Cognitive Topology and Lexical Networks[1]

Claudia Brugman and George Lakoff

University of California, Berkeley

This article is a rather long and detailed study of a polysemous lexical item—the English word *over*—paying specific attention to the character of the relations among its senses. Polysemy is a subtype of lexical ambiguity, contrasting with homonymy, wherein a single lexical form is associated with more than one meaning, and those meanings are unrelated. In the case of polysemy, one word is taken as having senses which are related. The distinction is an important one for the resolution of lexical ambiguity, for, as we will show in the bulk of this article, the way semantic information is stored in a lexical entry may differ depending on whether that lexical entry is taken as reflecting homonymy or polysemy. We will show that the common practice of giving a list of meanings of ambiguous items is neither the only way, nor, for polysemous words, the most efficient way, of storing such semantic information. We will argue instead that a network-style mode of storage is cognitively real, and that this allows for a maximum of shared, and otherwise related, information between senses.

Network-style representations are common in many areas of AI. But we use it here not by notational fiat but as part of a much more general conception of categorization, explored at length in Lakoff 1987. That work provides detailed empirical evidence and theoretical argument against the classical view of categories as collections of objects characterized by lists of necessary-and-sufficient conditions, and in favor of an enriched view of categories. On that view, categories may contain a great deal of internal structure—for instance, that one member of a category should be more exemplary of that category than some other member; that the boundaries of the category are not always clear-cut; that categories may be characterized in part with respect to their contrast with other categories. The category structure utilized here is called a "radial" structure, with a central member and a network of links to other members. Each noncentral member of the category is either a variant of the central member or is a variant on a variant. The theoretical claim being made is that a polysemous lexical item is a radial category of senses.

[1]This paper is a shortened version of a discussion of polysemy and image-schemas that appears as case study 2 in Lakoff 1987. It is published here with permission of the University of Chicago Press. That study, in turn, was based on Brugman's University of California M.A. thesis (Brugman 1981). Portions of this work were supported by grant no. BNS-8310445 of the National Science Foundation and a grant by the Sloan Foundation to the Institute of Cognitive Studies at the University of California, Berkeley.

What is important for our purposes is that the kind of network structure found here is not made up ad hoc to characterize this set of facts. Instead, this is a common category structure that occurs in domains other than the lexicon.

There is an important consequence of using the general theory of radial categories to characterize polysemy. In the general theory, the links between members of the network are not arbitrary. The theory of radial categories comes with a characterization of possible link types. In the case of polysemy, the link types are the types of relations linking the senses of the word. In general, some of the links may involve shared information, some may involve the relation between a general and a specific case, and some may be metaphoric. In the case under discussion, most of the links are what we have called "image-schema transformations." But overall there is only a small number of types of relations between senses of words, and this study is one of many that is being carried out in an attempt to figure out what they are.

Such studies are significant in a number of respects. They show that the relations between senses are not arbitrary, but are rather **principled, systematic,** and **recurrent** throughout the lexicon. Moreover, the relationships are natural, in the sense that they are either relationships that arise naturally within the cognitive system, or they are characterized by metaphors that have an independent existence in the conceptual system. From an explanatory point of view, the natural and independently motivated character of the links allows us to explain why polysemy should exist as a general phenomenon. From the point of view of language processing studies, it suggests that the lexicon has a structure that is made use of in processing.

Cognitive Topology versus Semantic Features

The traditional mode of representing lexical information, whether in linguistics or in cognitive science, has been the use of semantic features. Semantic features arise in general within the symbol manipulation paradigm in cognitive science: they are finitary symbols with no inherent meaning, but are to be made meaningful by being connected to things in the world. The features have no inherent structure, and any relationships among them are to be specified by a calculus characterizing permissible operations on the symbols.

One reason that we have chosen the example of *over* is that it demonstrates the inadequacy of feature analyses and shows the need for an oriented cognitive topology, which characterizes structures oriented relative to the human body that apply generally to spatial situations, structures like paths, bounded regions, tops, etc. Structures in a cognitive topology differ from semantic features in a number of ways: they are inherently meaningful (arising from sensory-motor operations), they have an inherent structure, they are analog rather than finitary, and the relationships among them arise naturally via the operation of the human sensory-motor system.

The evidence in this paper suggests that there are two respects in which cognitive topology is superior to feature analysis for *over*:

> The topological properties of the concepts are necessary to characterize "image-schema transformations" in terms that are cognitively natural, rather than in terms of an arbitrary calculus.

We will discuss this issue at the end of the paper. To facilitate the comparison between topological and feature analyses, we will use both types of representation in this paper. The drawings indicate the topological representation, while the names (such as 1.VX.C.E) indicate feature representations. We will demonstrate that, while feature representations might be useful in computer simulations, only topological representations characterize the cognitive reality of the meanings of words like *over*.

> Topological concepts are needed in order to account for how prepositions can be used to characterize an infinity of visual scenes.

The semantics of even the most basic spatial senses of *over* is such that a feature analysis simply will not do. Take, for example, the sense of *over* in "The ball went over the net." Given a scene with a ball moving with respect to a net, there is an infinity of trajectories of the ball relative to the net that *over* will fit and another infinity that it will not fit. To characterize those two infinities, one needs concepts that generalize over possible trajectories and properties of the landscape: this sense of *over* requires two bounded regions, and a path from one to the other that is oriented vertically relative to the net. In short, what is needed is an oriented cognitive topology with elementary structures (paths, bounded regions), orientations (vertical), and means of fitting them together into an overall gestalt.

Our reason for going into this issue in such detail is as follows: It is sometimes claimed that the symbol manipulation paradigm is necessary to account adequately for natural language. In fact, the reverse is true. The symbol manipulation paradigm cannot account for natural language semantics. For a lengthy discussion, see Langacker 1987.

Two Levels of Prototype Strcuture

The lexical representation of *over* actually constains two levels of topological structure. First, each sense of *over* is a complex topological structure. Second, all the senses together form a radial category, which is itself a complex topological structure. It is crucial to distinguish these two levels:

The first is the level of semantic content and the second is the level at which that content is structured in the lexicon.

Correspondingly, there are two levels of prototype structure—one at each level. At the second level, the level of lexical structure, the central sense in the radial category is the prototypical sense of *over*. But at the first level, the level of semantic content for each particular sense of *over*, the nature of prototypicality is quite different. At this level, prototypicality concerns the degree of fit of some real-world relation to an individual sense of the word. For example, consider "The plane flew over the mountain." The best fit is where the path goes right above the center of the mountain. As the path of flight moves away from the center of the mountain, the degree of fit lessens.

This introduction does not provide the whole story: we will not fully motivate the independent existence of all principles that relate senses, and we cannot do justice to the question of how contrasting lexical categories and the conventionalization of boundaries figure in a full semantic description of this item. The chief concern here is to provide a detailed example of how a lexical ambiguity of a specific kind can be given a reasonably complete description and to show the kinds of theoretical apparatus required for that description. To sum up: the critical features of this description are that feature-based descriptions are inadequate, that a topological representation appears to be needed, and that the senses of a polysemous item form a radially-structured lexical network.

run on next pg.

The Problem

To get some sense of the problem, let us consider a handful of the senses of *over*:

-The painting is *over* the mantle.
-The plane is flying *over* the hill.
-Sam is walking *over* the hill.
-Sam lives *over* the hill.
-The wall fell *over*.
-Sam turned the page *over*.
-Sam turned *over*.
-She spread the tablecloth *over* the table.
-The guards were posted all *over* the hill.

-The play is *over*.
-Do it *over*, but don't *over*do it.
-Look *over* my corrections, and don't *over*look any of them.
-You made *over* a hundred errors.

Even this small number of examples shows enormous complexity. The problem Brugman undertook was how to describe all these senses and the relations among them. The analysis we will be presenting is a minor refinement of Brugman's analysis. Let us begin with what Brugman found to be the central sense.

The Above-Across Sense

The central sense of *over* combines elements of both *above* and *across*.

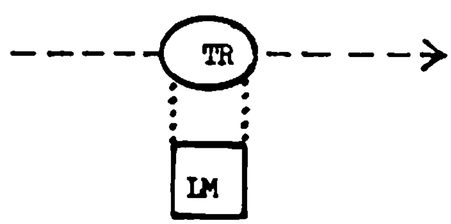

Figure 1: The plane flew over.
Name: Schema 1

In this example, the plane is understood as a trajector (TR) oriented relative to a landmark (LM). TR and LM are generalizations of the concepts figure and ground (Langacker, 1983). In this case the landmark is unspecified. The arrow in the figure represents the PATH that the TR is moving along. The LM is what the plane is flying over. The PATH is *above* the LM. The dotted lines indicate the extreme boundaries of the landmark. The PATH goes all the way *across* the landmark from the boundary on one side to the boundary on the other. Although the figure indicates noncontact between the TR and LM, the central sense is neutral on the issue of contact. As we will see shortly, there are instances with contact and instances without contact. In this respect the schema cannot be drawn with complete accuracy. Any drawing would have to indicate contact or the lack of it. The image-schema is neutral, and that is part of what makes it schematic. What we have here is an abstract schema that cannot itself be imaged concretely, but which structures images. We will return below to the question of what it means for an image-schema to structure an image.

Let us now turn to some special cases of the schema in figure 1. These are instances of the schema that are arrived at by adding information, in particular, by further specifying the nature of the landmark and by specifying whether or not there is contact. We will consider four kinds of landmark specifications: (1) LM is a point, that is, the landmark is viewed as a point with no internal structure. (2) LM is extended, that is, the landmark extends over a distance or area. (3) LM is vertical. (4) LM is both extended and vertical. For each such case, we will consider two further specifications: contact between TR and LM, and noncontact. Each schema will be named using the following abbreviations: X: extended, V: vertical, C:contact, NC: no contact. Thus, the schema name '1.VX.C' stands for the instance of schema 1 with vertical extended landmark (VX) and contact (C) between the LM and TR.

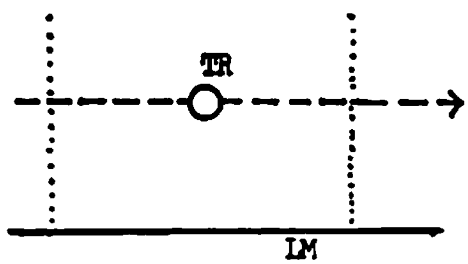

Figure 2: The bird flew over the yard.
Name: Schema 1.X.NC

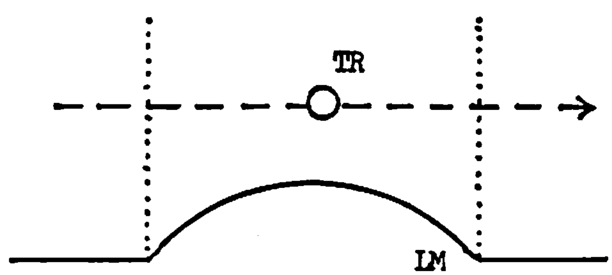

Figure 3: The plane flew over the hill.
Name: Schema 1.VX.NC

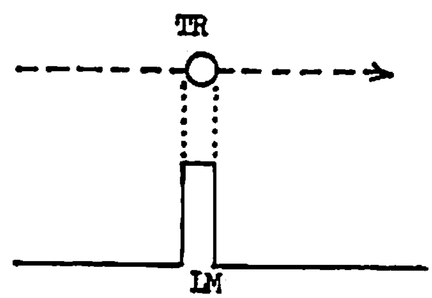

Figure 4: The bird flew over the wall.
Name: Schema 1.V.NC

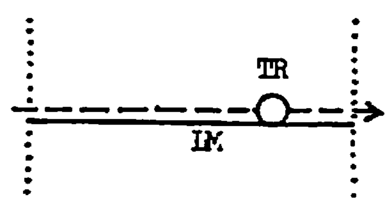

Figure 5: Sam drove over the bridge.
Name: Schema 1.X.C

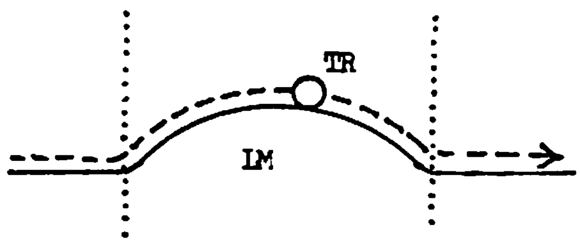

Figure 6: Sam walked over the hill.
Name: Schema 1.VX.C

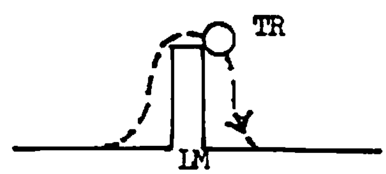

Figure 7: Sam climbed over the wall.
Name: Schema 1.V.C

The schemas in Figures 2 - 7 can be related by a diagram of the following sort:

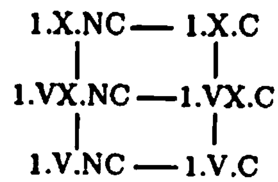

Figure 8: Links among schemas

These links indicate similarity. Thus, all the the schemas are linked, as are all the schemas that share noncontact. Moreover, each pair of schemas that share everything except for the contact parameter are linked. In addition, they are all linked to schema 1, since they are all instances of that schema.

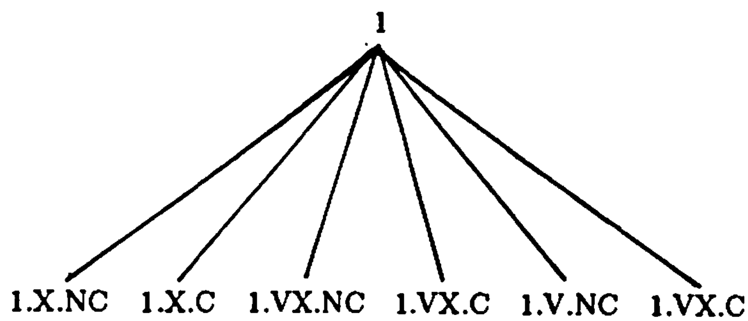

Figure 9

The schemas in Figures 2 - 7 can be viewed in two equivalent ways. Take, for example, a sentence like *Sam walked over the hill* in Figure 6. We can think of *over* in this sentence as being represented by the minimally-specified schema 1 of Figure 1, and we can think of the additional information as being added by the object and the verb. Thus, a hill is vertical and extended (VX) and walking requires contact (C) with the ground. Let us refer to this as the *minimal specification interpretation*. Equivalently, we can view the minimally-specified *over* of Figure 1 as generating all the fully-specified schemas of Figure 2 - 7. On this *full specification interpretation*, we can think of the *over* in *Sam walked over the hill* as having the full specification of schema 1.VX.C in

Figure 7. The verb *walk* would then match the contact (C) specification, and the direct object *hill* would match the vertical extended (VX) specification. The difference is whether the verb and direct object *add* the VX and C information or whether they *match* it.

These two interpretations make slightly different claims about the lexical representation of *over* in these sentences. On the minimal specification interpretation, only schema 1 exists in the lexicon; the other schemas all result from information added by the verb and direct object. On the full specification interpretation, there is a lexical representation for all these schemas; the more specific schemas are generated by schema 1 plus the general parameters we have discussed: C-NC and X-VX-V.

On the basis of what we have said so far, these two interpretations are completely equivalent; there is no empirical difference between them, and no a priori reason to choose between them. There is, however, additional evidence that favors the full specification interpretation, and we will be citing it throughout the remainder of this case study. We will be arguing that the senses of *over* form a chain with schema 1 at the center. On the full specification interpretation, the schemas in Figures 2 - 7 are part of that chain. Some of those schemas form links to other senses. The existence of such links suggests that the full specification interpretation is correct. Consider the following case, where there is a focus on the end-point of the path. We will use the abbreviation E in naming schemas where there is end-point focus.

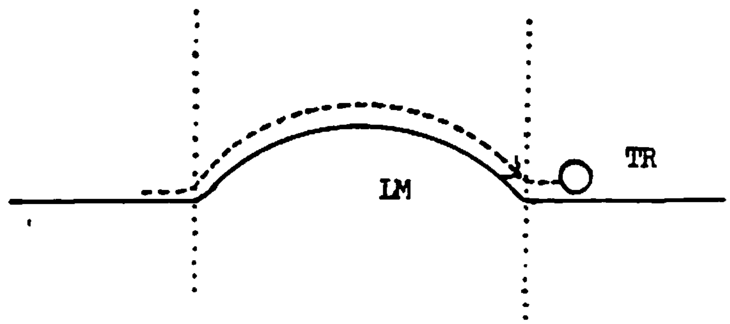

Figure 10: Sam lives over the hill.
Name: Schema 1.VX.C.E

In Figure 9, there is an understood path that goes over the hill, and Sam lives at the end of that path. The end-point focus is not added by anything in the sentence, neither *hill*, nor *lives*, nor *Sam*. *Over* here has an additional sense which is one step away from schema 1.VX.C -- a sense in which end-point focus (E) is added to yield schema 1.VX.C.E.

But end-point focus cannot be freely added to just any of the schemas in Figures 2 - 7. It can only be added to those with an extended landmark, as in the following case.

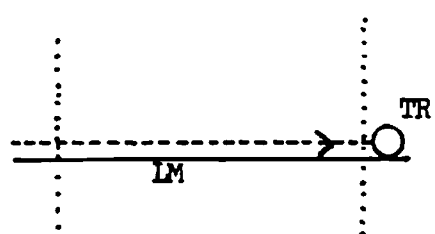

Figure 11: Sausalito is over the bridge.
Name: Schema 1.X.C.E

In these cases, *over* has the sense of 'on the other side of' as a result of end-point focus. However, *over* does not in general mean 'on the other side of'. For example, sentences like *Sam lives over the wall* and *Sam is standing over the door*, if they occur at all, cannot mean that he lives, or is standing, on the other side of the door and the wall. And a sentence like *Sam is sitting over the spot*, can only mean that he is sitting *on* it, not that he is sitting on the other side of it. Thus there is no end-point focus schema corresponding to schema 1.V.C of Figure 6.

Assuming the full specification interpretation, we can extend the chain in Figure 8 to include the schemas in Figures 10 and 11:

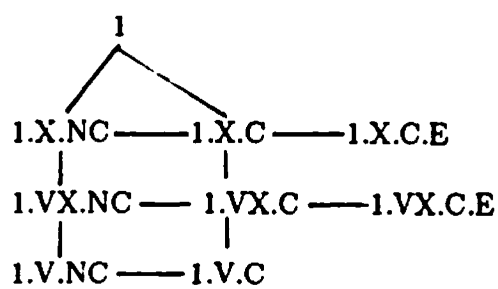

Figure 12: Links among schemas

So far, we have considered two types of links among schemas: *instance links* and *similarity links*. We can represent these as follows, letting 0 stand for null and A for some subschema.

Instance links: 0 <--> A

Similarity links: A <--> A

For example, the link between schema 1 and schema 1.V.C is an instance link, where A = V.C. And the link between schema 1.VX.NC and schema 1.VX.C is a similarity link, where A = 1.VX.

So far, we have looked only at instances of the *above-across* sense. And we have only looked at the least interesting links between schemas. Let us now turn to other senses and more interesting kinds of links.

The *Above* Sense

Over has a stative sense, with no PATH. It is roughly equivalent in meaning to *above*.

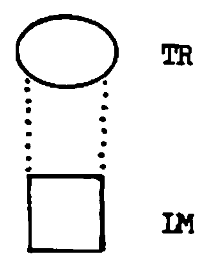

Figure 13: Hang the painting over the fireplace.
Name: Schema 2

Schema 2 has no particular constraints on either the TR or LM. It is linked to schema 1 in that it has the TR above the LM. However, it differs from schema 1 in two respects: First, it has no PATH and no boundaries; in other words, the *across*-sense is missing. Second, it does not permit contact between the TR and LM. The no-contact requirement can be seen in examples like he helicopter is hovering over the hill. If the helicopter lands, it is no longer *over* the hill, it is *on* the hill.

From time to time, linguists have suggested that schema 2 is *the core meaning* of the preposition *over*, that is, that schema 2 is present in all the uses of *over* as a preposition. It should be clear from what we have seen so far that this is false. Since schema 2 requires no contact, it cannot be present in those cases where contact occurs, for example, in schema 1.X.C exemplified by *Sam* *drove over the bridge*. Schema 2 also does not occur in the cases of end-point focus, such as schema 1.VX.C.E, which is exemplified by *Sam lives over the hill*. In this case, the TR is not above the LM.

One of the instances of schema 2 is the case where the TR is one-dimensional (which we will abbreviate as 1DTR).

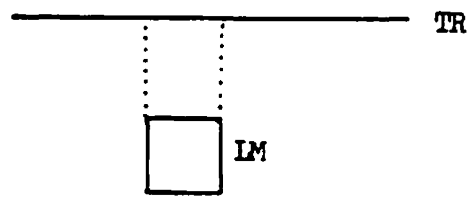

Figure 14: The power line stretches over the yard.
Name: Schema 2.1DTR

This schema is a minimal variant of schema 1.X.NC, exemplified by *The bird flew over the yard*, as shown in Figure 2. The extended path in Figure 2 corresponds to the one-dimensional solid trajector in Figure 14. We will call this kind of link between schemas a *transformational link*. This particular link between an extended path (X.P) and a one-dimensional trajector (1DTR) will be represented as:

X.P <—> 1DTR

This relationship is not directly reflected in the naming system for schemas that we have adopted. However, we can state the relationship more systematically if we do a little renaming of a sort that reflects image-schema decompositions. Let us use ABV for the 'above' subschema. And let us use PATH (P) for the 'across' subschema. Schema 1 would be renamed ABV.P, and Schema 1.X.NC of Figure 2 would be renamed ABV.NC.X.P. This name would reflect the fact that in this schema the TR is moving above (ABV) the LM, along a path (P), where the landmark is extended (X) and there is no contact between TR and LM (NC). Correspondingly, schema 2 would be renamed ABV.NC, and schema 2.1DTR in Figure 15 would be renamed ABV.NC.1DTR.

Schema 1.X.NC = ABV.NC.X.P
Schema 2.1DTR = ABV.NC.1DTR

This decomposition displays the relationship between the schemas directly. The schemas are transforms of one another, given the transformational link X.P <—> 1DTR.

It is important to bear in mind the difference between similarity links and transformational links. In the case of similarity links, the link is defined by shared subschemas. In the relationship described above, there are, indeed, shared subschemas: both schemas contain ABV.NC. But the transformational link is not a matter of *shared* subschemas, but of *related* subschemas.

The links among the schemas that we have described so far can be represented by the following diagram.

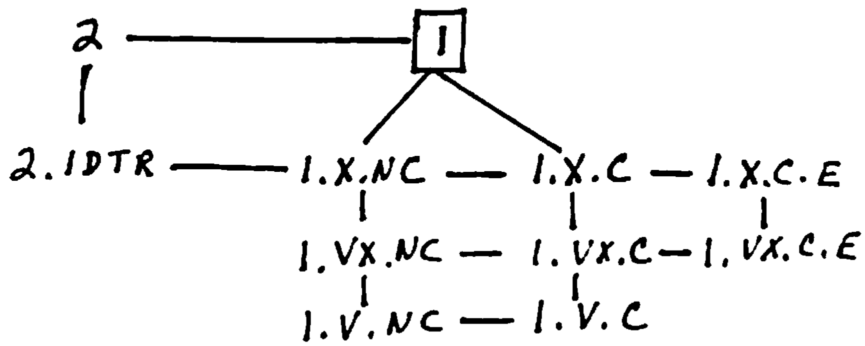

Figure 15: Links Among Schemas

The *Covering* Senses

There is a group of schemas for *over* that have to do with covering. This group is linked to the grid of Figure 15 in two ways. The basic covering schema is a variant of Schema 2, where the TR is at least 2-dimensional and extends across the boundaries of the LM.

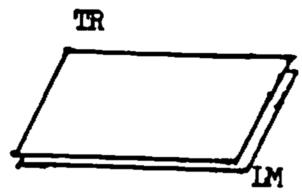

Figure 16: The board is over the hole.
Name: Schema 3

There are two differences between Schema 2 and Schema 3. The first is that Schema 2 is unspecified for the dimension of the trajector, while schema 3 must be at least 2-dimensional. But while schema 2 requires noncontact, schema 3 is neutral with respect to contact, allowing both contact and lack of it.

There is a minimal variant of schema 3 in which the TR moves to the position in schema 3. This schema is composed of schema 3 plus a path (P) indicating motion to the final position.

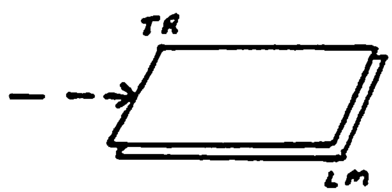

Figure 17: The city clouded over.
Name: Schema 3.P.E

Schema 3.P.E is linked to Schema 1. It shares motion of the TR above and across the LM. It also shares a lack of specification for contact. Schema 3.P.E differs from schema 1 in two ways. It is specified for the dimension of the trajector and it has end-point focus, which indicates that the final state is that of schema 3.

There are two covering schemas in which *over* is paired with a mass quantifier that quantifies regions of the landmark, e.g., *all, most, a lot of, entire,* etc. The quantifier *all* may combine with *over* in this sense to form the unit *all over*. The first of these two schemas has a *multiplex* (MX) trajector, that is, a trajector made up of many individuals.

-He has hair over most of his body.
-There are specks of paint all over the rug.
-There is sagebrush over the entire valley floor.

In these cases, the individuals -- the individual hairs, specks of paint, and bushes -- don't completely cover the part of the landmark quantified over. Rather, the landmark has small regions which jointly cover its surface (or most of it), and there is at least one trajector in each region.

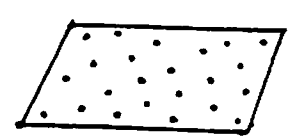

Figure 18: The guards were posted all over the hill.
Name: Schema 3.MX

The relationship between schema 3 and schema 3.MX is the relationship between a continuous region (or mass) and a multiplex entity. Such relationships are very common in language. Compare *cows* (multiplex) and *cattle* (mass). Quantifiers like *all* and *most* can occur with either masses (*all gold, most wine*) or multiplex entities (*all ducks, most trees*). The relationship between multiplex entities and masses is a natural visual relationship. Imagine a large herd of cows up close -- close enough to pick out the individual cows. Now imagine yourself moving back until you can no longer pick out the individual cows. What you perceive is a mass. There is a point at which you cease making out the individuals and start perceiving a mass. It is this perceptual experience upon which the relationship between multiplex entities and masses rests. The image transformation that relates multiplex entities and masses characterizes the link between schema 3 and schema 3.MX. We can characterize that transformational link as follows:

MX ⟷ MS

There is a second covering schema for *over* in which *over* is associated with a mass quantifier. It is a minimal variant on schema 3.MX in which the points representing the multiplex entity of 3.MX are joined to form a path (P) which 'covers' the landmark. Examples are:

-I walked all over the hill.
-We've hiked over most of the Sierras.
-I've hitchhiked over the entire country.

We can represent this schema as follows:

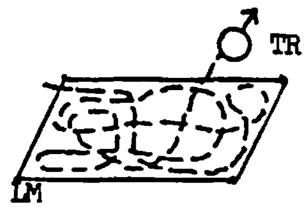

Figure 19: I walked all over the hill.
Name: Schema 3.MX.P

This schema is linked to schema 3.MX by an image transformation that forms a path through a collection of points. We will represent this transformational linkage as:

MX ⟷ MX.P

Schema 3.MX.P is also minimally linked to schema 3.P. In schema 3.P, the landmark is gradually covered as the trajector moves along the path. This is also true in schema 3.MX.P.

The covering schemas all have variants in which the relationship between the TR need not be above the LM. In all cases, however, there must be an understood viewpoint in which the TR is between the viewpoint and the LM. In other words, the TR is blocking accessibility of vision to at least some part of the landmark.

-There was a veil over her face.
-As the rain came down, it froze and ice spread all over the windshield.
-There were flies all over the ceiling.
-The spider had crawled all over the ceiling.

We will refer to these as *rotated* (RO) schemas, though with no suggestion that there is actual mental rotation degree-by-degree involved. One might suggest that instead of rotation from the vertical, there is simply a lack of specification of orientation. If there were, we would expect that the contact restrictions would be the same in all orientations. However, they are not. The rotated versions of the MX schemas -- 3.MX and 3.MX.P -- require contact, while the unrotated versions do not. Here are some typical examples that illustrate the distinction:

-Superman flew all over downtown Metropolis. (TR above LM, noncontact)
-*Superman flew all over the canyon walls. (TR not above LM, noncontact)
-Harry climbed all over the canyon walls. (TR not above LM, contact)

Thus, Superman's flying *alongside* the canyon walls does not constitute flying *over* them.

We will add 'RO' to the names of the unrotated covering schemas to yield names for the corresponding covering schemas. The rotated covering schemas have the following names: 3.RO, 3.P.RO, 3.MX.RO, and 3.MX.P.RO. Figure 20 is a diagram indicating the links among the covering schemas and the links to the other *over* schemas. And Figure 21 indicates the overall linkage among the schemas discussed so far.

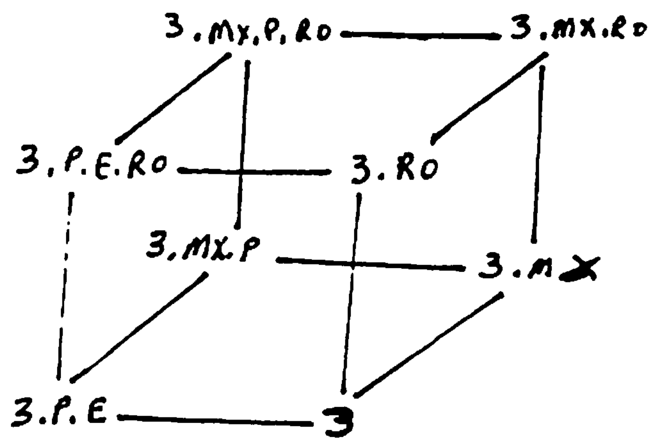

Figure 20: Links Among Covering Schemas

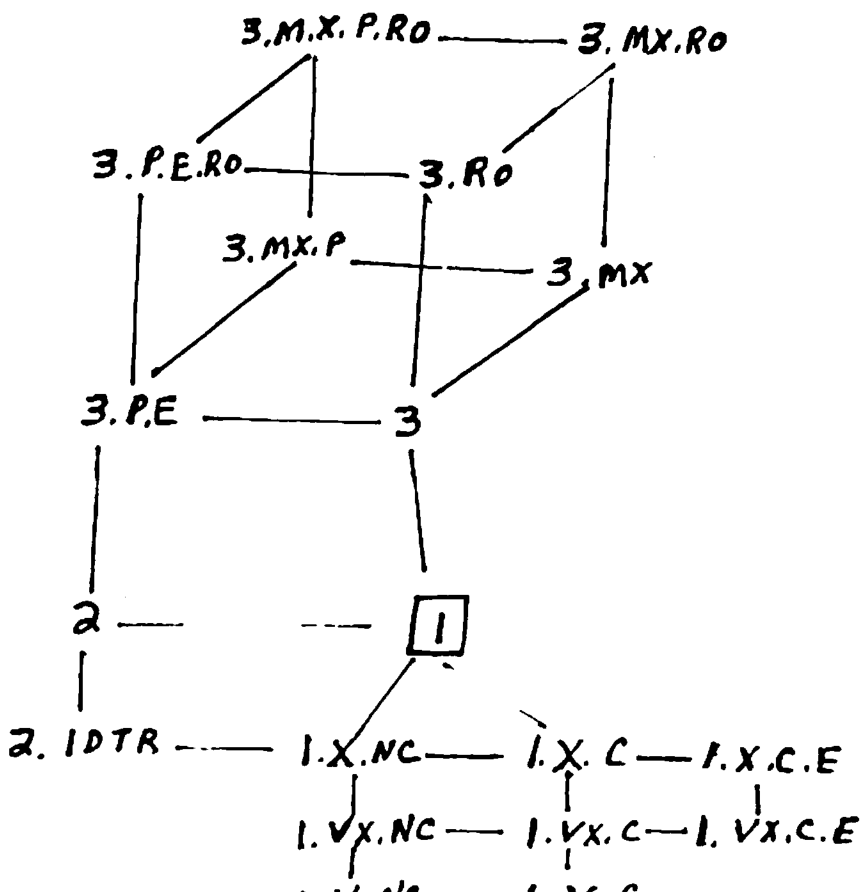

Figure 21: Links Among Schemas Discussed So Far

The Reflexive Schemas

Perhaps the most remarkable of the discoveries made by Lindner (1981, 1982) was the discovery of *reflexive trajectors*. The concept can be illustrated most simply using the example of *out*. The simplest use of *out* occurs in cases like *Harry ran out of the room*. We can represent this by the following schema.

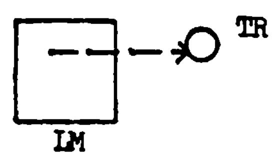

Figure 22: Harry ran out of the room.

Here the container (the room) is the landmark, and the trajector (Harry) moves from the interior to the exterior of the room. But this schema won't do for cases of *out* like:

-The syrup spread out over the table.
-The posse spread out over the hill.
-They stretched out the taffy.
-We rolled out the carpet.

Here the relevant trajectors are the syrup, the posse, the taffy, and the carpet. But they are not moving *out* with respect to any other landmark. Take the case of the syrup. Pour some syrup on a table. It will have a certain outer boundary at first. But the boundary moves. Some of the syrup that was inside the initial boundary is now outside that initial boundary. The syrup, or at least part of it, is moving 'out' relative to its own prior boundary. We can schematize this as follows.

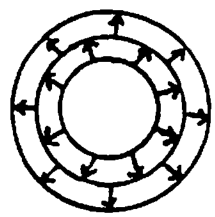

Figure 23: The syrup spread out.

In short, the syrup is its own landmark: TR = LM. Such a relation between a landmark and a trajector is called *reflexive*. Since there is only one entity under consideration, it is referred to as a *reflexive trajector*.

The '=' in 'TR = LM' is not strict identity; it is 'identity' of *part of* a bounded mass relative to itself as it *used to be* bounded. As we will see below, there are several ways in which 'TR = LM' can be realized. An important one is when parts of a single entity act as TR and other parts of the same entity act as LM. This kind of reflexive trajector occurs in the case of *over*. Consider examples like:

-Roll the log over.

Here a major part (roughly half) of the log is moving above and across the rest. That is, half the log is acting as landmark and the rest as trajector. The same is true in a case like

-Turn the paper over.

Both of these are variations on schema 1; they differ only in that LM = TR in the sense just described.

We can represent the schema for these cases as follows.

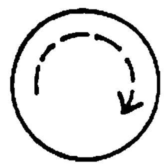

Figure 24: Roll the log over.
Name: Schema 4

Schema 4 can be viewed as a transform of schema 1, with schema 4 adding the condition TR = LM. We will represent such a transformational link as

NRF <—> RF

where 'NRF' means nonreflexive and 'RF' means reflexive. If we had chosen to name schema 4 according to its status as a variant of schema 1, we would have called it 1.RF.

The path of *over* in schema 4 traces a semi-circle above and across other parts of the thing being moved. We will refer to this as a *reflexive path*. There is a variant on schema 4 in which no part of the thing moving moves above or across any other part; instead, the entity as a whole traces the reflexive path. This occurs in cases like

-The fence fell over.
-Sam knocked over the lamp.

These are cases where the TR is initially vertical and moves so as to follow the last half of a reflexive path (RFP). The schema can be represented as:

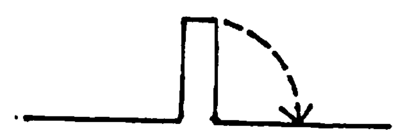

Figure 25: The fence fell over.
Name: Schema 4.RFP

The relationship between schemas 4 and 4.RFP can be stated as follows: In schema 4, half of the TR follows the whole reflexive path; in schema 4.RPF, all of the TR follows the last half of the reflexive path.

This schema is not only a variant of schema 4. It is also a minimal variant of one of the most common instances of schema 1, the instance that characterizes *over* in *The dog jumped over the fence.* In this case, there is a vertical landmark and the path of the trajector both begins and ends on the ground (G). This results in a semi-circular path.

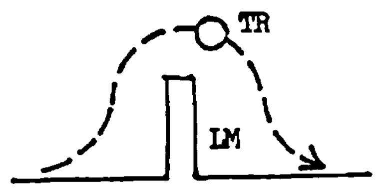

Figure 26: The dog jumped over the fence.
Name: 1.V.NC.G

If we take the reflexive transform of this schema, letting TR = LM, we get the schema of Figure 25 -- schema 4.RFP. Thus, schema 4.RFP has close links to two other schemas.

The Excess Schema

When *over* is used as a prefix, it indicates excess, as in:

-The bathtub overflowed.
-I overate.
-Don't overextend yourself.

The excess schema for *over* is a variation on one of the instances of schema 1, in particular, schema 1.X.C.E of Figure 11. In schema 1.X.C.E, there is an extended landmark; the trajector has moved across the boundaries of the landmark, and there is focus on the end-point, which is past the boundary. The excess schema has all these characteristics. It differs from schema 1.X.C.E in three respects. First, it is oriented vertically, with the end-point at the top. Second, it is understood as indicating amount via the MORE IS UP metaphor (cf. Lakoff and Johnson 1980). Third, the path is taken as indicating the course of an activity, via the ACTIVITY IS A JOURNEY metaphor. Fourth, the boundary point is taken as the upper limit of what is normal for that activity. Thus, being beyond that boundary point indicates excess.

The excess schema is therefore not merely an image-schema, but a complex schema that makes use of an image-schema (schema 1.X.C.E), an orientation transformation from horizontal to vertical, two metaphors (MORE IS UP and AN ACTIVITY IS A JOURNEY), and propositional information about what is normal. In link diagrams, we will refer to the excess schema as schema 5.

The Repetition Schema

One of the most common uses of *over* is to indicate repetition, as in

-Do it over.

Over here is used an an adverb. As in the case of the *over* of excess, the *over* of repetition makes use of a complex schema built on an instance of schema 1, namely, schema 1.X.C. This schema has an extended landmark, and indicates motion above and across it (cf. Figure 5). The repetition schema uses schema 1.X.C, and adds two metaphors to it. Again, the path is metaphorically understood as the course of the activity. This is via the very general ACTIVITY IS A JOURNEY metaphor. The landmark is understood metaphorically as an earlier completed performance of the activity. This is a special purpose metaphor, which is, to my knowledge, used only in this complex schema. This is the part of the repetition schema for *over* that is not motivated by an occurrence elsewhere in the conceptual system. It is what makes this sense of *over* somewhat idiosyncratic.

At this point, we are in a position to give a link diagram that shows a good deal of the complexity of *over*. In that diagram, we will refer to the repetition schema as schema 6. Figure 27 displays all the links we have discussed so far. A number of additional metaphorical links will be discussed below.

Figure 27 shows what is meant by a *radial* structure. Schema 1 occupies a central position; it and its instances are of primary importance in the system of links. The links correspond

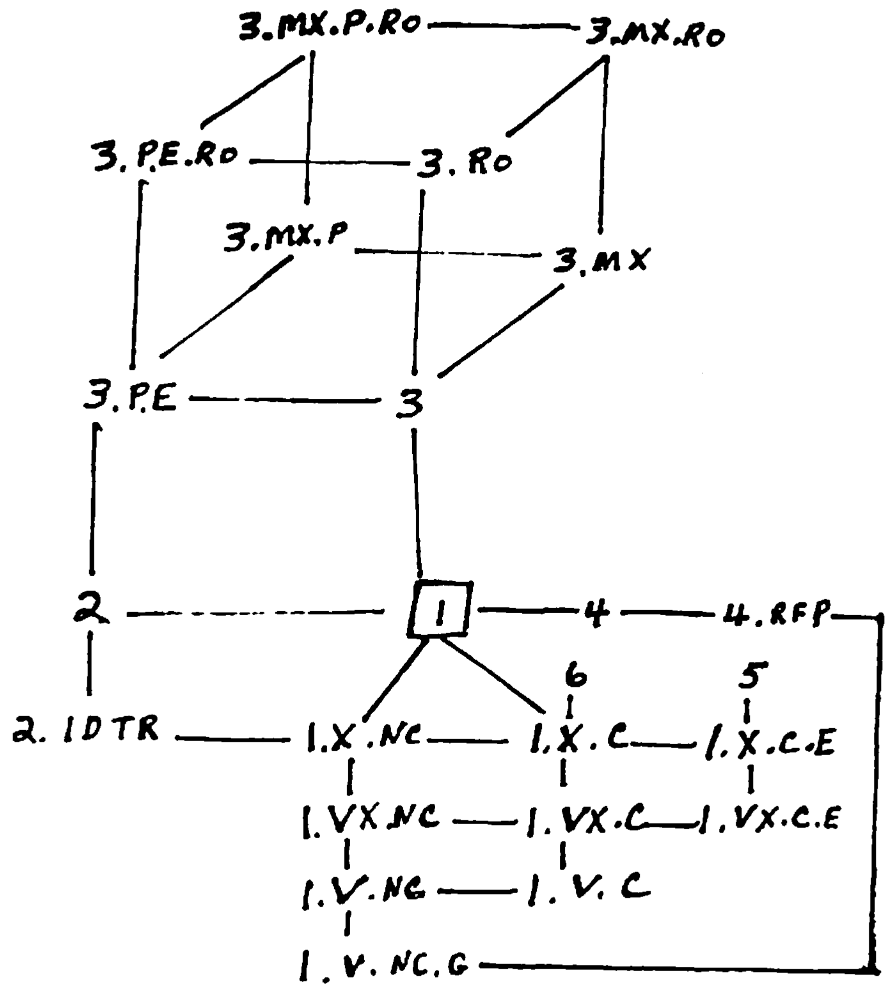
3.MX.P.Ro
3.MX.Ro
3.P.E.Ro
3.Ro
3.MX.P
3.MX
3.P.E
3
2
1
4
4.RFP
6
5
2.IDTR
1.X.NC
1.X.C
1.X.C.E
1.VX.NC
1.VX.C
1.VX.C.E
1.V.NC
1.V.C
1.V.NC.G

Figure 27

to what Wittgenstein called 'family resemblances'. The links are sometimes defined by shared properties. But frequently they are defined not by shared properties, but by transforms or by metaphors.

Some Metaphorical Senses

It is extremely common for metaphors to apply to image-schemas. A great many metaphorical models use a spatial domain as their source domain. Among the most common source domains for metaphorical models are containers, orientations, journeys (with paths and goals), vertical impediments, etc. In this section, we will give a number of cases where *over* has a metaphorical sense based on an image-schema discussed above.

-She has a strange power *over* me.

This is an instance of a very common metaphor: CONTROL IS UP; LACK OF CONTROL IS DOWN (cf. Lakoff and Johnson 1980). *Over* in this sentence is an extension of schema 2 (Fig. 14), where the trajector is simply above the landmark.

-Sam was passed *over* for promotion.

Here we have an instance of schema 1 (Fig. 1). Two metaphorical mappings apply to it. The first is CONTROL IS UP; LACK OF CONTROL IS DOWN. This entails that the person who passed over Sam was in control of Sam's status. The second metaphor the applies to this schema is another common one: CHOOSING IS TOUCHING. This occurs in such sentences as *He was tapped for service* and *The boss handpicked his successor.* Since the schema indicates that there is no contact, it is entailed that Sam was not chosen.

We are now in a position to make sense of the difference between *overlook* and *oversee.*

-You've *over*looked his accomplishments.
-We need to find someone who can *over*see this operation.

The *over* in *overlook* is based on schema 2.1DTR (Fig. 5). There are two metaphors involved. The first is a metaphor for understanding vision: SEEING IS TOUCHING. This occurs in examples like *I couldn't take my eyes off of her, Her eyes picked out every detail of the pattern, He undressed her with his eyes,* and *He fixed his gaze on the entrance.* According to this metaphor, one's gaze goes from one's eyes to what one sees. You see whatever your gaze touches. Under the metaphorical mapping, the path in schema 2.1DTR is the gaze. Since there is no contact in schema 2.1DTR, the metaphorical gaze doesn't touch the landmark; thus the subject of *overlook* is *not* looking at, and therefore does *not* see, the landmark. The second metaphor is the general MIND-AS-BODY metaphor (cf. Sweetser (1984)). The relevant aspect of that metaphor is the part in which LOOKING AT SOMETHING IS TAKING IT INTO CONSIDERATION. Accordingly, *I'll take a look at it* normally entails *I'll consider it.* Therefore, to overlook someone's accomplishments is *not* to take them into consideration.

The *over* in *oversee* is based on schema 2 (Fig. 15), in which the TR is above the LM. There a metaphor and a metonymy that are relevant to this example. The metaphor is CONTROL IS UP. Thus, the one who does the overseeing has control over the persons overseen. The metonymy is SEEING THAT SOMETHING IS DONE STANDS FOR MAKING SURE THAT IT IS DONE. This metonymy is based on an idealized model in which making sure of something typically involves seeing it. Because of this metonymic relation, *See that he gets his money* means *Make sure that he gets his money.* Thus, to *oversee* means to be in control and make sure that something is done.

We can now compare *overlook* to *look over.*

-Look *over* my corrections, but don't *over*look any of them.

The *over* in *look over* is based on schema 3.MX.P (Fig. 18), and the SEEING IS TOUCHING metaphor. The resulting complex schema is one in which the subject's gaze traces a path that 'covers' the direct object, *corrections.* In the resulting schema, the gaze does make contact with the landmark. The MIND-AS-BODY metaphor again yields a sense of *look* in which looking at something involves taking it into consideration. Thus, when one looks over X, one directs one's attention to a representative sampling that 'covers' X, and one takes into consideration each subpart that one directs attention to.

Motivation

Before we go on, it is worth commenting on what is and what is not being explained in these analyses. We are not explaining why *oversee, overlook,* and *look over* mean what they mean. Their meanings cannot be predicted from the meanings of *over, look,* and *see*. But their meanings are not completely arbitrary. Given the range of spatial meanings of *over*, and given the metaphors present in the conceptual system that English is based on, it *makes sense* for these words to have these meanings. We are explaining just why it makes sense and what kind of sense it makes.

In each of these cases, the metaphorical and metonymic models exist in the conceptual system independently of the given expression. For example, we understand seeing metaphorically in terms of of a gaze that goes out of one's eyes and touches the object seen. This metaphorical understanding is present regardless of whether any of the expressions just discussed have those meanings. Similarly, the schemas for *over* exist for expressions in the spatial domain independently of the existence of *oversee, overlook,* and *look over*. What one learns when one learns these words is which of the independently existing components of their meaning are actually utilized. Each of these expressions is a specialized 'assembly' of independently existing parts. The only arbitrariness involved is the the knowledge that such an assembly exists.

The psychological claim being made here is that it is easier to learn, remember, and use such assemblies which use existing patterns than it is to learn, remember, and use words whose meaning is inconsistent with existing patterns. What is being explained is not why those expressions mean what they mean, but why those are natural meanings for them to have. Thus, if one is going to have a word that means 'to fail to take into consideration', it is more natural to use *overlook* than to use an existing unrelated word like *sew*, or a complex word whose components are in conflict with the meaning, such as *underplan*, or *taste at*, or *rekick*. It is common sense that such expressions would not be used with such a meaning, and we are characterizing the nature of that 'common sense'.

As we have mentioned before, such an explanation requires going beyond the predictable/arbitrary dichotomy. It requires introducing the concept of *motivation*. Thus, the meaning of *overlook*, though not predictable, is motivated -- motivated by one of the spatial schemas for *over* and by two metaphors in the conceptual system. Similarly, all of the noncentral schemas for *over* in the chain given in Fig. 27 are motivated -- motivated by other senses and by principles of linking.

More Metaphorical Senses

The are some additional common metaphorical senses of *over* that are worth discussing. Take *get over* in

-Harry still hasn't gotten *over* his divorce.

This use of *over* is based on schema 1.VX.C (Fig. 6), and two metaphors. In the first metaphor, obstacles are understood in terms of vertical landmarks -- which may be extended or not. This metaphorical model is the basis for expressions such as *There is nothing standing in your way.* The second metaphorical model is one that understands LIFE as a JOURNEY. This occurs in sentences like *It's time to get on with your life*. In the above use, the divorce in an obstacle (metaphorically, a vertical extended landmark) on the path defined by life's journey.

-Pete Rose is over the hill.

Over the hill makes use of schema 1.VX.C.E (Fig. 10) and a metaphor for understanding a career in terms of a journey over a vertical extended landmark like a hill. In this metaphorical model of a career, one *starts at the bottom*, may *go all the way to the top*, and then *goes downhill*. Thus, *over the hill* means that one has already reached and passed the peak of one's career and will never have that high a stature again.

-The rebels overthrew the government.

This is an instance of schema 4.RFP (Fig. 25) and the CONTROL IS UP metaphor. Before the event takes place, the government is in control (metaphorically upright), and afterwards it is not in control (metaphorically down).

-He turned the question over in his mind.

This is an instance of schema 4 (Fig. 24), plus an instance of the MIND-AS-BODY metaphor in which THINKING ABOUT SOMETHING IS EXAMINING IT. This metaphorical model occurs in such sentences as *Let us now examine the question of factory chickens*. In examining a physical object, one turns it over in order to get a look at all sides of it. Questions are metaphorically understood as having sides, and when one turns a question over in one's mind, one is examining all sides of it.

-The play is over.

Here we have an instance of schema 1.X.C.E (Fig. 11). In general, activities with a prescribed structure are understood as extended landmarks, and performing such an activity is understood metaphorically as traveling along a prescribed path over that landmark. When one gets to the end, the activity is *over*. Thus, games, plays, and political campaigns can be characterized at their end as being *over*.

Image-Schemas As Links Between Perception and Reason

Two of our major sources of information are vision and language. We can gain information through either perceiving something directly or being told it. And we can reason about that information, no matter what its source. We can even reason using information from both sources simultaneously, which suggests that it is possible for us to encode information from both sources in a single format. I would like to suggest that image-schemas provide such a format.

It is my guess that image-schemas play a central role in both perception and reason. I believe that they structure our perceptions and that their structure is made use of in reason. The analysis of *over* that we have just given is rich enough for us to discuss such questions in some detail. Let us begin with the following question. Are the image-schema transformations we have discussed natural, and if so, what is the source of their 'naturalness'?

The Nature of Image-Schema Transformations

There are certain very natural relationships among image-schemas, and these motivate polysemy, not just in one or two cases, but in case after case throughout the lexicon. Natural image-schema transformations play a central role in forming radial categories of senses. Take, for example, the end-point-focus transformation. It is common for words that have an image-schema with a path to also have the corresponding image-schema with a focus on the end-point of the path. We saw this in *over* in cases like

-Sam walked *over* the hill. (path)
-Sam lives *over* the hill. (end-of-path)

Pairs such as this are common.

-Harry walked *through* that doorway. (path)
-The passport office is *through* that doorway. (end-of-path)

-Sam walked *around* the corner. (path)
-Sam lives *around* the corner. (end-of-path)

-Harriet walked *across* the street. (path)
-Harriet lives *across* the street. (end-of-path)

-Mary walked *down* the road. (path)
-Mary lives *down* the road. (end-of-path)

-Sam walked *past* the post office. (path)
-Sam lives *past* the post office. (end-of-path)

It should be noted that although such pairs are common, they are not fully productive.

-Sam walked *by* the post office. (path)
-Sam lives *by* the post office. (= *near*; ≠ end-of-path)

Here, *by* has a path schema, but no corresponding end-point schema.

-Sam ran *from* the house. (path)
-Sam stood three feet *from* the house. (end-of-path)

-Sam ran *to* the house. (path)
-*Sam stood (three feet) *to* the house. (≠ end-of-path)

From allows both path and end-of-path schemas, but *to* only allows a path schema.

Path schemas are so naturally related to end-point schemas that people sometimes have to think twice to notice the difference. The same is true of the schema transformation that links multiplex and mass schemas. It is natural for words that have a mass schema to also have a multiplex schema.

-*All* men are mortal. (MX)
-*All* gold is yellow. (MS)

-She bought *a lot of* earrings. (MX)
-She bought *a lot of* jewelry. (MS)

This schema transformation, of course, doesn't hold for all quantifiers:

-She bought *two* earrings. (MX)

-*She bought *two* jewelry. (MS)

There are also verbs which have both schemas:

-He *poured* the juice through the sieve. (MS)
-The fans *poured* through the gates. (MX)

This will also work for other verbs of liquid movement, such as *spill, flow,* etc.

-The wine *spilled* out over the table. (MS)
-The fans *spilled* out over the field. (MX)

There is a special case of the multiplex-mass transformation in which the multiplex entity is a sequence of points and the mass is a one-dimensional trajector. A variety of prepositions permit both schemas.

-There are guards posted *along* the road. (MX)
-There is a fence *along* the road. (1DTR)

-He coughed *throughout* the concert. (MX)
-He slept *throughout* the concert. (1DTR)

-There were stains *down* his tie. (MX)
-There were stripes *down* his tie. (1DTR)

There is a natural relationship not only between a one-dimensional trajector and a sequence of points. There is also a natural relationship between a one-dimensional trajector and a zero-dimensional moving trajector that traces a path.

-Sam *went* to the top of the mountain. (0DMTR)
-The road *went* to the top of the mountain. (1DTR)

-Sam ran *through* the forest. (0DMTR)
-There is a road *through* the forest. (1DTR)

-Sam walked *across* the stream. (0DMTR)

-There is a bridge *across* the stream. (1DTR)

Finally, there is a natural relationship between nonreflexive and reflexive trajectors. Here are some examples:

-He stood *apart* from the crowd. (NRF)
-The book fell *apart*. (RF)

-She walked *up* to me. (NRF)
-Let's cuddle *up*. (RF)

-She poured the syrup *out* of the jar. (NRF)
-The syrup spread *out* over the pancakes. (RF)

Let us consider for a moment what is natural about these image-schema transformations.

Path-focus / end-point-focus: It is a common experience to follow the path of a moving object until it comes to rest, and then to focus on where it is. This corresponds to the path-focus / end-point-focus transformation.

Multiplex / mass: As one moves further away, a group of individuals at a certain point begins to be perceived as a mass. Similarly, a sequence of points is perceived as a continuous line when viewed from a distance.

0DMTR / 1DTR: When we perceive a continuously-moving object, we can mentally trace the path it is following. This capacity is reflected in the transformation linking zero-dimensional moving trajectors and a one-dimensional trajector.

NRF / RF: Given a perceived relationship between a TR and a LM which are two separate entities, it is possible to perceive the same relationship between different parts of the same entity.

In short, these schema transformations are anything but arbitrary. They are direct reflections of our visual experiences.

> **The fact that image-schemas are a reflection of our visual experience is hardly surprising, yet it plays a very important role in the theory of image-schemas. Perhaps we can see that significance most easily by contrasting the image-schema transformations we have described with the names we have given to them. Take the transformation name 'MX↔MS'. The names 'MX' and 'MS' are relatively arbitrary relative to the character of what they name: a group of individual entities and a mass. Our perceptions of things named are naturally related in our experience. The names are not naturally related in any way. The transformation is a natural relationship, but the name of the transformation is just a bunch of arbitrary symbols.**

This distinction is important because of certain versions of the Symbol Manipulation Paradigm. On one theory of image representation--the "propositional theory"--visual scenes are represented by arbitrary symbols which are linked together in network structures. Arbitrary symbols such as X and Y are taken as standing for some aspect of a scene, such as a point or an edge or a surface or an entire object. Other symbols are used to express relations among these symbols, for example, 'ABV(X,Y)' and 'C(X,Y)' might represent relations which are supposed to correspond to 'X is above Y' and 'X is in contact with Y', but which, so far as the computer is concerned, are just symbols. Such a symbolization describes how various parts--points, edges, surfaces, etc.--are related to one another. Objects in a scene are described using such symbolizations.

According to the Symbol Manipulation Paradigm as applied to visual information and mental imagery (Pylyshyn, 1981), only such propositional representations are mentally real, while images are not real. This view stems from taking the Symbol Manipulation Paradigm very seriously. Since digital computers work by the manipulation of such arbitrary symbols, the strong version of the computational theory of mind requires not only that visual perception and mental imagery be characterizable in such a 'propositional' form, but also that such symbolic representations, and only those, are mentally real.

The names that we have given to the image schemas, as well as to the image-schema transformations, are in keeping with feature representations. They have the properties that (1) they are arbitrary, in the sense that the internal structure of symbols plays no role in how the symbols interact or what they mean; (2) they are inherently meaningless, and have to be assigned meanings; and (3) they are finitary in nature. Such feature-style names are in these respects opposite from the corresponding image-schemas, which (1) are nonarbitrary, in the sense that they have an internal structure that plays a crucial role in what they mean and how they interact; (2) they are inherently meaningful; and (3) they are analog in nature.

Suppose that, instead of merely using such symbols as convenient names, we chose to take such a use of symbols seriously, as one would have to if one were to adopt the Symbol Manipulation Paradigm for Cognitive Science. According to that paradigm, topological representations such as image-schemas are not available as cognitive representations; all that is available are symbolic representations, which would look like the symbolic names we have given and would have their properties. The Symbol Manipulation Paradigm would thus make the implicit claim that the cognitively natural image-schema transformations of the sort we described did not exist. In their place would be arbitrary transformations relating the names we have given. Instead of a natural explanation of types of polysemy, we would have no explanation at all, but only an arbitrary description. We consider the lack of such explanatory force intolerable.

For instance, consider the relationship between the *over* of *The bird flew over the yard* and that of *The power line stretches over the yard.* These are adjacent senses in the network, linked by the natural relationship between a zero-dimensional moving trajector and a corresponding one-dimensional stationary trajector. If there were no such natural relation between the senses, they would not be adjacent in the network. Thus, the configuration of the network is anything but arbitrary; it is determined by an account of what constitutes a "natural" link type. Since it is the topological character of the representations that makes such relationships natural, it is that topological nature that makes the configurations of the networks nonarbitrary. If symbolic feature representations are substituted for the topological representations, then the naturalness of the relationships disappears and, with it, the explanation for what is linked to what. In short, there are natural reasons for the extensions of certain senses to other senses, and those reasons must be given in a cognitively adequate account of polysemy.

Conclusion

Systematic polysemy of the sort we have just seen is a pervasive phenomenon in language. It is so common and automatic that we often do not even notice it. Yet, as we have seen, simple basic processes like image-schema transformations and metaphors can interact to form rather large networks that characterize the natural relationships among the senses of polysemous words.

Our primary conclusion is that lists will not do; networks of the kind we have described are needed to characterize the relationships among the senses of polysemous words.

Our secondary conclusion is that features will not do. Oriented topological representations using image-schemas are necessary for two reasons: to account for the range of scenes that concepts like *over* can fit, and to account for the naturalness of image-schema transformations. This suggests that the Symbol Manipulation Paradigm, which cannot tolerate the existence of such representations using a cognitive topology, is inadequate for natural language semantics.

References

Brugman, Claudia. 1981. *Story of* Over. University of California, Berkeley M.A. Thesis. Available from the Indiana University Linguistics Club.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind.* University of Chicago Press.

Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By.* Chicago: University of Chicago Press.

Langacker, Ronald. 1983. Remarks on English Aspect. In P. Hopper, ed., *Tense and Aspect: Between Semantics and Pragmatics.* Amsterdam: John Benjamins. pp. 265-304.

Langacker, Ronald. 1987. *Foundations of Cognitive Grammar,* Vol. 1. Stanford: Stanford University Press.

Lindner, Susan. 1981. *A Lexico-Semantic Analysis of Verb-Particle Constructions with Up and Out,* University of California, San Diego Ph.D. Dissertation. Available from the Indiana University Linguistics Club.

Lindner, Susan. 1982. What Goes up Doesn't Necessarily Come Down: The Ins and Outs of Opposites. In *Papers from the Eighteenth Regional Meeting, Chicago Linguistic Society.* pp. 305-323.

Pylyshyn, Zenon. 1981. The Imagery Debate: Analogue Media versus Tacit Knowledge. *Psychological Review* 87,16-45.

Sweetser, Eve E. 1984. *Semantic Structure and Semantic Change.* University of California, Berkeley Ph.D. Dissertation.

Roman