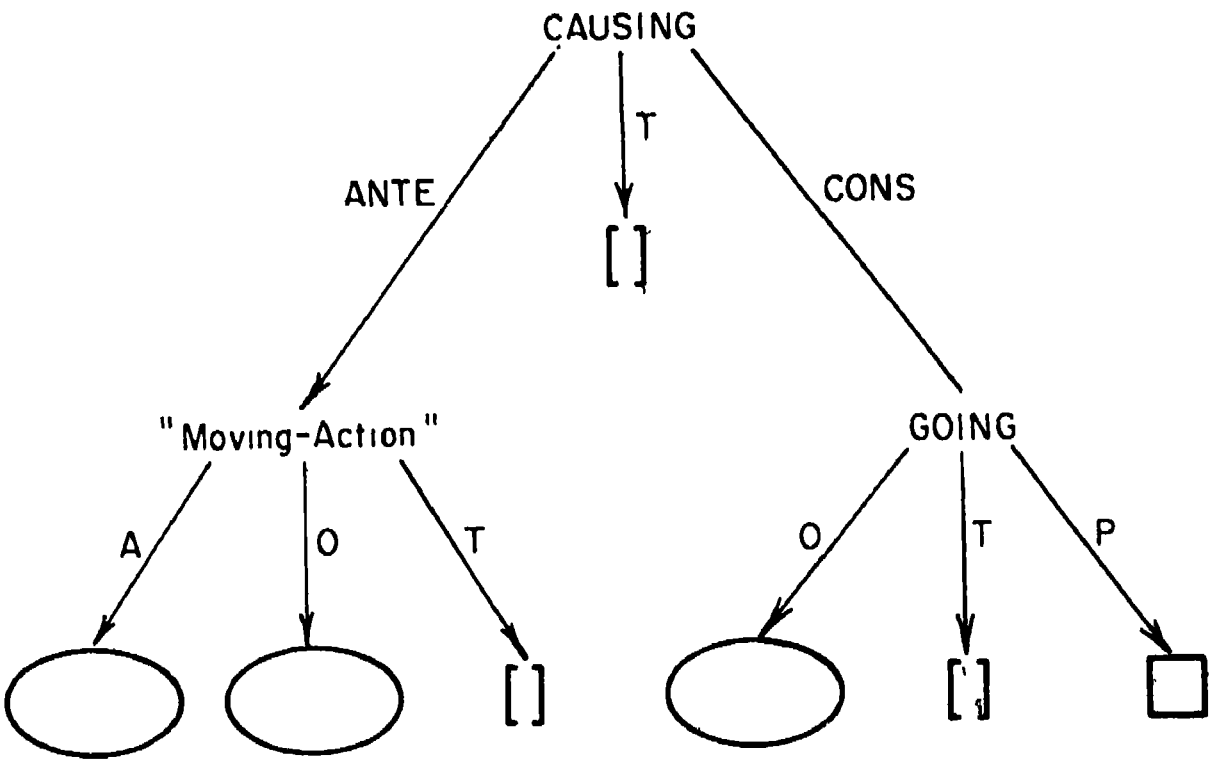

FIGURE 7.1 The basic structure for movement locative object sentences.

Only "take" and "bring" show the same pattern. For this reason, an abstract concept of pure motion, called GOING, will be used in our analysis. Figure 7.1 shows the sentential structure into which most movement SR's will fit. The structures for "take" and "bring" will have Going and Coming, respectively, in place of the abstract form. For "go" and "come" themselves, the semantic structures will match the motional event/state shown with the other verbs with the exception of the type of event/state. The place objects of all the motional events can be considered the same, as can the way SR's apply to the different types of motion. We can also think of motional qualifiers as analyzable with the same structure. Because of this, the structure of movement predication will be considered in general and isolated from other forms.

VII.2 Thinking About Motion

As was pointed out in Section II, one reason that motional SR's are difficult is the multiple predications of different types which must be orderable in time. These problems can be overcome with appropriate consideration of the motion and the place objects of motional events.

The insight for a better analysis comes from considering answers to questions of where motion occurs. Consider the answer to where the first Marathon was run. It is probably something like "in Greece" or "from Marathon to Athens". These tend to place the entirety of motion. It is unlikely to be just "from Marathon" or "to Athens". These just place part of the motion. People tend to locate motion as if it were a single thing, "a motion" so to speak. This is how we propose to think of the place object of motional events.

Place objects of motional events can be thought of as showing that motion, essentially showing a trace of the path of motion. This trace would be similar to the trace a piece of chalk leaves as it crosses a blackboard. But it should be the marks that would be made by the entire chalk if space was a three-dimensional blackboard and the entire chalk could write. This idea is displayed pictorially in Figure 7.2 with another example where something approaching an overexposed photograph of a rolling ball shows a solid cylinder tracing a ball's movement. It is this type of cylinder that motional place objects represent.

This trace idea has one great merit. It allows direct analysis of the most troublesome class of movements SR's:

7.9 He walked through the puddle.
7.10 He walked across the puddle.
7.11 He walked around the puddle.
7.12 He walked over the puddle.

As was pointed out in Section II, the above require a representation that considers every instance of movement. The trace idea does this in such a

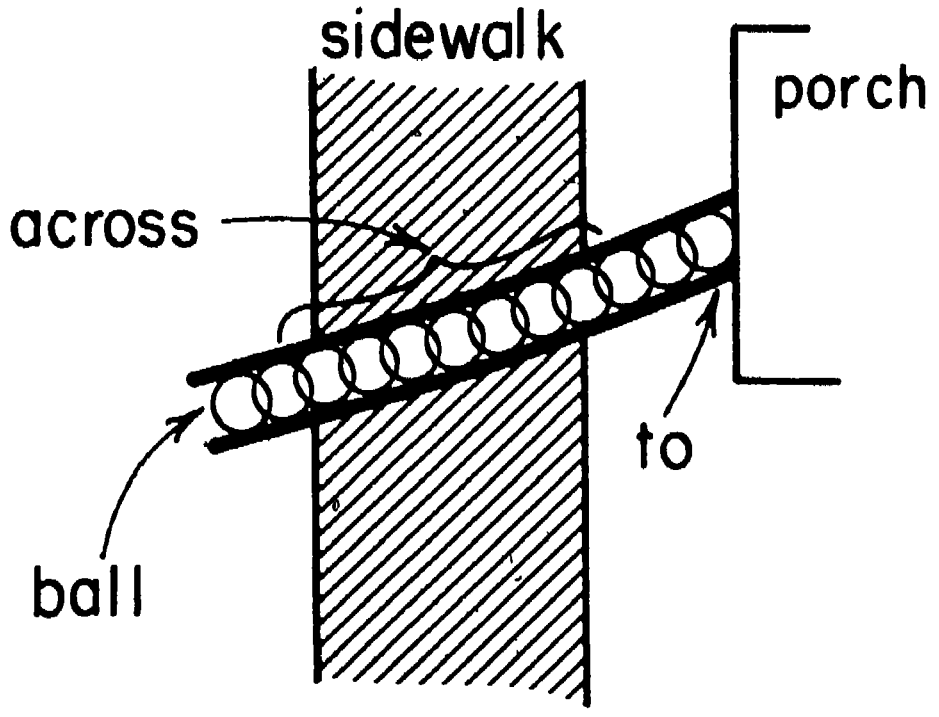

FIGURE 7.2 A ball rolling across a sidewalk to a porch.

way that the SR's can be shown applying to the trace directly. Further, it does it in a way that allows the basic static use of the preposition to be used in the representation:

7.13 The bridges across the Mississippi are closed.

This was pointed out in Section IV to be the same sense that applied in the "across-from" form:

7.14 The man stopped across the street from here.

Hence three usages collapse into one with this representation.

This concept can be extended to allow for differentiating "up" and "down" by considering the solid traces to have an inherent ordering based on the direction of motion:

7.16 He walked up the hill.
7.17 He walked down the hill.

Hence, the traces in 7.16 and 7.17 could be exactly the same except for the ordering and the preposition could be sensitive to this. This ordering

sensitivity shows up with other uses of the prepositions and other prepositions:

7.18 The carotid arteries extend up the neck to the head.
7.19 A woman stood at the front of the line while a man stood at the rear.

Hence its use is not arbitrary.

The trace or path idea does not provide an immediate explanation for other movement SR's, those that reference instantaneous change:

7.20 He hit the ball into the corner.
7.21 He walked out of the house.

With the above we can not say that the overall path of motion was either "into" the house or "out of" the house in the static sense of these prepositions However, there is a way we could use the static sense. If we could refer to positions achieved by the moving object as it followed the path, we could say that there were positions where the object first got to be "into the corner" and "out of the house". This would be like allowing reference to the position of the individual balls displayed in Figure 7.2. We can conclude that we ought to be able to reference parts of place objects.

Being able to reference parts of motion actually leads to a solution of the problem of temporal ordering inherent in multiple SR, such as "... across the yard up the stairs ...". If these durational forms are also thought of as modifying discrete, bounded parts of the kind of place objects that are being discussed, then they too can be compared. For the phrase just mentioned, a part of the motional object that was across the yard could be compared to a part that was up the stairs as being less further along it. The same could be done to compare the parts involved with instantaneous reference.

To summarize, the idea is to think of movement as a trace of the event over time, which has an inherent orientation and which can be predicated in part. We can now almost present our representation. We will first present a slightly incomplete proposal and then revise it.

VII.3 Semantic Structures for Motion

Tentatively, we propose two different functions to produce parts from complete place objects. These are called SEGMENT and UNIT. They will be used with durational and instantaneous references, respectively. The durational function can be taken as picking off bounded parts of a place object. The instantaneous function can be assumed to pick off part of the trace beginning at the earliest point, and going up to the point of change. Both functions will have the place object they accept identified by an S link and the produced space identified by a VALUE link. To distinguish the two outputs, the SEGMENTized place object will have a colon inserted, and the UNITized one a period. The segments will be shown as ordered through "numeric" comparisons. Figure 7.3 therefore gives a tentative analysis for the sentence "The cat came across the yard up the stairs into the house". One SEGMENT function picks out the motion across the yard while another picks out motion up the stairs. A UNIT function picks out motion into the house. The segments are all ordered by less-than-or-equal links.

The temporal ordering of the partial traces is the one tentative part of the analysis. To have it be sensible, some scale of comparison must exist. The appropriate choice appears to be the temporal scale. When the locations were achieved is, of course, what is being ordered. There must also be conventions on application of the comparison. This is because there must be a way to force the comparison on only the appropriate end points of segments. We might develop a way of making these conventions inherent, but I propose to make them explicit.

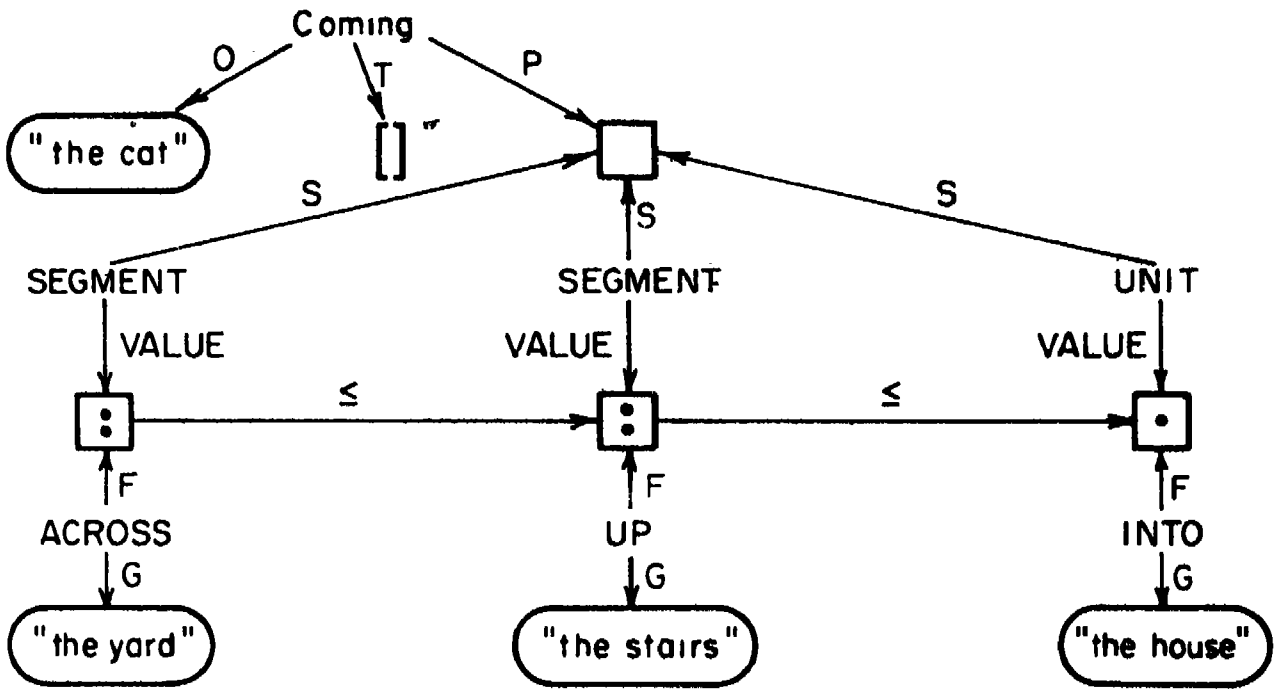

FIGURE 7.3 A tentative analysis of "The cat came across the yard up the stairs into the house."

Our final proposal for the structure of motional SR's is to include time parameters with the functions. In this way, both the end points of the segments and the temporal scale can be identified. For the SEGMENT function, two links, T1 and T2, will identify the times that initial and final points were occupied. For the UNIT function, one link, T, will identify the time the final position was achieved. These structures are shown in Figure 7.4 and 7.5.

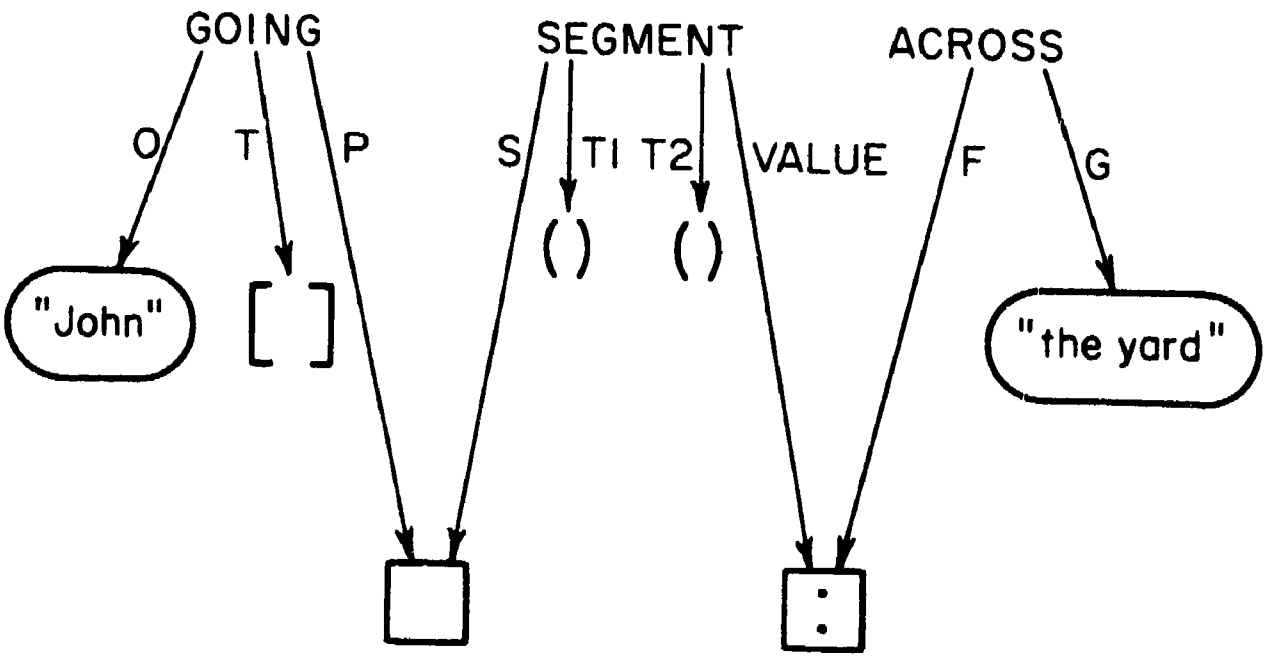

FIGURE 7.4 The motional elements in "John walked across the yard."

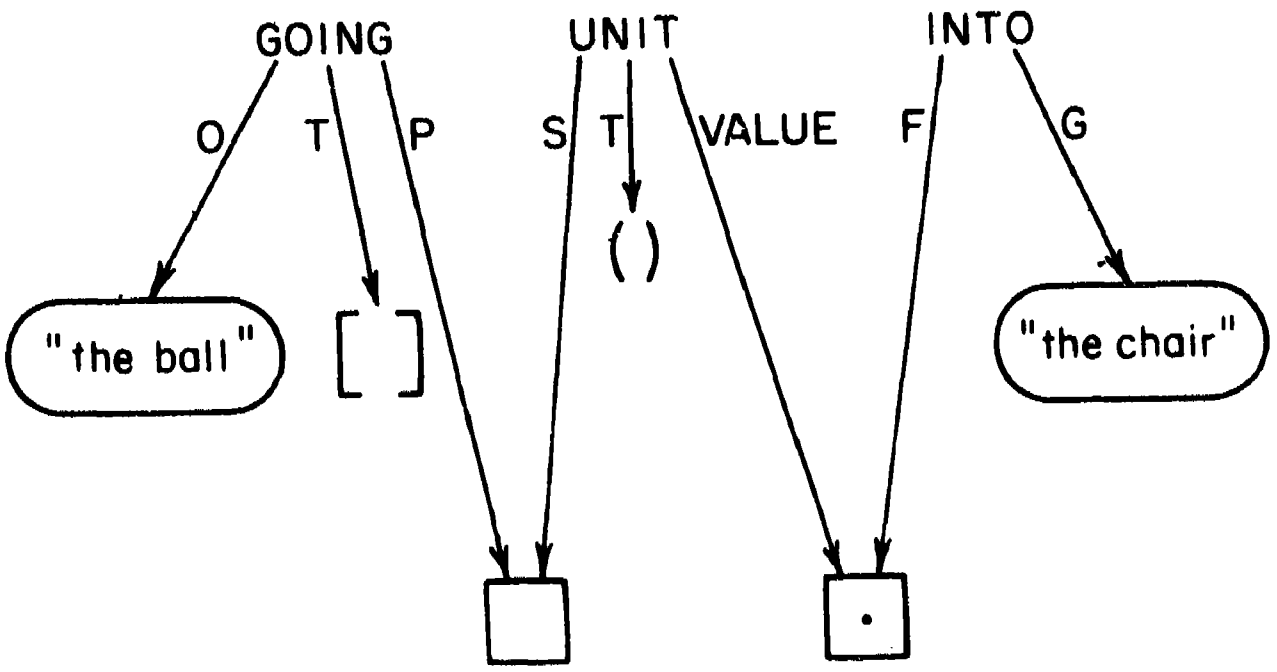

FIGURE 7.5 The motional elements in "I hit the ball into the chair."

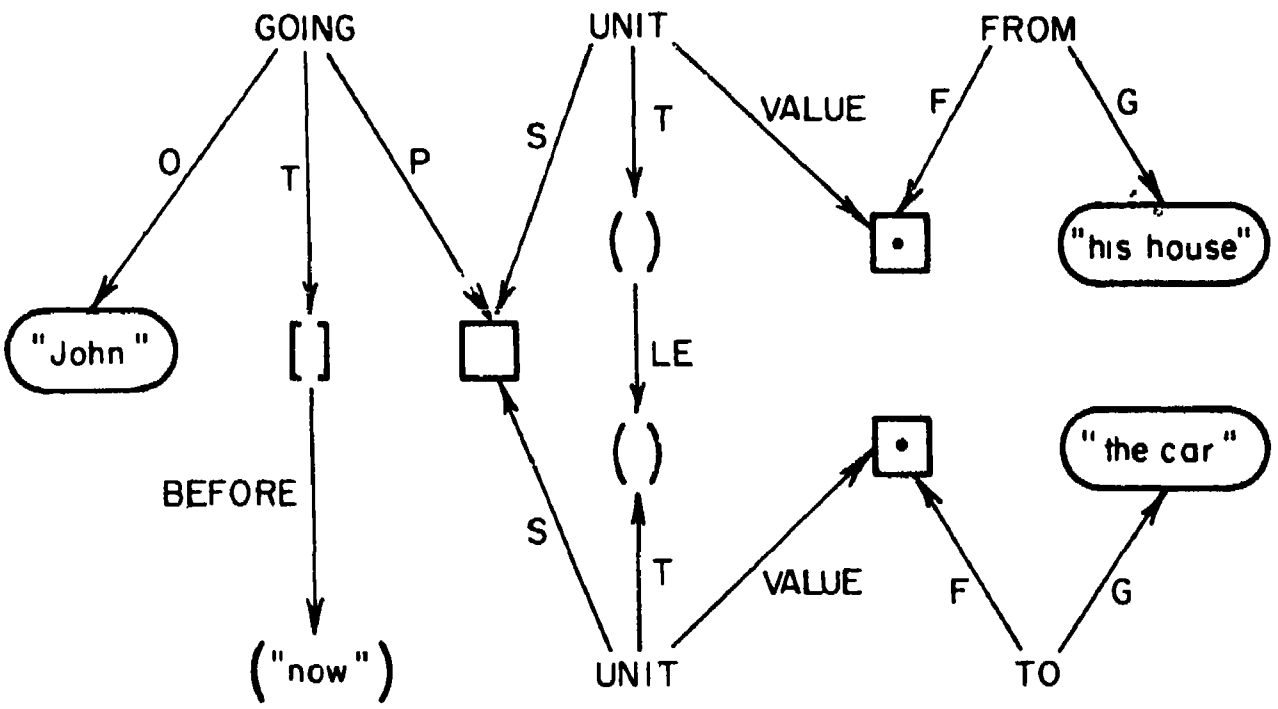

FIGURE 7.6 The motion component of "John walked from his house to the car.".

Now, in order to allow for multiple motional locative objects, two time instances can be related with a temporal relation, LE, for less than or equal. This is done in Figure 7.6.

An interesting aspect of the semantic structure of Figure 7.6 is the static representation of "from". It is to be understood as showing that up to some point

in the journey the moving object was not away from the house, but that it eventually got to be away from it. "Out of" and "off of" are analyzed similarly.

VII.4 Static Spatial References Applying to Motional Events

Besides the durational and instantaneous predications of motion, there can be overall predications of moving objects. These come in two forms. Adjuncts in movement sentences place the entirety of motion:

7.22 In Chicago, he walked around the downtown.
7.23 John came to Chicago in a plane.

One class of verbs, which allows both movement and non-movement locative objects, allows the moving object to be statically placed during movement:

7.24 He carried the dog onto the bus in a box.
7.25 He brought John to Chicago in a plane.

This class is the portability verbs left over from the last section. These verbs take causative analyses with a motional event/state as the caused event. In both of these kind of examples, the motional event must have its motional properties represented at the same time as its static properties. Instantaneous and durational SR's must be shown predicating special place objects which are parts of whole place objects. Therefore, we must show the overall predication applying to different forms. These must be the complete place objects representing the entirety of motion. This is consistant with our other anlayses, as will shortly be seen in more detail. It will also simplify the inference rules that bring down overall spatial predications from higher levels to the motional place objects. This analysis is seen in Figure 7.7 which essentially summarizes this section.

We have introduced two new functions and types of place objects. These have allowed for movement locative objects. We must, however, realize that there are other uses for this analysis. We will see why in the next section.

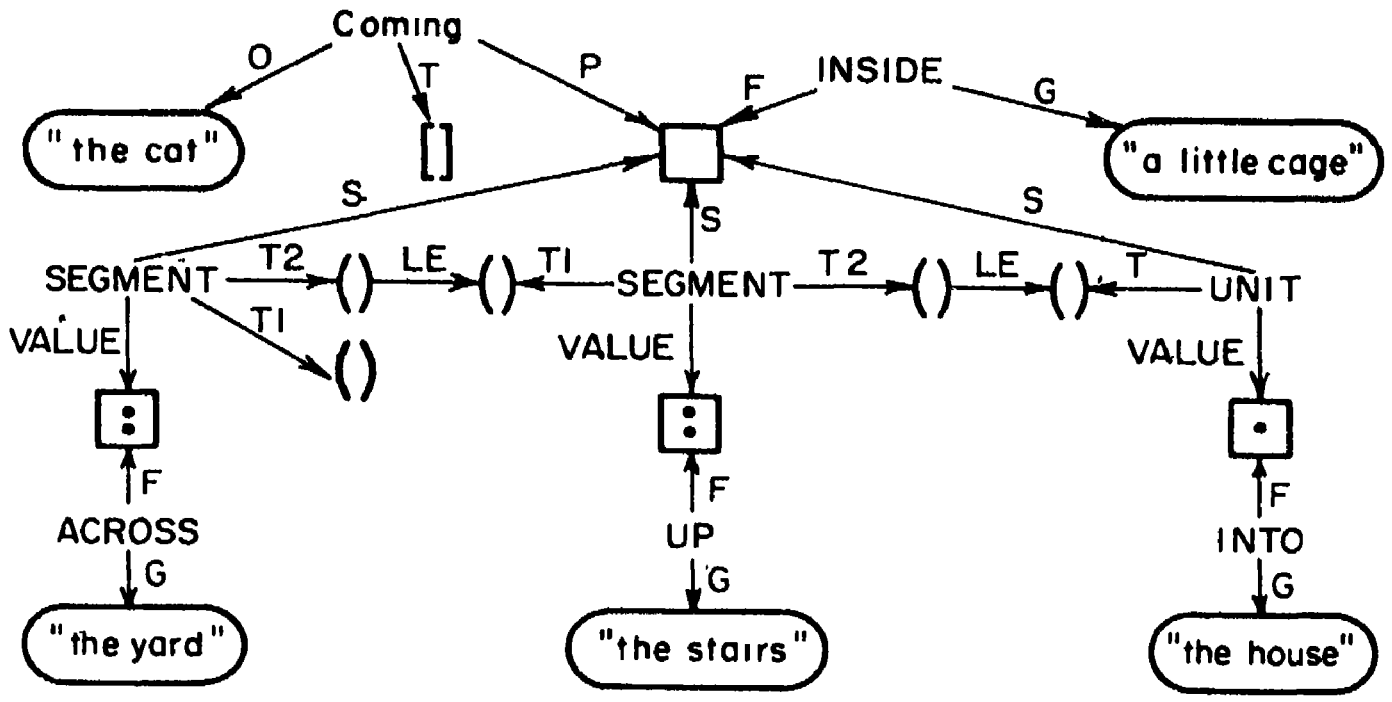

FIGURE 7.7 The motional component of the "The cat was brought across the yard up the stairs into the house.".

VIII. Extending the Motional Analysis to Other Spatial References

The last section may have given the reader the impression that the analyses for motional SR's are really different from those given other SR's. Motional place objects have been set out as a history of movement in space and time. Nonmotional place objects are left as "just" the location of certain events and states of affairs. In this section, we argue that this should definitely not be assumed. The place object of nonmotional SR's must be seen to have the same space-time structure as motional place objects. These are several arguments for this point.

Relative motion has been considered only for sentences with movement verbs, but relative motion and references to motion are common as adjuncts of "non-motional" sentences:

8.1 John held the ice bag to his head in the moving car.

8.2 Jane sat on her purse from New York to Los Angeles.

In each of the above examples two objects are statically related, i.e., John and the ice bag, and Jane and her purse, respectively. However, all are moving. One pair moves but remains static with respect to a car. The other pair is moving and changing with respect to two cities. Hence, motion must somehow be allowed for in these "nonmotional" analyses. Further, change of relative position must be allowed for in at least one. No hint was given of how this last problem is to be solved in any of our discussions of nonmotional SR's.

Even when motion is not overt, time may have to be considered with SR's:

8.3 He died in his car.

As we have seen in the analyses of sentences like 8.3, the car is to be related to the event of the dying. Consider the fact that the car is moveable. If we were to check to see if this were true, we must have either a history of the car's location or have the ability to find its location at the time of death. In oher words, time must be available for even instantaneous events.

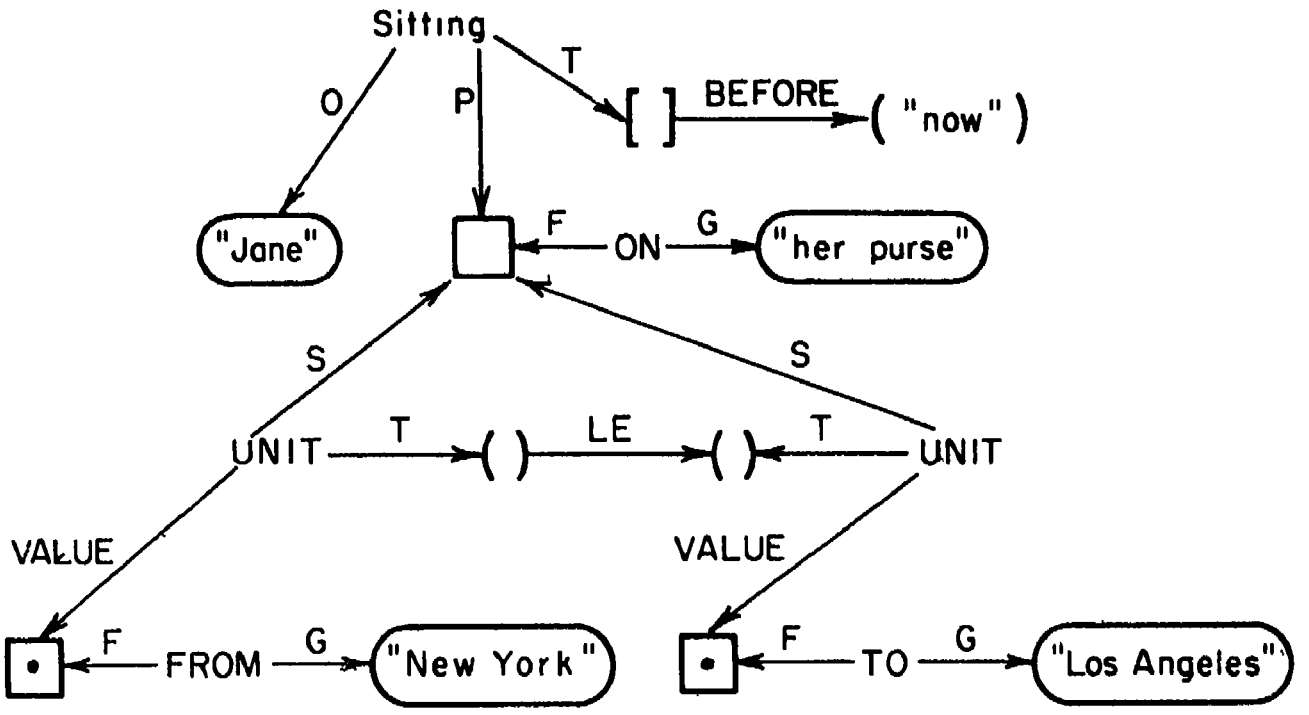

FIGURE 8.1 "She sat on her purse from New York to Los Angeles.".

The way to extend our analysis to cover these facts is to recognize the connection between motional and nonmotional place objects. Motional events involve the location over time of moving objects. Nonmotional events and states of affairs do not necessarily involve moving objects, but they can involve location over time. This location is the space of spaces occupied by an event or state of affairs during its holding. To analyze the cases of relative motion and motion with nonmotional events and states of affairs, these locations-over-time must be taken as the locations of the events and states of affairs. Note that this does not change any of the analyses presented earlier, only the way they are understood. With the same type of place object in all SR's the problematic examples that began this section can be elegantly allowed for with the use of the motional functions. This is shown for example 8.2 in Figure 8.1. So in conclusion we propose an analysis that treats all SR's the same.

IX. Interpreting the Representation

The last several sections presented the "syntax" of our semantic analysis. The term syntax is appropriate since the form of the analysis was presented. The semantics or interpretation to be given the proposed structures was only informally discussed, as when the trace or path analogy for motion was introduced. We noted in Section III that semantic nets do not just allow for the syntactic aspect of meaning structures, but also for the representation of the interpretation or definition of the concepts used in these structures. In this section, this property will be used to help formalize an interpretation of our analyses. This is only one of many possible interpretations, but showing it will help clarify the semantic structures. The formalism for the conceptual definitions is based on Brachman (1977). Again, many abbreviations of a complete formalism are used.

The center of our previous discussions was the place objects. This must also be true in discussing conceptual definitions. The nature of the place objects must first be defined, followed by the definition of everything that relates to place objects. Event/states will be discussed first, then the SEGMENT and UNIT functions, and then the prepositional concepts. The definition of the WAY function will not be attempted.

IX.1 The Place Object

Our interpretation of the place object will be based on a discrete representation of time. Time can be considered as composed of arbitrarily densely packed time instances. A place object can show the location of an event/state at one time for instantaneous events/states, a set of consecutive instances for durational event/states, and any set of instances for intermittent event/states. Structurally, we can take a place object to be a set of what we can call

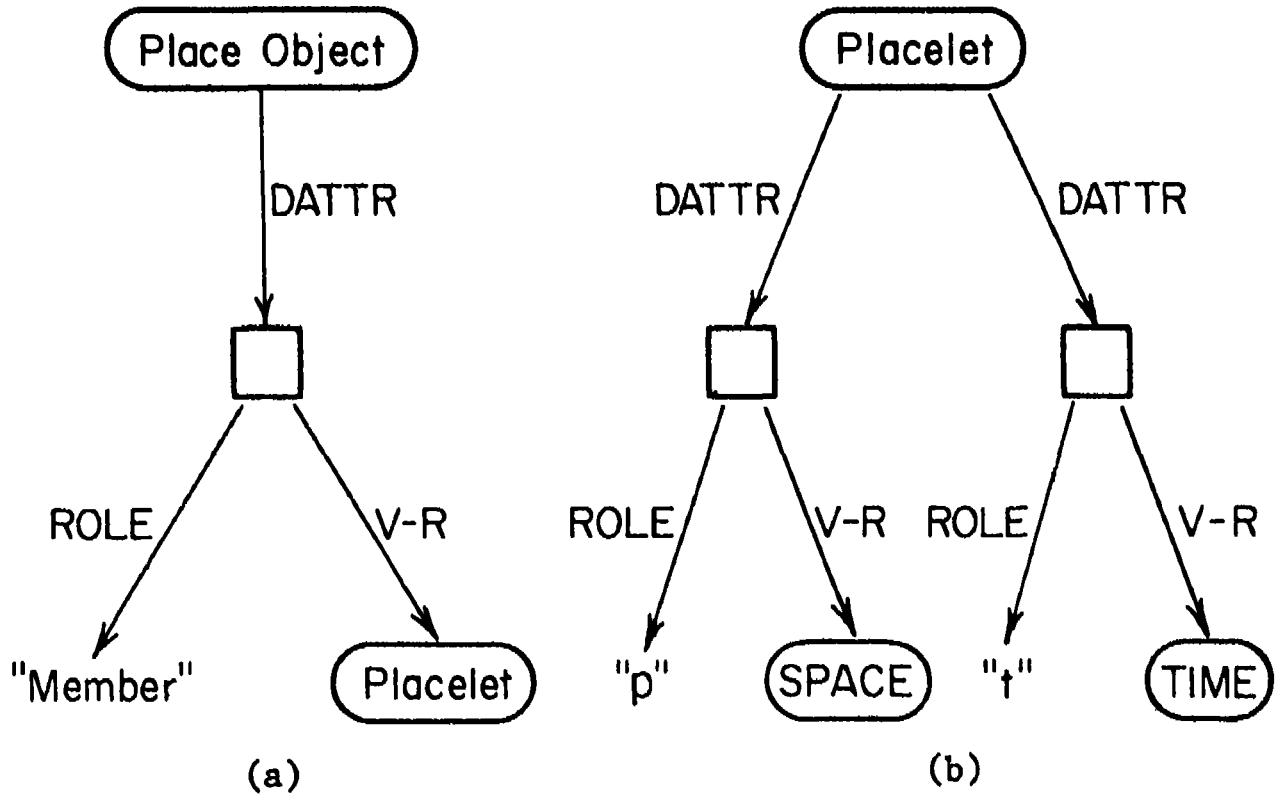

FIGURE 9.1 a) The definition of Place Object.
b) The definition of Placelet.

placelets, each of which is an ordered pair whose first element is a volume in space and whose second is an instant in time. This is formalized in Figures 9.1a and b.

In Figure 9.1a, the node labelled Place Object stands for the concept of a place object. An arc with the special label DATTR points to its one defining attribute. This attribute and all others in this section are shown by a special node shaped as a square. The fact that placelets are members of the set that compose place objects is the defining attribute of place objects. This node captures this by using an arc labelled ROLE to point to the special name Member and one labelled V-R for value-restriction to point to the restriction on any member, namely that it must be a Placelet. In Figure 9.1b, the concept Placelet

is defined. In this case, the concept has two defining attributes since a placelet must have a space and a time. The two attributes are shown accordingly with the one in the role p restricted to be a SPACE and one in the role called t being a TIME. The concepts of SPACE and TIME will be treated as primitive, here.

IX.2 Event/States and Place Objects

The structure we defined for place objects will be referenced whenever place objects are used. One reference will be in event/states where place objects are involved with the case P. Hence, with the conceptual definition of every type of event/state that has a location, there will be a defining attribute with role P and value-restriction Place Object.

It is also the case that with each event/state, there will be a way to show how the place object fits in with the definition of the event/state. This will include the way in which the place object will be related to the participants in the event/state and structural restrictions on the place object Consider the abstract event/state GOING. GOING requires of its place object that the placelets show where the moving object was at each instance during the movement. The placelets must refer to the time of the GOING. Since a discrete representation of time is used, placelets for successive instances of time during the movement must show an overlap in positions occupied. Further, since movement is

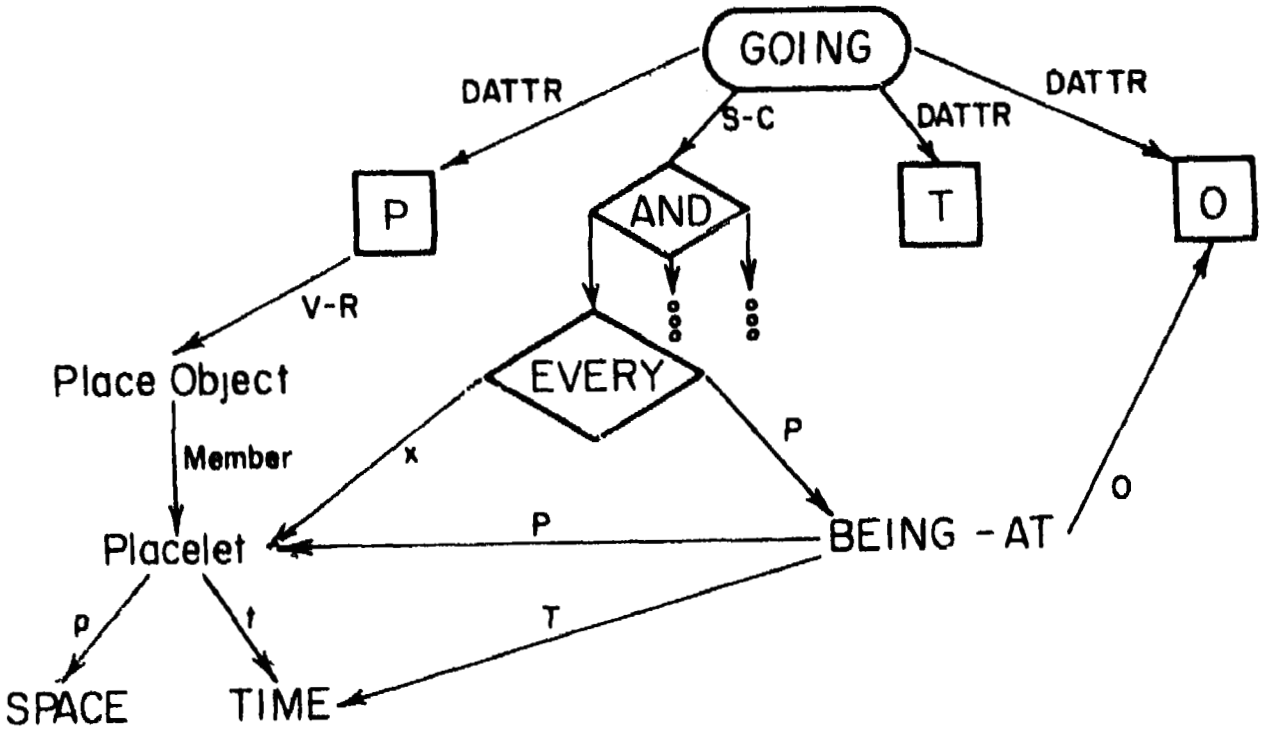

FIGURE 9.2 Partial Definition of GOING.

necessary at least two positions among the placelets must be different. All these facts will have to be shown in the definition of GOING.

In order to show the flavor of event/state definition, we show in Figure 9.2 the relation between moving object and place object for GOING. The definition shows that the event/state has three attributes corresponding to the cases, P, T, and O. Names have been added to the attribute nodes to make this easier to see. The event/state also has a structure identified by a special S-C link, for structural conditions, which is used to identify how the event/state is structured. The conditions for GOING are a set of conjuncts identified by the label in the diamond shaped node. This shape is an aid to the reader and indicates a logical operator. The structure necessary for the spatial relation is the leftmost of the conjuncts. It essentially takes the form of an implication statement saying that for every placelet in the place object, the moving object, which is identified by the O role, will have a BEING-AT holding for the place and

time in the placelet. The statement begins with a logic node, labelled EVERY, identifying a universal quantification. The domain of the quantified variable is shown by the link labelled x. By pointing to the appropriate attribute node, the restriction is to the set of placelets in the place object. The link labelled P identifies the proposition within the scope of the quantifier which shows that for each placelet BEING-AT is the case for the entity in the 0 role at the place and time of that placelet. The representation of this last depends on the ability to focus in on attributes of entities being quantified over, for which see the FOCUS-SUBFOCUS mechanism in Brachman (1977). This ability is indicated here by the special representation of the P attribute.

IX.3 The SEGMENT and UNIT Functions

Formalizing the definitions of the SEGMENT and UNIT functions is fairly straightforward. Both can merely identify subsets of sets of placelets. The structural conditions for both can be shown with the same function which can be called GENERATE-RANGE. It can be assumed to apply to any set and to produce the subset that fits a range defined by two limits and a measure. Since the definitions are similar, only the UNIT function will be shown, see Figure 9.3.

The GENERATE-RANGE function for the UNIT definition can take as its input, identified by the SOURCE link, the place object marked by the S link of the UNIT function. The scale for measurement can be established by reference to a special Temporal scale and apply to the time values in the set of placelets being operated on. The boundaries of the subset to be generated can be shown by FROM and TO links. The FROM value would be produced from the set of placelets by a special LOWEST function to produce the lowest time value from among the placelets. The TO value can be a placelet with time specified by the T role in the UNIT function and unspecified space. The RESULT link can show that the generated value should be connected to the VALUE role of the UNIT function.

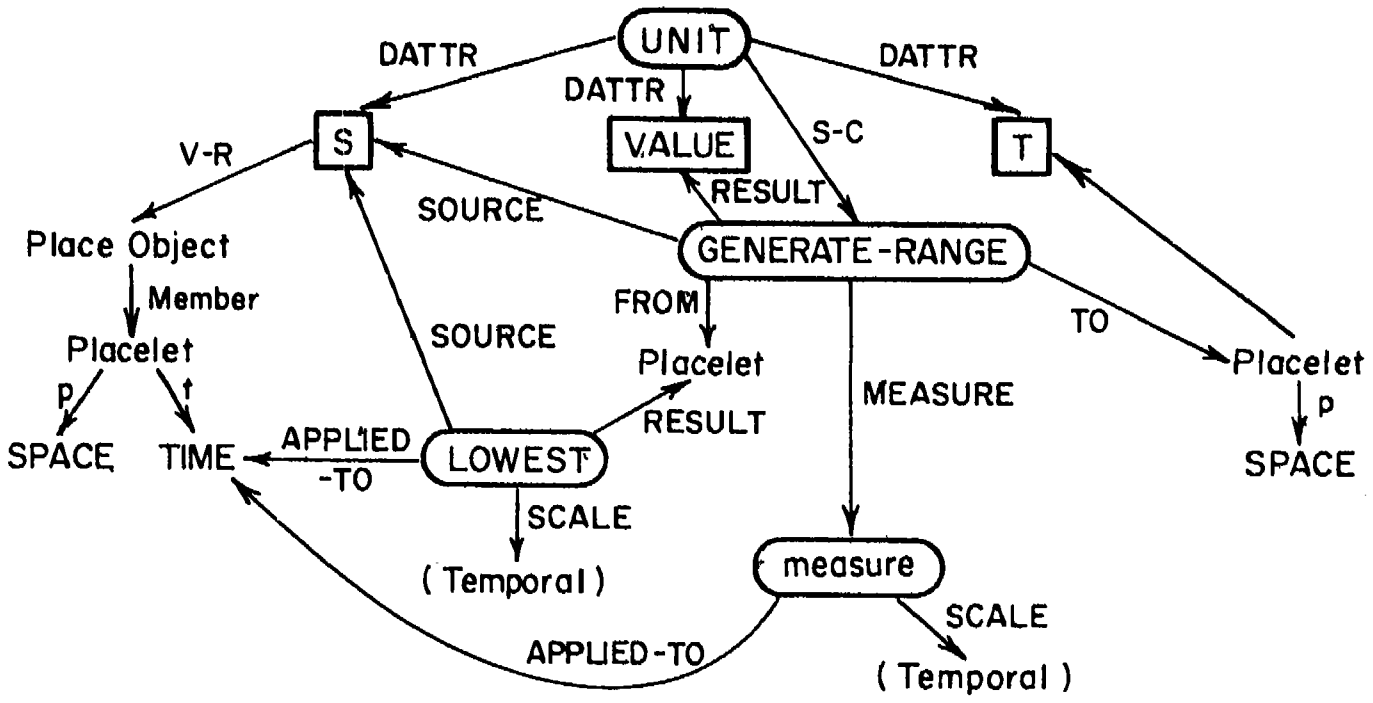

Figure 9.3 Definition of the UNIT Function.

IX.4 Prepositions

Prepositions are the final concepts whose relation to place objects will be considered. It is always the case that prepositional concepts relate place objects which are locations over time to simple objects. The suggestion in the last section was that the locations of the place objects are related to the position of the referenced object at the times of the place object. The nature of this relationship depends on the source of the place object being predicated.

Consider first predication of place objects which directly show the location of event/states:

9.1 In his new shoes, John walked through the barnyard.

The above example asserts that at each instance during the walking, the walker was "in" with respect to the position of his shoes. Such examples require that the object's position at each instant during the event/state be compared to the location of the event at that instant.

Allowing for the prepositional concepts applying to place objects produced by the UNIT function must be done differently:

9.2 John went into the car.

Example 9.2 is of this class with the semantic structure showing concept INTO relating a part of the going to the location of the car. Here the position of the simple object must again be compared at each instant of the place object to the location of the place object at that instant. However, only at the last instant must the relation be shown as holding.*

Prepositions predicating place objects produced by the SEGMENT function are more complex:

9.2 An ant is crawling up your arm.

A simple interpretation of the prepositional concept in the above may be problematic. Since your arm could be in motion, a stationary observer would include some motion attributable to your arm in the ant's path. Further, even if we wanted to take the position of the arm at some one instant it is unclear which to take. These problems, however, disappear with the realization that the motion referenced is not with respect to an arbitrary observer but to one on the arm. For him, the SR can be treated as involving not a moving arm but one essentially static in space. This can be allowed for by requiring the conceptual definition of the prepositions to project the referenced objects' positions shown in the place object in the F case onto the base object shown in the G case. This will be like taking the base as a static ground and the referenced object as a figure seen against it.

*Since the change into the final state may be gradual and not dramatic, fuzzy relations (Zadeh, 1973) might be used. For instance, the degree of "into"-edness could be quantified with the analysis showing that a certain degree was reached.

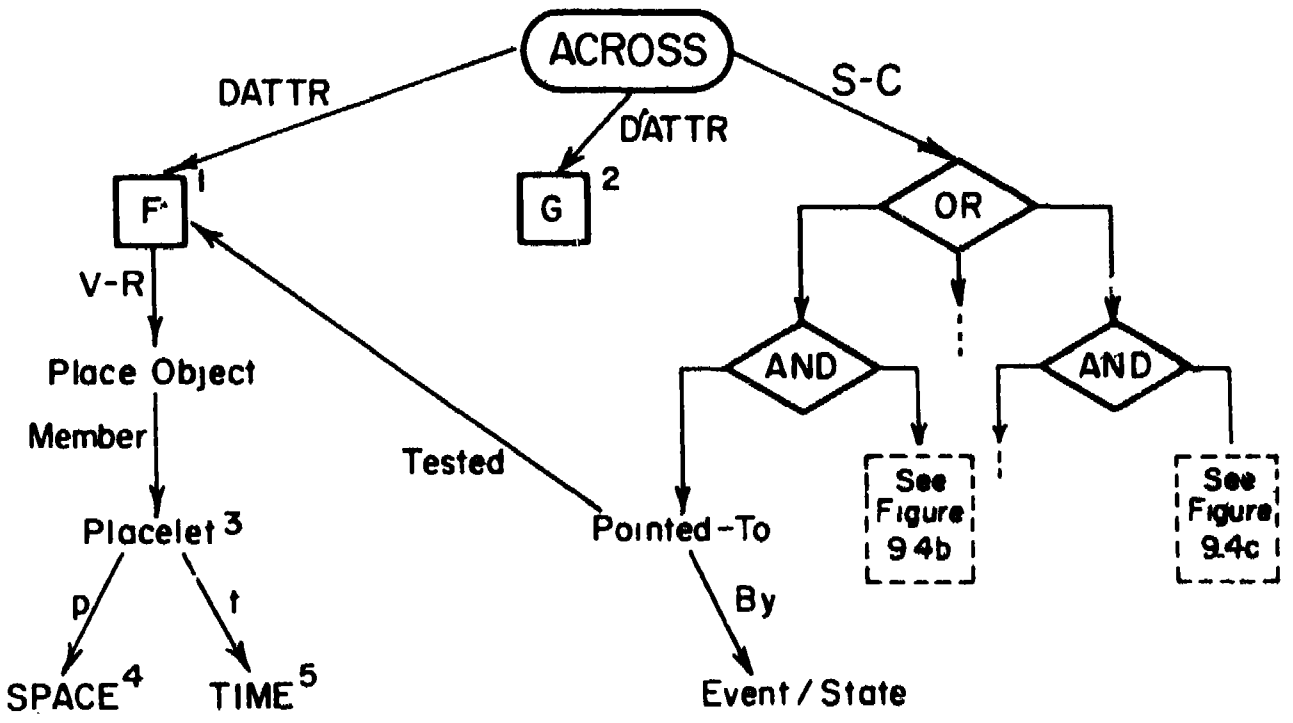

FIGURE 9.4a The definition of ACROSS: highest level.

The conceptual level definition of the prepositional concept ACROSS is sketched in Figures 9.4a, b, and c. Here again, it should be remembered that we are not trying to show the entire meaning of the prepositional form, only its relation to the place object. Accordingly, a number of unanalyzed forms will be used. Perhaps the most curious one of the forms will be one labelled across, this can be seen as the physical part of the concept ACROSS. It would have to be the next form developed if we are analyzing the meaning of the preposition.

The definition of ACROSS begins in Figure 9.4a with an indication of the F and G nodes. The structural condition again shows the required connection between these two elements. It has three alternative opportunities for satisfaction, one for direct predication, one for UNIT functions, and one for SEGMENT functions. The three choices are reflected by the three arcs projecting from the OR node. Which case applies should be shown by one of the two arcs projecting from the corresponding AND node. The test for direct predication is identified by the object in the F role being pointed to by an Event/State.* This is shown by

*It could just as well be pointed to by a WAY node. This could be tested for with the addition of a disjunct.

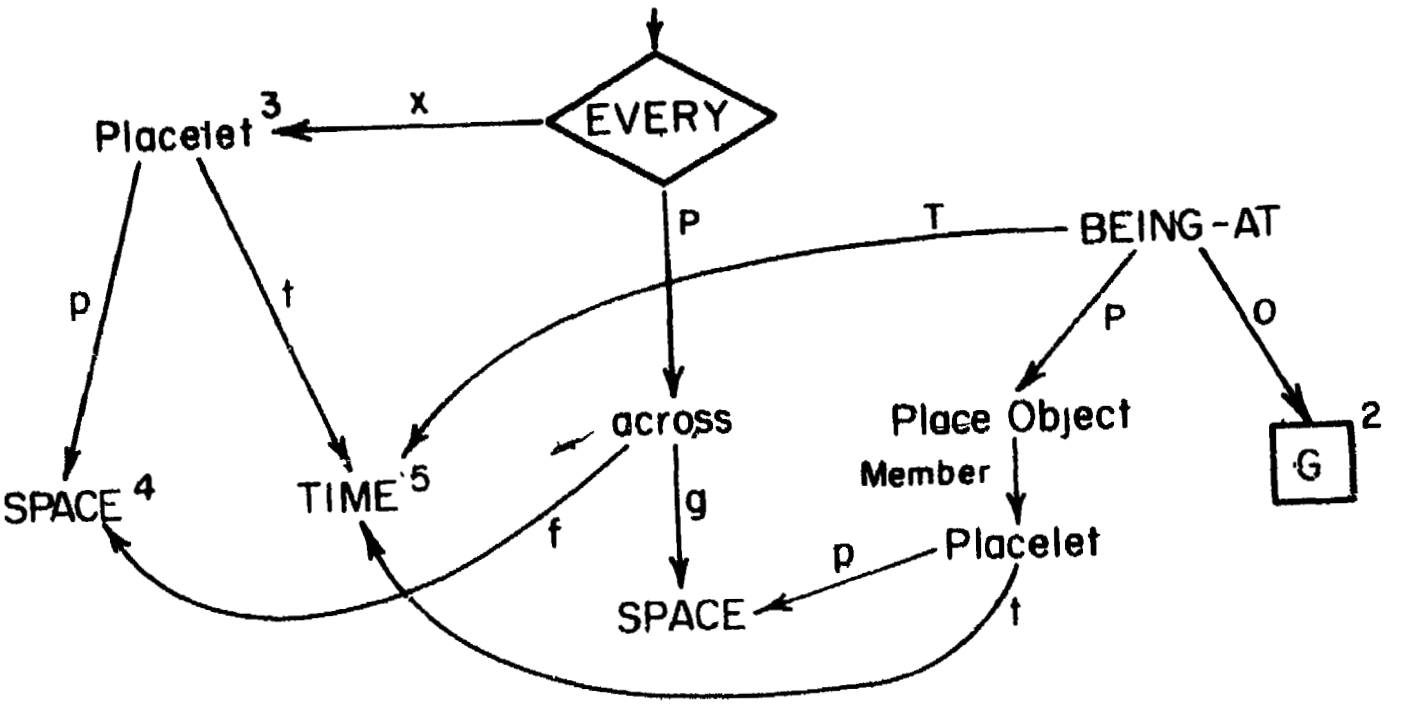

FIGURE 9.4b The definition of ACROSS: direct predication.

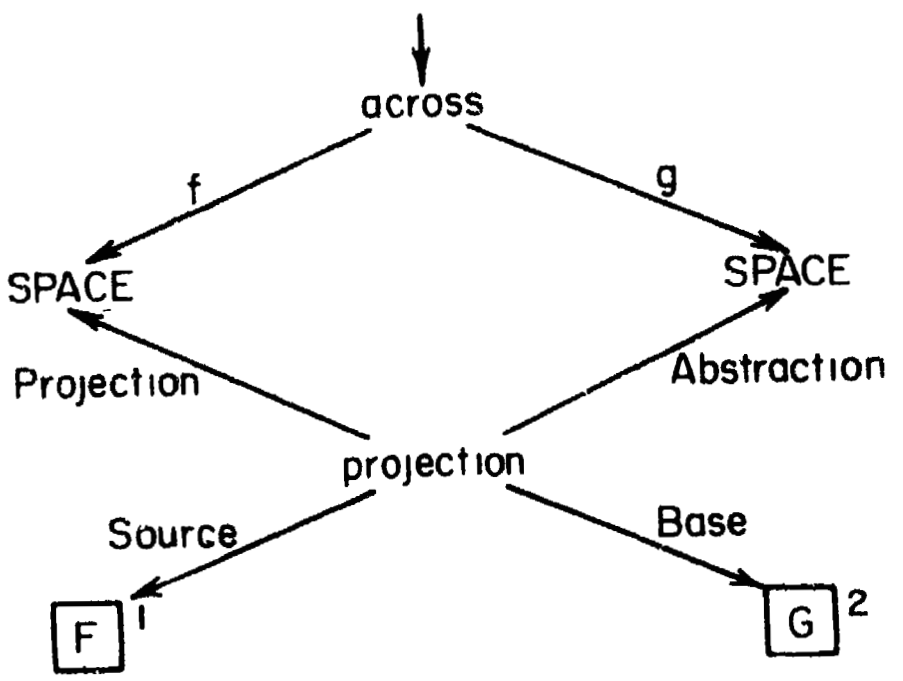

FIGURE 9.4c The definition of ACROSS: SEGMENT predication.

a test labelled with these names. The condition that must hold if we do have direct predication is shown in Figure 9.4b. It states that for every placelet the physical across must hold between the space of the placelet and the space of the place object of a BEING-AT which locates the object identified by the object in the role at the time shown in the placelet. The superscripts on nodes in the figure establish co-reference between the different parts of Figure 9.4. The structure for the UNIT case is fairly similar and not shown. For the SEGMENT case, the condition is based on one space being physical across from another. This is shown in Figure 9.4c. Both spaces are shown being produced by a special

projection function which takes the place object in the F role and the object in the G role and produces the projection, shown by an arc of that name, and an abstract space to compare it to, shown by the Abstraction arc.

To summarize, the section has shown how conceptual level interpretation can be given the semantic structures proposed earlier in the paper. Any system that uses the semantic structures can also use the interpretations. Of course, the interpretations are based on one way of structuring place objects. Since there are other ways, other interpretations are possible.

X. Limitation, Summary, and Conclusions

There are definite limits to the claims we wish to make. In this concluding section, we point out several half-solved and unsolved problems, one area where we could conceivably expand our claims, and then end with a summary and final defense.

Metaphorical usages are important but difficult subjects for semantic representation. Things like "climbing the ladder of success" are far enough away from spatial reference to be ignorable. However, some SR phenomena appear to be metaphors:

10.1 John yelled his greetings to John.

In the above, an imaginary object, "his greetings", seems to be sent through space. In the following, a hypothetical journey is referenced:

10.2 The bridge goes from New York to New Jersey.

Any direct representation of these phenomena using the definitions from the last section is unlikely since a non-instantaneous time interval must be present, while these sentences are basically instantaneous. My best suggestion is to represent these using the motional structures but to indicate by a function or operation applied to every appropriate form that the actual sense is metaphorical. Unfortunately, this leads to odious complexity. It is probably better to say our claims stop at this point.

Many adverbials can apply to modify SR's or show spatial-like properties of events and states. We have no definite analyses for these, either. An incomplete study indicates that these may be analyzable within our model. For example, some forms can be shown as modifiers of a prepositional concept, such as in the following:

10.3 I put the ball completely under the car.

Some seem to predicate place objects directly:

10.4 I walked two miles.

Others seem to coordinate with SR's:

10.5 Go straight into the house.

Here "straight" can be shown as predicating the part of the journey up to the time the house was entered. However, I do not know how many other terms remain to be considered.

Two problems remain completely unsolved. The first involves relative motion.

10.6 The ant walked over the rising pile of dough.

Now, it is actually possible that 10.6 can be true but that the dough changed during the walking. Since there is no one static pile of dough, this would make problematic the use of the dough as the object into which the motion of the ant is projected. Secondly, it appears that inferences vary in habitual sentences:

10.7 He bought a present for her in New York.
10.8 He always buys a present for her in New York.

The last examples differ because the former says an event occurred in New York but the latter says that a certain type of event must occur when the person is in New York. I have solutions for neither problem. We can only appeal to the fact that these phenomena do present problems in many other areas of semantics.

Switching from difficult to promising areas, one strong possibility exists for expanding the analysis and corroborating it. As the analysis was developing, time could be seen to become more closely associated with place. In the end, time was claimed to fit in every place object. Perhaps with our place object, no separate time attribute needs to be associated with events or states of affairs. We may be able to claim that, to quote E. J. Lemmon (1967), we can associate events with "space-time zones" instead of times and spaces. How this would be done remains to be seen. However, if we have not already met our goal of putting space on a par with time, that would certainly do it.

In summary, this paper has shown how the semantic structure of spatial references can be shown as locating events and states of affairs. Within a semantic net, this has the form of showing a location as an attribute of event/state nodes. In line with this, the concept of a place object, showing where events and states of affairs held at instances of time, was developed. Several functions were developed for use in predicating locations. Inferencing of spatial facts, the use of prepositional-like concepts for showing spatial relationships, and the overall semantic structure of utterances was also discussed.

Throughout the paper, the main justification has been that the analysis handles phenomena that other analyses do not. However, there are other justifications. Only one source for space simplifies the modeling of spatial phenomena. Using only static forms simplifies the interpretation of spatial terms. Also, the use of static forms fits in with proposals for state-based semantic representations (Cercone and Schubert, 1975). Finally, we can see that the analysis of semantic structures, in general, fits in with "deeper" analyses of semantic structure such as Schank (1973) and Norman and Rumelhart (1975). In sum, there appears to be a strong case for the analysis.

Acknowledgements

I wish to acknowledge the help of David Brown on an earlier draft of this paper and the even earlier advice of Richard L. Venezky and Peter S. Schreiber who launched me into time and space.

Bibliography

Abrahamson, Adele A., "Experimental Analysis of the Semantics of Movement," in Explorations in Cognition, eds., Donald A. Norman and David E. Rumelhart, and the LNR Research Group, San Francisco: W. H. Freeman and Company, 1975, 247-276.

Badler, Norman, "The Conceptual Description of Physical Activities," American Journal of Computational Linguistics, Microfiche 35, (1975), 70-83.

Bennett, David C., Spatial and Temporal Uses of English Prepositions: An Essay in Stratificational Semantics, London, England: Longman Group, 1975.

Brachman, Ronald J., A Structural Paradigm for Representing Knowledge, Ph.D. Dissertation, Harvard University, 1977.

Bruce, Bertram, "Case Systems for Natural Language," Artificial Intelligence, 6, 4, (1975), 327-360.

Cercone, Nick and L. K. Schubert, "Toward a State-Based Conceptual Representation," in Advance Papers of the 4th International Joint Conference on Artificial Intelligence, Cambridge: Artificial Intelligence Laboratory, Massachusetts Institute of Technology, 1975, 83-90.

Coles, L. Stephen, "An On-Line Question-Answering System with Natural Language and Pictorial Input," in Proceedings - ACM National Conference, 1968, 157-167.

Cooper, Gloria S., A Semantic Analysis of Engish Locative Prepositions, Report No. 1587, AFCRL-68-0056, Bedford, Massachusetts: Bolt, Beranek and Newman, Inc., January 1968.

Davidson, Donald, "The Logical Form of Action Sentences" and "Reply to Comments," in The Logic of Decision and Action, ed., Nicholas Rescher, Pittsburgh: University of Pittsburgh Press, 1967, 81-95, 115-120.

Denofsky, Murray E., How Near is Near?: a "near" specialist, AI Memo No. 344, Cambridge: Artificial Intelligence Laboratory, Massachusetts Institute of Technology, February 1976.

Fillmore, Charles J., "Deictic Categories in the Semantics of 'Come'," Foundations of Languages, 2, (1966), 219-227.

Fillmore, Charles J., "The Case for Case," in Universals in Linguistic Theory, eds., Emmon Bach and Robert T. Harms, New York: Holt, Rinehart and Winston, 1968, 1-88.

Fillmore, Charles J., "Some Problems for Case Grammar," in Georgetown University Monograph Series on Languages and Linguistics 24, ed., R. J. O'Brien, Washington D. C.: Georgetown University Press, 1971, 35-56.

Geis, Michael L., "English Time and Place Adverbials," in Working Papers in Linguistics, No. 18, Columbus: Department of Linguistics, The Ohio State University, 1975a, 1-11.

Geis, Michael L., "Two Theories of Action Sentences," in Working Papers in Linguistics, No. 18, Columbus: Department of Linguistics, The Ohio State University, 1975b, 12-24.

Geis, Michael L., "What do Place Adverbials Modify?," in Working Papers in Linguistics, No. 18, Columbus: Department of Linguistics, The Ohio State University, 1975c, 25-29.

Gruber, Jeffery S., Studies in Lexical Relations, Ph.D. Dissertation, Massachusetts Institute of Technology, 1965 (distributed by the Indiana University Linguistic Club, 1970).

Harmon, Gilbert, "Logical Form," Foundations of Language, 9, (1972), 38-75.

Hendrix, Gary G., Partitioned Networks for the Mathematical Modeling of Natural Language Semantics, Technical Report NL-28, Austin: The University of Texas at Austin, Department of Computer Sciences, December 1975.

Hobbs, Jerry R., "A General System for Semantic Analysis of English and its Use in Drawing Maps from Directions," American Journal of Computational Linguistics, Microfiche 32, (1975), 21-41.

Jacobs, Roderick A. and Peter S. Rosenbaum, English Transformational Grammar, London: Ginn and Company Ltd., 1968.

Kochen, Manfred, "Automatic Question-Answering of English-Like Questions About Simple Diagrams," Journal of the ACM, 16, 1, (1969), 26-48.

Kuipers, Benjamin, Representing Knowledge of Large-Scale Space, AI-TR-418, Cambridge: Artificial Intelligence Laboratory, Massachusetts Institute of Technology, July 1977.

Lakoff, George, "Pronominalization, Negation, and the Analysis of Adverbs," in Readings in English Transformational Grammar, eds., Roderick A. Jacobs and Peter S. Rosenbaum, Waltham, Massachusetts: Ginn and Company, 1970, 145-165.

Lemmon, E. J:, "Comments on D. Davidson's 'The Logical Form of Action Sentences'," in The Logic of Decision and Action, ed., Nicholas Rescher, Pittsburgh: University of Pittsburgh Press, 1967, 96-103.

Lyons, John, Introduction to Theoretical Linguistics, Cambridge, England: Cambridge University Press, 1968.

Miller, George A., "English Verbs of Motion: A Case Study in Semantics and Lexical Memory," in Coding Processes in Human Memory, eds., A. W. Melton and Edwin Martin, Washington, D. C.: V. H. Winston and Sons, 1972, 335-372.

Norman, Donald A., David E. Rumelhart, and the LNR Research Group, Explorations in Cognition, San Francisco: W. H. Freeman and Company, 1975.

Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech and Jan Svartvik, A Grammar of Contemporary English, New York: Seminar Press, 1972.

Schank, Roger C., "Identification of Conceptualizations Underlying Natural Language," in Computer Models of Thought and Language, eds., Roger C. Schank and Kenneth M. Colby, San Francisco: W. H. Freeman and Company, 1973, 187-247.

Schubert, L. K., "Extending the Expressive Power of Semantic Networks," Artificial Intelligence, 7, (1976), 163-198.

Shapiro, Stuart C., "A Net Structure for Semantic Information Storage, Deduction, and Retrieval," in 2nd International Joint Conference on Artificial Intelligence: Advance Papers Papers of the Conference, London: British Computer Society, 1971, 512-523.

Simmons, Robert F., "Semantic Networks: Their Computation and Use for Understanding English Sentences," in Computer Models of Thought and Language, eds., Roger C. Schank and Kenneth M. Colby, San Francisco: W H. Freeman and Company, 1973, 63-113.

Sondheimer, Norman K., The Computational Semantics of Locative Prepositions, Ph.D. Dissertation, The University of Wisconsin-Madison, 1975.

Sondheimer, Norman, "A Semantic Analysis of Reference to Spatial Properties," Linguistics and Philosophy, 2, (1978).

Sondheimer, Norman and Doyt Perry, "SPS: A Formalism for Semantic Interpretation and its Use in Processing Prepositions that Reference Space," American Journal of Computational Linguistics, Microfiche 34, (1975), 49-63.

Talmy, Leonard, "Semantics and Syntax of Motion," in Syntax and Semantics, 4, ed., John P. Kimball, New York: Academic Press, 1975, 181-238.

Tsotsos, John K., A Prototype Motion Understanding System, Technical Report No. 93,. Toronto: Department of Computer Science, University of Toronto, June 1976.

Winograd, Terry, Understanding Natural Language, New York: Academic Press, 1972.

Woods, William A., "What's in a Link: Foundations for Semantic Networks," in Representation and Understanding: Studies in Cognitive Science, eds., Daniel G. Bobrow and Allan Collins, New York: Academic Press, 1975, 35-82.