Extending the Lexicon by Exploiting Subregularities*

Robert Wilensky
Division of Computer Science
University of California, Berkeley

1. Introduction

A standard question one asks in judging the scope of a NLP system is the size of its vocabulary. However, most assertions about the number of lexical items that a given AI natural language system can handle are misleading: Virtually every system includes at most a few senses of any given word. However, most words have quite a large number of senses, and quite a lot of information about each one is needed for the system to function adequately. By this measure, there is probably no system that has the full command of a single word.

We believe that underlying this inadequacy is an important theoretical shortcoming. To address it we advance a position that is in sharp contrast to the rather pervasive view of language as a system governed by general rules with memorized exceptions. In this view, the lexicon is relegated to the status of a large group of arbitrary facts which simply need to be memorized. While we acknowledge, and indeed, underscore, the need for this large collection of stored facts, we believe that this knowledge is neither unstructured nor arbitrary. Rather, the structure of lexical knowledge is complex and abundant, and, we believe, yields important insights into the nature of cognition.

Such a position may seem at first to be only of theoretical interest to those concerned with the traditional dichotomy between rules and memorized facts. From this view, if one must learn a body of knowledge in order to function competently as a language user, there can be no role for any alleged structure purported to be found in this knowledge: A fact is either arbitrary, or it is predictable. However, we believe that a compelling role does exist for this knowledge. It is exploited during language learning.

In this paper, we first examine the nature of word senses. We discuss some of our previous work on representing and exploiting an important subclass of the structure to be found there, namely, that based on metaphoric relations. This portion of our work has led to the construction of a system whose learning of new word senses is facilitated by its previous knowledge. We then discuss what other structure there may be in the lexicon and in the conceptual system of a language, and how we might generalize our approach to accommodate it.

In particular, we propose a new method of analogical reasoning, and apply it to hypothesize new polysemous word senses. This method is the basis for one of a number of knowledge acquisition devices to be included in DIRC (Domain Independent Retargetable Consultant), an intelligent, natural language-capable consultant kit that can be retargeted to different domains. DIRC is essentially ‘‘empty-UC’’ (UNIX Consultant, Wilensky et al., 1988); DIRC is to include the language and reasoning mechanisms of UC, plus a large grammar and a general lexicon. The system builder must then add domain knowledge, user knowledge and lexical knowledge for the area of interest.

To facilitate the system builder’s job, we provide a number of tools for knowledge acquisition. This paper

*The research reported here is the product of the Berkeley Artificial Intelligence and Natural Language Processing seminar; contributers include Michael Braverman, Narcisco Jaramillo, Dan Jurafsky, Eric Karlson, Marc Luria, James Martin, Peter Norvig, Michael Schiff, Nigel Ward, and Dekai Wu. This work also benefitted from discussions with George Lakoff, Eve Sweetser and Paul Kay. James Martin pointed the use of metaphor knowledge in learning; Peter Norvig’s lexical network work demonstrated the existence of other, useful lexical relations. This research was sponsored by the Defense Advanced Research Projects Agency (DoD), monitored by Space and Naval Warfare Systems Command under Contract N00039-88-C-0292 and by the Office of Naval Research, under contract N00014-89-J-3205.

describes a theory underlying just one of several approaches we are exploring to one part of this problem, that of lexical acquisition. While building DIRC is the context for the research described below, our primary interest is in developing this theory as a cognitive model of human language use and acquisition.

2. The Nature of Lexical Entries

Since we are concerned with extending the lexicon, it is necessary to first specify what assumptions we are making about lexical entries in general. It seems generally agreed upon that words often have multiple senses, can serve as different parts of speech, and manifest variable valence. Moreover, these facts about a given word go together. To use a trite example, the word ‘‘bank’’ can be a verb whose meaning has to do with using a financial institution, and when it does, it subcategorizes for a subject expressing the customer and a prepositional complement beginning with ‘‘at’’ that expresses the financial institution. Alternatively, when ‘‘bank’’ is a verb meaning to bang a billiard ball off a cushion, it subcategorizes for a subject expressing the agent and a direct object expressing the ball. Of course, the word has at least two meanings as a noun, none of which subcategorizes for anything.

We assume that each lexical entry contains this kind of information. We call each basic unit of lexical information a *lexical grammatical construction* (or *lex-con* for short). Each lex-con specifies a phonological form (i.e., a ‘‘word’’), grammatical information such as part of speech and valence, and a meaning.

This terminology reflects the intuition that words are just a kind of grammatical construction, each of which has grammatical and meaning information associated with it. However, nothing that we say in this paper depends strongly on this intuition. Indeed, we believe that lexical entries in many systems more or less resemble the lex-cons we describe.

The meaning in a lex-con is simply a concept that the word denotes. These concepts are supposed to be wholy non-linguistic, that is, they are objects in some internal representation language, and everything we know about these concepts (other than associated linguistic information) is spelled out in the conceptual domain. We call a concept associated with a word via a lex-con a *sense*: A sense is just a concept that happens to be lexicalized, i.e., that has gotten attached to a word.

We will use the term *lexical relation* to mean any systematic relation between lex-cons. We divide relations between lex-cons into three classes:

1) *Grammatical relations* – These are meant to capture the trivial case of grammatical variations of the same stem.

2) *Core relations* – Lex-cons are *core-related* when their senses share some component of meaning.

3) *Transforming relations* – These are relations between lex-cons one of whose senses is thought of as derived from the other through some process of meaning transformation.

Each class of relations is now described in greater detail.

(1) Grammatically-related Lex-cons

Lex-cons as described so far require a rather fine level of granularity. In particular, we require separate lex-cons for the cross-product of all grammatical, semantic, and phonological distinctions. So, for example, there is a distinct lex-con for the word ‘bank’ corresponding to the first person singular verb meaning to use a financial institution, another for the same word used as the second person plural verb with the same meaning, and others for all combinations of tenses, person, number and senses of this word and its variants.

However, the semantic relations between these lex-cons are in some sense trivial. Some of these lex-cons differ only in grammatical features like person, so their senses are identical; others, such as the tenses or number have different but completely generative meanings. Thus, even though the phonological relation between the words might vary, we do want to recognize these lex-cons as not having an independent status. We call two lex-cons *grammatically related* if their word senses are in one of a set of ''grammatical meaning relations''. Grammatical meaning relations are just semantic relations that are grammatically realized, such as singular-multiple, present-past, and especially, the identity relation.

Simply having grammatical relations between lex-cons does not allow one to state certain basic abstractions, however. For example, there is no entity corresponding to what we might intuitively think of as one of the verbs ''bank''. Instead, there is only a large collection of lex-cons whose senses are grammatically related. This is unfortunate, since most of the generalizations we would like to make are at a more abstract level.

Thus, we posit the existence of lex-cons that are incomplete in the sense that they are lacking information necessary for their participation in syntactic constructions*. Let us call these *abstract lex-cons*. For example, corresponding to what we think of intuitively as a verb (e.g., the intransitive verb 'bank' meaning to use a financial institution), there would be an abstract lex-con that omits person and number information, and specifies only the stem of the verb, but which still specifies the (timeless) sense of the verb, and its thematic roles. Such abstract lex-cons would be related to concrete (i.e., fully-specified) lex-cons by simple grammatical lexical relations. Thus, the concrete lex-con for the third person singular present tense of some verb might be related to the abstract lex-con for that verb by some ''third person singular present'' relation; of course, additional levels of abstraction may be desirable. Similarly, rather than relate a singular lex-con for ''boy'' directly to a plural lex-con for ''boys'', both might be related to an abstract lex-con whose sense is the concept of ''boy'', but which is neutral with respect to number.

We will call the most abstract lex-con of a hierarchy of lex-cons associated with a word a *basic lex-con* for that word; those lex-cons that are grammatically related to it we will call *grammatically derived* lex-cons. For example, assuming for the time being that there is a single sense of ''give'' involving transfer of possession, we would have a basic lex-con specifying the root and this sense; we would also have many derived lex-cons, for example, one for the first person singular past form, whose sense is that of the basic lex-con plus a predication involving time.

Almost all of what we are interested in talking about in this paper concerns relationships among basic lex-cons; the reader can assume that all subsequent uses of the term refers to lex-cons of this sort. (Actually, much of what we have to say later on seems to apply as well to the relations between basic lex-cons and grammatically derived lex-cons, but we shall not develop this point here.)

We note that while we speak of some lex-cons as being derived from others, of course the derivation is only semantically generative; the phonological forms may vary. Also, the basic lex-cons of even polysemous senses of the same word might be in the same grammatical relation to different concrete lex-cons. Thus, while the past tense of ''bank'' is ''banked'', no matter which verbal lex-con is in question, the past tense of ''hang'' is ''hanged'' or ''hung'', depending on the associated sense; similarly, ''work'' alternates between ''worked'' and ''wrought''; ''fly'' between ''flew'' and ''flied'' (as in ''The first batter flied out.''); ''ring'' between ''rang'' and ''ringed'' (as in ''The Margarita was ringed with salt.''). (Of course, the fact that these examples are all unusual in some way suggests a tendency towards normalization, a phenomenon we intend out theory of learning to explain.)

Finally, we note that a word that has a grammatically-derived lex-con may also have lex-cons that are not grammatically derived. For example, every verbal lex-con seems to have a gerund which describes the same situation as does the verb. Thus, the verb ''kill'' has a gerund ''killing'' which also denotes a killing event, as in ''The killing of the nuns was indefensible''; the verb ''fill'' has the gerund ''filling'' which

*This idea is adapted from Paul Kay, personal communication.

describes a filling event, as in ‘‘The filling of the huge pit took all day’’. However, both words have lex-cons that are not determined grammatically. Thus, ‘‘John made a killing in coffee futures’’ and ‘‘These doughnuts have a delicious filling’’ involve different lex-cons of these terms whose relation to their other lex-cons is of the sort we find more interesting.

(2) Core-related Lex-cons

One way in which the relation between two lex-cons may be more interesting is when their meanings intersect in ways that are not completely predictable. We call this shared meaning a *core meaning* (or *core concept*) of the senses, and say that these senses (and thereby, the lex-cons which provide them) are *core-related*. Such relations are assumed to hold between basic lex-cons of the same word, and between those of different words. When we have occasion to refer to all the senses that extend a core meaning, we will use the term *meaning complex*.

As an example, there are lex-cons of ‘‘have’’, ‘‘give’’, ‘‘take’’, ‘‘get’’ and ‘‘receive’’ whose senses are core-related. The core meaning is the notion of possession, which happens to be a sense of the word ‘‘have’’. The senses of (the lex-cons of) the other terms extend this core by specifying a change in possession, along with different relations between their roles; in addition, these lex-cons have different valences. For example, ‘‘give’’ has at least two lex-cons relevant to this possession meaning complex; let us temporarily call the sense specified by both of these lex-cons Give. The Give concept has a Giver, a Givee, and a Given, say; the meaning of these terms is established in the conceptual domain, i.e., by representing facts predicating their relation to other concepts in the system, e.g., that each Give is an Action by the Giver resulting in the Givee possessing the Given. One of the two lex-cons specifying this sense would have the valence information that the Given is specified by the direct object and the Givee by a prepositional complement beginning with ‘‘to’’, the other that the Givee and the Given are specified by the two consecutive noun phrases (e.g., a ditransitive), in that order.

Similarly, ‘‘get’’ may also have two lex-cons, in this case with different but related senses: One sense specifies that the recipient is also the active agent of the transfer, the other that the recipient plays a passive role. Indeed, the first of these senses may be identical to a sense of ‘‘take’’, and the second identical to the sense of ‘‘give’’ just mentioned. However, the lex-cons are quite different. For example, the lex-con of ‘‘get’’ that would have the same sense as one of ‘‘give’’ would specify that the syntactic subject is the Givee, rather than the Giver, which can be specified by an optional ‘‘from’’ complement.

Of course, if a single concept is shared by a lex-con of ‘‘give’’ and a lex-con of ‘‘get’’, and another shared by a lex-con of ‘‘give’’ and one of ‘‘take’’, we should name them more neutrally. For example, instead of Give, we should have named this concept something like Agent-causes-recipient-to-possess; instead of Take, we should have Agent-causes-self-to-possess. We should similarly rename the various thematic roles. In any case, the semantics of these terms is established on the conceptual level, e.g., by stating in one’s knowledge representation language that an Agent-causes-other-to-possess event is a kind of Cause-to-possess event, etc.

We illustrate a few of these lex-cons below. For example, here are the lex-cons for ‘‘give’’ and ‘‘get’’ that share the same sense, and the ones for ‘‘get’’ and ‘‘take’’ that share the same sense:

Lex:	root-give
Sense:	Agent-cause-recipient-to-possess
POS:	V
Val:	[Subj: Agent, DO1: Recipient, DO2: Patient]

Lex:	root-give
Sense:	Agent-cause-recipient-to-possess
POS:	V
Val:	[Subj: Agent, DO: Patient, P[to]: Recipient]

Lex:	root-get
Sense:	Agent-cause-recipient-to-possess
POS:	V
Val:	[Subj: Recipient, DO: Patient, (P[from]: Agent)]

Lex:	root-take
Sense:	Agent-cause-self-to-possess
POS:	V
Val:	[Subj: Agent, DO: Patient, (P[from]: Donor)]

Lex:	root-get
Sense:	Agent-cause-self-to-possess
POS:	V
Val:	[Subj: Agent, DO: Patient, (P[from]: Donor)]

These are basic lex-cons, so the lexical entries contain roots rather than word forms per se. The valence structure is given as a list of syntactic functions and associated thematic roles, with parentheses indicating optionality, and ‘‘P[*word*]’’ indicating a prepositional phrase headed by *word*.* On the conceptual level, we would have facts representing the relation of the concepts in these lex-cons to other concepts. For example, assuming that we have a predicate ‘‘AIO’’, representing the relation between an individual and a category, we might have the following:

```
V g E a,p,d,c,i,t
   AIO(g,Agent-cause-self-to-possess)
    –> Agent(a,g) & Patient(p,g) & Donor(d,g)
          & AIO(c,Causing)
             & Cause(i,c)
             & Effect(t,c)
          & AIO(i,Action)
             & Agent(a,i)
          & AIO(t,Transfer-of-possession)
             & Donor(d,t)
             & Recipient(a,t)
             & Patient(p,t)
```

That is, every instance of a Agent-cause-self-to-possess event, there is an Agent, a Patient and a Donor, and, moreover, the Agent does some action that causes the Transfer-of-possession of the Patient from the

*We are appealing rather casually in this discussion to some notion of thematic role, or conceptual case frame. Actually, the thematic roles are specified too generally in these examples. In the spirit of this analysis, we might have specialized thematic roles for each distinct concept, e.g., an Agent-of-Agent-cause-self-to-possess, rather than simply Agent. That every such role is also an Agent role would be captured on the conceptual level. We generalized to more common case roles primarily for expositional purposes, but shall return to related issues below. See Wilensky (1990) for a more detailed analysis of the nature and status of such objects.

Donor to that Agent. (This particular style of representation, in which there are objects corresponding to categories of events, is discussed further in Wilensky (1990), although nothing in the present discussion hinges of these details.)

''Meaning complex'' and ''core meaning'' are relative notions. In particular, a given sense might be a member of more than one meaning complex, each with its own core; the various meaning complexes pertaining to an individual meaning may intersect, or they may be disjoint. For example, in the case of the possession terms listed above, all the terms but ''have'' involve a notion of change of possession. Thus we might define ''change of possession '' to be a core meaning of this more restricted semantic field. This core happens to contains the core of the larger class. On the other hand, we might also define a class whose core is simply the concept ''change''; such a core would not contain the core of the original complex.

We have no a priori reason to prefer one core to another. However, we believe that some semantic intersections will be empirically justifiable in terms of how a language and conceptual system is structured, and will be reflected in speaker intuitions and linguistic data. For example, the possession uses of ''take'' and ''give'' will probably be judged by most speakers as being strongly related, while uses of ''give'' and ''destroy'' would not be. On the other hand, ''steal'' and ''murder'' might reflect a core relation at least as compelling as that underlying ''give'' and ''steal''. Indeed, one might consider core-relatedness to be a matter of degree, strong for those cores which, like possession, select a small percentage of concepts, and weak for those cores, like change, which select a relatively large proportion of concepts. It is generally only very strongly core-related concepts that are of interest to us here.

As another example, consider ''buy'' and ''sell'', and related terms. According to Fillmore (1978), the relatedness of these terms can be captured by postulating a single ''commercial transaction frame'', plus different valence descriptions for each word. In contrast, we would prefer to stipulate that this commercial transaction frame is just another concept, which happens to serve as the core for the relevant senses of ''buy'' and ''sell''. I.e., a Buy concept and a Sell concept are postulated, which relate to the Commercial-transaction concept by information expressing facts such as that the agent of a buying event is the person with the money in the commercial transaction scenario. The lex-cons for ''buy'' and ''sell'' just specify the Buy and Sell concepts, respectively, and add valence information, e.g., that, for the word ''buy'', the Buyer is realized as the syntactic subject.

We do not view our proposal as opposed to Fillmore's as much as an extension of it. Terms that Fillmore views as frame-related are certainly core-related in our terminology. However, we can also easily talk about the relation of terms in which a frame is difficult to postulate (for example, the ''have'', ''give'', etc., example). Also, we can make some finer distinctions. For example, the concept ''paying'' is certainly related to ''buying'' and ''selling'', but it is not obviously frame-related: One can (and, unfortunately, needs to) pay for items other than goods and services delivered; taxes and parking fines are two painful examples. But paying and buying are still related, presumably because the latter involves the former, even if the former does not appeal to the commerical transaction scenario per se. This observation does not rule out the utility of the commercial transaction frame, as much as diminish its status to just another concept serving as the basis of core-relatedness.

Two issues raised by this discussion are just how many lex-cons there are for any particular lexical item, and just how many senses these lex-cons distinguish. For example, consider the uses of ''tie'' in the following sentences:

(1) Jan tied a knot in the rope.
(2) Jan tied the boat to the dock.
(3) Jan tied the rope around the tree.
(4) To remember your notes, tie a knot around your finger.
(5) Tie your shoelace.
(6) Jan tied up Lynn.
(7) Jan tied two ropes together.

(8) This new rope ties easily.

One's naive intuition may be to attribute a single, rather general lex-con to cover all of these usages, and assume that the differences between the usages can be computed via some general principles. But such general lex-cons are evasive. First of all, there are semantic differences between the usages that do not appear derivable via general principles from a single meaning: If there were a general meaning, it could not include ''connect'' or ''secure'' as part of it, since that would not apply to ''tie a knot''. Instead, this aspect would have to be inferred from the other constituents, plus, say, the world knowledge that one typically ties knots to secure or connect things. The problem with this line of reasoning is that it seems not to apply to other cases in which it should. For example, the verb ''knot'' also means to make a knot, but one cannot say ''Lynn knotted Jan to the chair'', meaning, to secure her to a chair by making a knot in a rope; ''?Lynn knotted the rope around Jan'' is questionable. ''bind'' is similar to ''tie'', but while one can ''bind Jan to the chair with a rope'', one cannot ''?bind the rope around Jan.''

This is just a variation of the familiar argument that meaning, at least in its grosser forms, seems not to completely determine grammatical properties. Thus we have ''Jan gave the United Way some money'', but not ''*Jan contributed the United Way some money''; ''Jan is likely to leave'' but not ''*Jan is probable to leave''; ''Jan made Lynn leave'' but ''Jan forced Lynn to leave''. We are merely stressing here that similar arguments apply to valence alternations of the same term. While there have been many attempts to capture some of these relations (notably, lexical redundancy rules, grammatical transformations, and reference functions), there always seem to remain exceptions that must simply be stipulated. (Or, rather, as we will discuss below, if we do attribute these differences to differences in meaning, they are subtle differences that are hard to justify on independent grounds, so that the differences in meaning must be stipulated just to account for the grammatical variation.) Since neither syntactic, semantic nor pragmatic principles seem to predict these facts, we will assume they each must be specified by separate lex-cons. Obviously, such lex-cons will share significant components. The primary focus of this paper will be to capture and exploit such regularities.

Even assuming that a distinct lex-con is required for each of the sentence (1)-(8), we still are left with the question of whether each of the proposed lex-cons specifies a distinct sense, or whether several of them share a common sense and differ only in valence. Whatever the senses are, they appear to involve the core concept of making a knot. Indeed, this core seems to be the sense used in (1) and (8). Arguably, the other sentences involve the notion of securing something by making a knot in a rope-like object. Thus, one way to accommodate these examples is to propose two senses: One corresponding to the use of ''tie'' in the first and last sentences, the other to all the rest. The first sense, say Tie-knot, has a Tier, a Knot-tied, and a Rope-obj. The second, say Secure-by-Tie-knot, has the same roles plus a Tie-object and a Tie-to-object.

Let us example the second sense, which is obstensibly used by lex-cons in sentences (2)-(7). For example, in the lex-con corresponding to (2), the Tie-object is the syntactic direct object, and the Tie-to-object is expressed by a ''to''-complement (an optional ''from''-complement could be used to express the Rope-obj); in (3) the direct object expresses the Rope-obj, which, perhaps, also happens to be the Tie-object, and the prepositional complement expresses the Tie-to-object (i.e., Jan secured the rope to the tree by tying a knot in the rope); in (4) the Knot-tied in the direct object, and the Tied-to-object is expressed by the prepositional complement. In (5), we might interpret the direct object as the Rope-obj, the Tie-obj, and the Tie-to-obj, i.e., the shoelace is secured to itself by tying a knot in it. Likewise, in (6), direct object is both the Tied-obj and the Tied-to-obj. In (7), the two ropes each play the role of the Rope-obj and the Tie-obj and the Tie-to-obj.

In the approach just suggested, we assume that sentences (2)-(7) all share some basic scenario; any additional facts about each scenario must be inferred by some general principles. For example, if instead of (3) we had

(9) Jan tied the rope around Lynn.

it would be an inference that Lynn (in addition to the rope) was secured by this action. Similarly, in (6) and (7), anything special about securing something to itself would have to be inferred.

An alternative to this analysis is to recognize each of these lex-cons as referring to a slightly different situation. For example, one situation might involve securing one object to another; another situation involves securing an object to itself. In this case, we end up with different senses for each lex-con. One advantage of this formulation is that we are left with no delicate inference problem: Each sense just contains whatever information is appropriate. Also, the senses are all core-related, so we have not lost any obvious generalization. Indeed, we might still posit a rather more abstract concept in which two (possibly identical) objects are bound to each other, and then differentiate this concept by more specialized ones that serve as word senses. For example, in one sense, one of the two objects that are coupled together is profiled, i.e., viewed as the object being secured, while the other is viewed as background. Another sense has the two objects being identified. This alterative is quite similar to the one we have already suggested in the ‘‘buy-sell’’ example above, in which the senses of ‘‘buy’’ and ‘‘sell’’ differ from one another primarily in how they schematize the common Commercial-transaction concept.

Postulating subtle semantic differences between senses seems to offer other potential benefits. For example, sentences like

(10) Jan gave Lynn a book.
(11) Jan gave a book to Lynn.

are often seen as paraphrases, and therefore, would appeal to a single sense as we assumed above. But postulating that there are two closely related senses, one of which schematizes a participant as a goal, and another which schematizes that participant as a recipient, offers an interesting advantage. Specifically, we might further suppose that various thematic roles are expressed systemically: In particular, agents, recipients and patients can be expressed non-obliquely, and in order of subject, first direct object, second direct object; other roles are expressed obliquely, using various standard prepositional complements: goals by a prepositional ‘‘to’’ complement, and donors by a ‘‘from’’ complement. Then we could postulate the following lex-cons for ‘give’, in contrast to the corresponding ones proposed previously:

Lex:	root-give
Sense:	Agent-cause-recipient-to-possess
POS:	V
Val:	[Agent, Recipient, Patient]

Lex:	root-give
Sense:	Agent-cause-goal-to-possess
POS:	V
Val:	[Agent, Patient, Goal]

The senses that appear in these lex-cons are core-related, i.e., we would stipulate facts relating both Agent-cause-recipient-to-possess and Agent-cause-goal-to-possess to some concept that neutralizes the distinction they create, e.g., Agent-cause-someone-to-possess. Since the various thematic roles are all realized according to general principles, we need not specify any grammatical information in the valence description: The first lex-con will correspond to sentences of the form ‘‘Agent gives Recipient Patient’’, and the second to sentences of the form ‘‘Agent gives Patient to Goal’’.

In addition to accounting for sentences like (10) and (11), this approach will also account for the fact that, in the passive lex-cons for such words, the recipient may appear as subject but not the goal. Namely, the passive lex-cons have the agent role demoted to be realized by an optional oblique ‘‘by’’ complement, thus allowing other roles to be expressed in subject position. Corresponding to the first lex-con, in American

English, the Recipient may be expressed as subject, but in British English, either the Recipient or the Patient; for the second lex-con, only the Patient could appear as subject, since Goals are presumed to always be expressed obliquely. For verbs like ‘‘contribute’’, we would have just a single sense that provides a goal rather than a patient. The lex-cons referring to this sense would sanction sentences like ‘‘I contributed money to the church’’ and ‘‘Money was contributed to the church’’ since goals are ordinary expressed obliquely, but not ‘‘*I contributed the church money’’ or ‘‘*The church was contributed money’’, since only recipients (along with other ‘‘higher’’ thematic roles like patient and agent) can be expressed non-obliquely.

Our conclusion, then, is to err on the side of proposing too many word senses. In the worse case, this ‘‘sense per lex-con’’ position can accomodate all the data with some redundancy; but having lots of close word senses does not miss any important generalization as long as the core-relatedness is captured on the conceptual level.

We emphasize that we are not making a deep philosophical point here, although perhaps we are refuting one. Superficially, our argument is reminiscent of Wittgenstein’s (1953) argument that there are concepts all of whose members share no common properties. One of his well-known examples is the notion of ‘‘game’’. Wittgenstein claimed that the set of all games shares no common properties, as some games involve competition and some do not, some involve skill while others do not, etc. In the worse case, we regard this argument as simply false: Wittgenstein was merely debunking bad definitions of games, not showing that there were no good ones. For example, the definition ‘‘an activity involving the pursuit of a goal following a set of gratuitous rules’’ seems adequate for cases Wittgenstein felt problematic. However, in the best case, should we acknowledge some related activity not covered by this definition (e.g., ‘‘war games’’) we can simply deny the existence of a single game concept, and instead propose additional senses related to each other by some lexical relation.

(3) Transforming Lexical relations

We now come to those basic lex-cons of a word that can be related to other basic lex-cons through some non-core relation. We give a few brief examples here to illustrate the idea. For example, the basic lex-cons of ‘‘have’’, ‘‘give’’, and ‘‘get’’ mentioned above are each metaphorically extended to the domain of being infected with a disease. Thus, ‘‘have a cold’’ means being infected with a cold; ‘‘give Lynn a cold’’ means to infect Lynn with a cold, etc. (This example of metaphoric extension is discussed in Martin (1988).) In our terminology, for each of the initial set of lex-cons, there is another lex-con which is metaphorically based on the first. Similarly, the lex-con of ‘‘tie’’ whose sense means to associate or abstractly connect (as in ‘‘What ties these facts together’’) is presumably a metaphoral extension of one of the senses of ‘‘tie’’ discussed above.

We will call such relations *transforming* lexical relations. The intuition is that transformating lexical relations extend one sense to another that does not contain the original. We say that one of the senses is *based* on the other, and that such a sense is an *extended sense* of another. We call a meaning of a word that is not based on another a *central sense*. A transforming relation is a synchronically logical relation, that is, it pertains when a speaker cognizes one sense as being based on another in some fashion. It does not imply that one sense is derived from another during ordinary language use, although our theory of learning presumes that central senses would ordinarily be learned prior to extended ones. In cases in which a synchronic priority of one word sense relative to another cannot be established, we will consider both senses central.

Central senses appear to be the ones that are most easily recalled. (Of course, recall is presumably conditioned by many factors, but central senses seem so easy to retrieve that our naive intuition seems to equate them with the meaning of the word.) Thus, the possession meaning of ‘‘give’’ and the knot meaning of ‘‘tie’’ are central senses, while the infect lex-con of ‘‘give’’ and, say, the join-in-marriage lex-con of ‘‘tie’’ appeal to extensions of these senses via metaphoric relations.

While metaphorical relations are quite common, and perhaps worthy of special attention, many other kinds of relations can be discerned. For example, we allege that the verb ‘‘tie’’ has a central sense of ‘‘make a knot’’ (as in ‘‘tie a knot’’). This sense is core-related to the meaning ‘‘secure/connect by making a knot’’ (as in ‘‘tie Jan to the chair’’). However, it is a transforming relation that extends it to the sense of physically connect (as in ‘‘tie the beam to the joist’’), since the latter does not involve tying in the first sense. (These senses are related to nominal senses in an interesting fashion. The metaphorical sense of the noun ‘‘tie’’, meaning the abstract connection (as in ‘‘family ties’’), seems to be most readily available; the sense of physical connection is rather restricted (as in ‘‘railroad tie’’); and there appears to be no general nominal senses related to the central sense of the verb (i.e., a nominal lex-con ‘‘tie’’ meaning the rope used to connect two things).)

As a less obvious example, consider first the central sense of ‘‘gate’’, which we presume to be a movable structure controlling entrance through a barrier. While there are certainly metaphorical extensions of this concept (e.g., in electronics), consider the sense involved in a sporting arena, airport, and some university campuses, in which the gate refers to a passageway, and not to a movable structure controlling access. This relation appears to be some kind of ‘‘frame-complementation’’, in which the term used for central component is extended to a nearby component. Since this relation seems somewhat curious, we make the point of arguing that it is not an isolated case. Consider for example, the use of the term ‘‘hole’’ in the sense of ‘‘doughnut hole’’, which refers centrally to the empty space, and via frame-complementation to the dough ball formed creating this empty space; ‘‘window’’ can refer to the space in the wall as well as the object that fills it; ‘‘bed’’ can refer to the frame or the mattress and box spring that it supports.

We note that the notions like ‘‘central sense’’ are not universally accepted. Most notably, Nunberg (1978) suggests that such notions are ill-conceived. Instead, he proposes that uses of a term can be related to one another through a relatively small number of referring functions without necessarily having to decide which use is more primary. For example, ‘‘type of’’, ‘‘cause of’’, ‘‘possessor of’’, ‘‘location of’’, ‘‘contents of’’ and ‘‘product of’’ are examples of referring functions. These are aimed at explaining at least the metonymic use of natural language terms. Thus, the word ‘‘Picasso’’ used to refer to a painting is explained by appealing to a ‘‘product of’’ function between the person and his work. Nunberg’s uses do not correspond to word senses per se. Instead, he suggests a scale of conventionalization upon which given uses might be found.

We feel that the flexibility of pragmatic interpretation does not necessarily undercut the notion of central senses and lexical relations, however. In particular, it seems clear that some usages are quite general and depend neither on the individual lexical items nor on conventionalization. It would simply be inappropriate to compile these into lexical entries, despite our inclination to err in this direction. For examples, it seems that all concrete object count nouns can be interpreted in at least three ways: as the entire class, as a type of the class, or as an individual of that class:

(12a) Lynn bought the convertible she test-drove,
(12b) The convertible Ford just came out with is quite an improvement.
(12c) The convertible changed the way people drive.

However, the availability of these interpretations has nothing to do with the lexical item per se. For example, the same readings are possible in all cases when ‘‘convertible’’ is replaced with ‘‘car with a soft top’’. So these possibilities of interpretation appears to be a general fact about the interpretation of noun phrases, rather than facts about the lexicon. Of course, this is not necessarily true of all of Nunberg’s referring relations, some of which re-emerge as kinds of lexical relations below. Also, some particular usages do appear to get lexicalized, and others do not. For example, Nunberg points out that ‘‘window’’ can be used to refer variously to the glass, the entire structure, and the enclosing space, and appears to have this property in all languages. At the same time, ‘‘gate’’ does not ordinarily have this function.

A Test for Centrality

Speakers seem to judge central senses as *literal* or *actual* uses of a word (although, as we have argued elsewhere (Wilensky 1989), literal should not be confused with lexicalized). This provides us with a convenient, but limited, tool for judging whether a sense is central or not. For example, consider the following:

(13a) You didn't literally cut your class.
(13b) ?You didn't literally cut your hand.

(14a) That's not really a doughnut hole; it's the part they removed to make the hole.
(14b) ?That's not really a doughnut hole; it's the space created from removing the hole.

(15a) You didn't buy an actual bed; you just bought a frame.
(15b) ?You didn't buy an actual bed; you just bought a frame, box spring and mattress.

While it is useful, this test is limited in several ways. First, it seems to apply to only polysemous senses of the same parts of speech; second, it seems difficult to apply to some parts of speech, for example, adverbs and connectives; third, it is unclear what the test says for the most closely related hypothesized senses, which of course are the most problematic. For example, it is unclear that there are strong judgments about cases like the following

(16) You didn't literally take Lynn to dinner.

An Example

We illustrate in some more depth the kinds of lexical relations of interest to us by considering some of the lex-cons of a single word, namely, the word ''open.'' We emphasize that we are by no means claiming to explicate all the lex-cons of this word; in particular, we ignore the important cases of verb-particle construction. However, examining a number of these lex-cons helps demonstrate the nature and scope of lexical relations.

The central sense of this word involves a core concept that is, roughly, physical access to a region. This core is most directly manifest in the adjectival form, e.g., ''the open door''. It is also the basis of a meaning complex that includes related senses of the same term. For example, the core meaning is extended to the concept ''change to enable physical access'' for the sense used in the intransitive, as in ''The door opened''; it is extended further to a sense, used in the transitive, specifying that access is created by acting upon an object, as in ''Jan opened the door''. However, as before, there are several lex-cons involved in each case, differentiated from one another by valence. For example, consider only the transitive verb ''open'' whose use involves the central sense. One lex-con has the object moved as the patient, and hence as the direct object of the transitive verb (as in ''open the door'' or ''open the lid [to a jar]''); another uses the container itself as the direct object (e.g., ''open the jar''); perhaps another involves some sort of aperture that widens (e.g., ''open your mouth'', and perhaps, ''open the throttle''). Each of these lex-cons corresponds to a different intransitive lex-con (e.g., ''the door opened'', ''Bill's mouth opened'', and a different adjectival lex-con (''open door'', ''open jar'').

Other lex-cons of ''open'' appeal to meanings that appear to be derived from this central sense in various ways. One such meaning is causing an information-containing item to come into existence (e.g., ''open a bank account'' or ''open a file on someone''). This meaning seems to stand in a metaphorical relation to the central sense. Also, it seems to have adjectival lex-cons but no corresponding intransitive ones (i.e., ''an open bank account'' is fine, but ''?The bank account opened.'' is awkward at best).

Another metaphorical extension is the computer terminology ''open a file''. Note that, unlike the previous

meaning, this one does not necessarily imply creating the object (i.e., one can open an existing file for reading, but it would make no sense to open an existing bank account). Rather, here the physical access dimension of the meaning is extended to the more abstract notion of enabling access of information.

A similar contrast between creation and enabling access seems to exist in an extension of the central sense along a social dimension. Thus, in ‘‘McDonald’s opened 100 new stores this year’’, each store came into existence, whereas in ‘‘They open the store at 9:00’’, social access is enabled. Another related meaning involves beginning a scripted activity, as in ‘‘open the bidding’’ and ‘‘open the meeting’’. These meanings seem to appear in intransitive verbs and adjectival lex-cons as well. A perhaps related extension is in the direction of permission, as in ‘‘They opened the meeting to all interested participants’’, or, ‘‘They opened the [previously closed] road’’. This meaning occurs in the adjectival but not the intransitive.

Interestingly, while one can reopen a legal case that has been closed, one does not open the case to begin with.

While there are other related metaphorical meanings, there are related meanings of quite different sorts. For example, there are other physical meanings, as in ‘‘open the lock’’ (which seems to facilitate core opening of whatever is locked). Arguably, the unfolding use of ‘‘open’’, as in ‘‘open the umbrella’’, ‘‘the bird opened its wings’’ (but not ‘‘*The bird opened’’ with the same meaning), is a separate but related physical sense, as is that underlying ‘‘open the venetian blinds’’.

We will count as separate senses what are probably metonymic uses, such as ‘‘open the wine’’, ‘‘open your eyes’’ and ‘‘The stock opened at 34’’. The former sense seem awkward at best in the intransitive; the latter has only this form; the adjectives corresponding to the first usage seems marginally acceptable, but ‘‘?open stock’’ or ‘‘?That stock is open’’ does not.

There are, of course, many adjectival uses that do not obviously correspond to any verbal lex-con, such as ‘‘open country’’, ‘‘open book’’ (in the metaphoric sense), ‘‘open set’’ (in mathematics), ‘‘open question’’, ‘‘That position is still open’’, ‘‘open letter’’, ‘‘open manner’’, ‘‘open vowel’’, etc.

The nominal ‘‘open’’ apparently has only the quite restricted lex-con of a kind of sports tournament, while ‘‘the open’’ refers to either open space or the metaphorically related notion of public knowledge. More of the meanings discussed above are related to lex-cons of the nominal ‘‘opening’’, which, in addition to the standard grammatically-derived meaning of an event of opening corresponding to any verbal lex-con, can also denote an aperture, an open space, the beginning of an event that has been conventionally sectioned (i.e., the opening of a play or chess game), an opportunity, and an unfilled position.

Finally, other words, e.g., ‘‘close’’, have senses that we presume to be core-related to the various lex-cons of ‘‘open’’ just discussed. Indeed, every verbal lex-con of ‘‘open’’ has a corresponding lex-con of ‘‘close’’, with the following exceptions: The metonymic extension of central ‘‘open’’ (e.g., ‘‘open the wine’’) does not seem to have a corresponding lex-con of ‘‘close’’ (i.e., ‘‘?close the wine’’); also, ‘‘?close the lid’’ seems unacceptable. The adjective corresponding to ‘‘open’’ is ‘‘closed’’ (rather than ‘‘close’’), and seems to correspond to its ‘‘open’’ counterpart less systematically, so we do not have ‘‘*closed country’’ or ‘‘*closed position’’. There do not appear to be any lex-cons of ‘‘closing’’ other than the grammatically-derived meaning (e.g., there is no lex-con of ‘‘closing’’ corresponding to a closed space, the end of a chess game, the loss of opportunity, or the like.)

We will not defend the details of this particular analysis here. But we interpret the data as suggesting the following: Words have many lex-cons whose details must simply be learned by the language user. However, while these facts may not be completely predictable from other facts, there is much about them that is not arbitrary either. Different lex-cons share the same meaning, or extend the same core meaning, or involve the transformation of a meaning, in reasonably regular ways. Thus, if the central sense of open is extended to opening a store, it should be no surprise that closing is extended to this domain in a consistent way. While we probably have no way of knowing, for example, from all the other facts we might possess,

that chess games have an opening but not a closing, given that they have an opening, we probably could predict that this is the beginning of the game and not, say, the end. To use the terminology advocated by Fillmore (1978), the relations between word lex-cons are *motivated*, that is, there may be some relatively transparent reason for them, even if that reason cannot necessarily be used in advance to predict the phenomenon.

Also, our proposal seems consonant with that of Vennemann (1974), who suggested that rules are generalizations about the lexicon, and might help to organize or structure it, rather than to simplify it. In addition, we interpret as evidence supportive of a position like ours psycholinguistic results such as those of Cutler (1983) and Butterworth (1983), which suggest that words are represented in their full ‘‘undecomposed’’ form, along with some sorts of relations between related words.

2.1. Kinds of Lexical Subregularities

We view the various lexical relations just discussed as one kind of linguistic subregularity. By subregularity, we mean any phenomenon that is systematic but not predictable. To qualify as a useful lexical relation, the relation must hold between a number of different examples; at the same time, most of these relations do not predictably hold wherever they might. For example, the systematic metaphoric structuring alluded to above, and studied extensively by Lakoff and Johnson (1980), comprise one important class of subregularities. Thus, it is worth recognizing the metaphoric lexical relation between central sense ‘‘open’’ and the ‘‘open’’ of ‘‘The play opens Thursday’’ because it appears in other places, such as ‘‘He opened with P-K4’’ and ‘‘The play closes Thursday’’. But the use of the relation is not predictable: ‘‘?He closed with QxB’’ is awkward at best.

The question arises as to just what kinds of subregularities there might be. One set of candidates is approached in the work of Brugman (1981, 1984) and Norvig and Lakoff (1987). In particular, Norvig and Lakoff (1987) offer six types of links between polysemous word senses in what they call *lexical network theory*. Links in a lexical network map senses to one another, starting from some central sense and encoding some minimal variation. The six links identified are as follows, with the accompanying sentences illustrating the link:

(1) Image-schema transformation links.

Jan ran over the hill.
Jan lives over the hill.

(2) Metaphoric links.

Jan took the book to Lynn.
Jan took a punch at Lynn.

(3) Metonymyic links.

Budweiser produces a lot of beer.
Give me another Budweiser.

(4) Frame-addition links.

Jan took one suitcase to Bali.
Jan took Lynn to the movies.

(5) Semantic role differentiation links.

The baby took the toy from its mother.
The baby took the book to its mother.

(6) Profile shift links.

Jan took the book to Lynn.
Jan took one suitcase to Bali.

We note that while these authors do not explicitly offer these relations as subregularities, this appears to be their intent. That is, these links are presumed to occur in other places, but not predictably so. Also, their theory is concerned only with lex-cons of one word, while inter-lexical generations are at least as importance to us* (see Norvig 1989)). Finally, we feel that there are many, many more links than these, or at least, many subcases which are necessary to distinguish. Indeed, we have no reason to believe that the number of such subregularities is bounded in principle, and thus do not offer a theory that pithily summarizes them.

Instead, we present a partial list of the subregularities we have encountered that appear to have some predictive value. Note that most of these are either equivalent to transforming lexical relations, or core-relations between lex-cons of the same lex-con. This is because, to have predictive value, a relation must specify phonological, syntactic and semantic components. In particular, most core relations between lex-cons of different words are not predictive even as subregularities. For example, while ‘‘open’’ and ‘‘close’’ are obviously core-related, there is no way to use just this semantic and syntactic relation (say, verbs with the same valence denoting opposite actions) to predict the meaning of one from the other (although, of course, a system trying to infer the meaning of a word from context might be aided by the notion that the opposite of a term in a meaning complex is likely to be lexically realized). On the other hand, knowing that ‘‘open’’ is metaphorically extended in a give manner is useful for hypothesizing that a core-related term like ‘‘close’’ will be extended analogously. Similarly, ‘‘open the lid’’ can help to hypothesize a meaning for ‘‘open the jar’’. Thus transforming lexical relations such as metaphoric relations, and core-related lex-cons of the same term correspond to useful subregularities, while core-relatedness in general does not.

The list below uses a rather informal rule format, and gives a couple of examples of words to which the rule is applicable. It is hoped that explicating a few examples below will let the reader infer the meanings of the others:

2.1.1. Subregularities as Rules

(1) function-object-noun –> primary-activity-‘‘determinerless’’-noun
(‘‘the bed’’ –> ‘‘in bed, go to bed’’; ‘‘a school –> at school’’; ‘‘my lunch –> at lunch’’; ‘‘the conference –> in conference’’)

(2) noun –> resembling-in-appearance-noun
(‘‘tree’’ –> ‘‘(rose) tree’’; ‘‘tree’’ –> ‘‘(shoe) tree’’); ‘‘tiger’’ –> ‘‘(stuffed) tiger’’, ‘‘pencil’’ –> ‘‘pencil (of light)’’)

(3) noun –> having-the-same-function-noun
(‘‘bed’’ –> ‘‘bed (of leaves)’’)

*Interlexical generalizations are certainly not antithetical to these researchers; they have simply not provided a specific account of it as they have of the intra-lexical.

(4) noun –> concreted-noun
(‘‘opening’’ (meaning opportunity) –> ‘‘opening’’ (meaning job opportunity); ‘‘professor’’ (meaning professor of any rank) –> ‘‘professor’’ (meaning full professor); ‘‘dog’’ (meaning any dog) –> ‘‘dog’’ (meaning male dog))

(5) noun –> generalized-noun
(‘‘man’’ (meaning male human) –> ‘‘man’’ (meaning human); ‘‘cow’’ (meaning female bovine) –> ‘‘cow’’ (meaning any bovine))

(6) noun –> ‘‘involve-concretion’’-verb
(‘‘a tree’’ –> ‘‘to tree (a cat)’’; ‘‘a knife’’ –> ‘‘to knife (someone)’’)

(7) verb –> verb-w-role-splitting
(‘‘take a book’’ –> ‘‘take a book to Mary’’, ‘‘John shaved’’ –> ‘‘The barber shaved Bill’’)

(8) verb –> verb-w-role-joining
(‘‘tie the rope to the tree’’ –> ‘‘tie a shoelace’’; ‘‘glue the handle to the cup’’ –> ‘‘glue the broken vase’’)

(9) verb-w-differently-profiled-roles –> verb-w-equally-profiled-roles-coordinated
(‘‘John married Mary’’ –> ‘‘John and Mary married’’; ‘‘The Honda hit the Jaguar’’ –> ‘‘The Honda and the Jaguar hit [each other]’’; ‘‘tie the boat to the dock’’; ‘‘tie the two pieces together’’; ‘‘I can’t tell one from the other’’ –> ‘‘I can’t tell them apart’’)

(10) verb –> alternative-profiling-verb
(‘‘give Lynn a book’’ –> ‘‘give a book to Lynn’’; ‘‘tie the boat to the dock’’ –> ‘‘tie a rope around the tree’’)

(11) verb –> frame-imposition-verb
(‘‘take a book’’ –> ‘‘take someone to dinner’’, ‘‘go’’ –> ‘‘go dancing’’)

(12) cause-state-verb-i –> do-cause-state-verb-t
(‘‘The door opened’’ –> ‘‘Jan opened the door’’, ‘‘The window broke’’ –> ‘‘Jan broke the window’’)

(13) activity-verb-t –> concretion-activity-verb-i
(‘‘eat an apple’’ –> ‘‘eat [a meal]’’, ‘‘drink a coke’’ –> ‘‘drink [alcohol]’’, ‘‘We fed our guests’’ –> ‘‘the dog feeds’’)

(14) activity-verb-t –> dobj-subj-middle-voice-verb-i
(‘‘drive a car’’ –> ‘‘the car drives well’’)

(15) activity-verb-i –> activity-verb-t-primary-category
(‘‘John dreamed’’ –> ‘‘John dreamed a dream’’; ‘‘John slept’’ –> ‘‘John slept the sleep of the innocent’’)

(16) activity-verb-i –> do-cause-activity-verb-t(patient as subject)
(‘‘John slept’’ –> ‘‘The bed sleeps five’’)

(17) activity-verb –> activity-of-noun
(‘‘to cry’’ –> ‘‘a cry (in the wilderness)’’; ‘‘to punch’’ –> ‘‘a punch (in the mouth)’’)

(18) activity-verb –> product-of-activity-noun
(‘‘copy the paper’’ <–> ‘‘a copy of the paper’’; ‘‘film the event’’ –> ‘‘a film of the event’’; xerox,

telegram, telegraph)

(19) activity-verb –> object-used-for-activity-noun
(‘‘tie together’’ –> ‘‘railroad tie’’; ‘‘film the scene’’ –> ‘‘film for the camera’’)

(20) activity-verb –> used-for-‘‘gerundive’’-noun
(‘‘fill the doughnut’’ –> ‘‘filling for the doughnut’’)

(21) activity-verb –> result-of-‘‘gerundive’’-noun
(‘‘open [up] the wall –> ‘‘an opening in the wall’’)

(22) activity-verb –> experienced-as-result-of-‘‘gerundive’’-noun
(‘‘open [up] the wall’’ –> ‘‘an opening in the woods’’)

(23) activity-verb –> activity+purpose-verb
(‘‘tie a knot’’ –> ‘‘tie a Lynn to the chair’’; ‘‘take a book’’ –> ‘‘take a book to the Cape’’; ‘‘Jan screamed’’ –> ‘‘Jan screamed at Lynn’’; ‘‘Jan wandered’’ –> ‘‘Jan wandered over to the local pub’’)

(24) activity-verb –> purpose-of-activity-verb
(‘‘tie a knot’’ –> ‘‘tie the beam to the joist’’; ‘‘open the jar’’ –> ‘‘open the road’’; ‘‘open the jar’’ –> ‘‘open the meeting’’)

(25) activity-verb –> enable-purpose-of-activity-verb
(‘‘open the door’’ –> ‘‘open the lock’’)

(26) functional-noun –> use-function-verb
(‘‘the telephone’’ –> ‘‘telephone John’’; machine, motorcycle, telegraph)

(27) object-noun –> central-component-of-object
(‘‘a bed’’ –> ‘‘bought a bed’’ [=frame with no mattress]; ‘‘an apple’’ –> ‘‘eat an apple [=without the core]’’))

(28) central-object-noun –> frame-complement-noun
(‘‘gate’’ –> arena gate, campus gate, ‘‘hole’’ –> doughnut hole)

(29) domain reason –> speech act reason
(‘‘The party will be great BECAUSE they’re having a live band.’’ –> ‘‘I hope you’re not busy, BECAUSE there’s a great party tonight.’’ ‘‘IF you’re not doing anything, come to the party tonight.’’ –> ‘‘IF you’re not doing anything, there’s a party tonight I know about’’)

(30) word –> metaphoric mapping of word
(‘‘give a book’’ –> ‘‘give a cold’’, etc.)

Consider the first rule. This rule is intended to capture the idea that, for some noun whose central meaning is a functional object, there is another nominal lex-con of that word that occurs without determination, and means something like the primary activity associated with the central sense of the term. For example, the word ‘‘bed’’ has as a central sense a functional object used for sleeping. However, the word can also be used in utterances like the following: ‘‘go to bed’’, ‘‘before bed’’, and ‘‘in bed’’, (but not, say, ‘‘*during bed’’). In these cases, the noun is determinerless, and seems to mean something akin to being in bed for the purpose of sleeping for a significant period of time (i.e., to retire). Note, for example that it would be infelicitous to say I went to bed if I merely went over to a bed and sat down on it for a few minutes, or even if I took a short afternoon nap in one.

Other examples include ‘‘jail’’, ‘‘conference’’, ‘‘school’’ and virtually all the scheduled meal terms of English, e.g., ‘‘lunch’’, ‘‘tea’’, ‘‘dinner’’. For example, it would be misleading to say that I was ‘‘in jail’’ yesterday if I were visiting a relative, but acceptable to say that I was ‘‘in *the* jail’’ under such a circumstance. Note further than I can ‘‘send my children to school’’, but not ‘‘*to school down the street’’, while ‘‘to the school down the street’’ is acceptable. (Indeed, non-referential use of the noun presumably motivates its determinerless nature.) Also, which words conform to this subregularity is apparently a function of dialect. British English allows ‘‘in hospital’’ and ‘‘in university’’, while American English does not.

The dialect difference underscores the point that this is truly a subregularity: concepts that might be expressed this way are not necessarily expressed this way. Similar examples to the same point abound. For example, the more bound use of ‘‘bathroom’’ in the phrase ‘‘go to the bathroom’’ might as well have been determinerless; it just happens not to be.

We belabor this example not because it in itself is a particularly important generalization about English, but precisely because it is not. That is, there appear to be many such facts of limited scope, and each of them may be useful for learning analogous cases.

For another, rather different example, consider rule 6, which relates function nouns to verbs. Examples of this are ‘‘tree’’ as in ‘‘The dog treed the cat’’ and ‘‘knife’’ as in ‘‘The murderer knifed his victim’’. The applicable rule states that the verb means some specific activity involving the central sense (e.g., ‘‘knife’’ and ‘‘tree’’, respectively). I.e., the verbs are treated as a sort of conventionalized denominalization. Note that the activity is presumed to be specific, and that the way in which it must be ‘‘concreted’’ is assumed to be pragmatically determined. Thus, the rule can only tell us that ‘‘treeing’’ involves a tree, but only our world knowledge might suggest to us that it involves cornering; similarly, the rule can tell us that ‘‘knifing’’ involves the use of a knife, but cannot tell us that it means stabbing a person, and not say, just cutting.

As a final illustration, consider rules 7, 8 and 9, so-called ‘‘role splitting’’ (this is the same as Norvig and Lakoff’s semantic role differentiation link) ‘‘role joining’’, and ‘‘role grouping’’. Rule 7 suggests that a verb in which two thematic roles are realized by a single complement may have another lex-con in which these two complements are realized separately. For example, in ‘‘John took a book from Mary’’, John is both the recipient and the agent. However, in ‘‘John took a book to Mary’’, John is only the agent, and Mary is the recipient. Thus, the lex-con of ‘‘take’’ involved in the first sentence, which we suggest corresponds to a central sense, is the basis for the lex-con used in the second, in which the roles coinciding in the first are realized separately. A similar prediction might be made from an intransitive verb like ‘‘shave’’, in which agent and patient have a single realization, to the existence of a transitive verb ‘‘shave’’ in which the patient is realized separately as the direct object. Here the primacy of one sense over the other is unclear. (Of course, the tendency of patients to get realized as direct objects in English should also help motivate this fact, and should presumably also be exploitable during learning.) In rule 8, however, the intuition is that roles that are realized separately in the central sense are realized by a single object playing both roles: Thus, the use of ‘‘tie’’ in ‘‘tie one’s shoelace’’ is accounted for by postulating the existence of a lexical entry in which a single item fills two roles. In rule 9, two roles are similarly realized by a single complement, but in these cases, the two objects are notionally separate and may be found within a group specified by the single complement. Thus, ‘‘marry’’ is both a transitive verb, and an intransitive verb expressing the same roles with a subject denoting the pair of individuals.

Rules will often apply to lex-cons that are the product of other rules. Thus, a physical ‘‘opening’’ in a fence may be derived from the central lex-con of the verb ‘‘open’’; the lex-con of ‘‘opening’’ meaning opportunity is a metaphorical extension of the physical opening, and the lex-con of ‘‘opening’’ meaning job opportunity a concretion of the opportunity sense.

Note that while we have just listed these subregularities, both more general and more specific subregularities can be specified. For example, rules 4 and 5 have some special cases. A gender-specific category term

often becomes the term for the entire category; the name for an animal is often the name for the family containing that animal. Similarly, some of the rules for relating verbs to nouns (14-19) might be thought of as special instances of more general rules. In general, then, we recognize the existence of ‘‘sub-subregularities’’, or, a hierarchy of rules.

As the previous discussion suggests, we view the lexicon as containing conventionalized instances of what would otherwise be described as quite imaginative processes, e.g., metaphor, denominalization, and the like. If such a view is valid to some degree synchronically, then the lexicon is a very complex object indeed.

In sum, we have distinguished three kinds of relations between lexical entries: core relations, grammatical relations, and transforming relations. Core-relations capture the commonality of meanings within a meaning complex; grammatical relations hold between different forms of a word whose semantic relation is completely productive; transformation relations map meanings into new meanings. Meanings not produced by transformation relation are central senses. The amount of knowledge to be contained in the lexicon is large, but we can take solace in its systematicity, and hope to exploit this systematicity in acquisition.

3. Lexical Acquisition

The focus of the previous section has been the nature of lexical entries, and in particular, the nature of the subregularities among them. In this section, we propose how these subregularities can be gainfully exploited for language acquisition. We begin by reviewing some previous work on the acquisition of the lexicon, especially our own work on acquiring metaphorical word sense extensions. We then propose an extention of this work to the general problem of exploiting subregularities. The form this extention takes is that of a theory of analogical reasoning.

There have been numerous attempts to build systems that automatically acquire word meanings. One body of work attempts to acquire rather superficial accounts of lexical items by processing dictionary entries. However, work involving attempts at more substantial semantic accounts is more pertinent to our concerns. For example, Selfridge (1982) models acquisition of language, including lexical knowledge, in children. Granger (1977) attempts to hypothesize word meanings from the use of a word in content. Such systems attempt to learn basic word meanings assuming no prior knowledge about the new words they have encountered; rather they attempt to infer word meanings from the context of word use, in particular, from the semantic constraints imposed on words by the utterances in which they are found, and the world knowledge about the situations to which the words are being applied.

In our view, such work has produced valuable insights on the importance of world knowledge in the task of learning language. However, this approach by itself has not been notably successful to date. This is not very surprising, as the problem formulated in this way is a very difficult one. Instead, we focus on the problem of acquiring lex-cons that are related to ones already known to the system. This version of the problem does not allow one to ignore the importance of world knowledge altogether, but there are good reasons to think that it is somewhat more tractable.

One reason why prior lexical knowledge might make the acquisition problem more tractable is that, while the details of the meaning or usage of language forms may in general not be not predictable; but they are often not arbitrary either. If one has knowledge of a motivation underlying several linguistic forms, and instances of new uses, the motivation might be successfully generalized to apply to the new uses.

The approach to lexical knowledge acquisition we develop here grew out of our previous work on metaphoric knowledge (in particular, from Martin (1988), of which much of the discussion to follow is a summary). More specifically, as Lakoff and Johnson (1980) have argued, many linguistic subregularities seem to emerge if one assumes some underlying metaphoric structure. For example, the existence of the metaphor ‘‘actions done for the benefit of someone are objects transferred to that person’’ might be advanced to

sanction sentences like ‘‘John gave Mary a kiss’’ and ‘‘John got a kiss from Mary’’; however, the presence of such a metaphor would in itself not predict that ‘‘John did Mary a favor’’ is acceptable, but that ‘‘*John gave Mary a favor’’ and ‘‘*John did Mary a kiss’’ are not. It would also not explain many subtler phenomena. For example, if Mary asked John to kiss Lynn for her, and John did so, John would have done an action for the benefit of Mary, but it would be unacceptable to describe this action by the utterance ‘‘John gave Mary a kiss.’’

Even in technical domains, the use of systematic metaphors abounds (cf. Carbonell 1980). For example, in the UNIX domain, the following are typical utterances:

> How do I *enter* the editor?
> How can I *get out* of lisp?
> What command *exits* emacs?
> You can *leave* that program by typing ctrl-z.
>
> Typing ctrl-c will *kill* a process.
> My process *died*.
> The *parent* process *spawned* another process.

We describe the systematicity underlying such utterances as *conventional metaphors*, and attribute a significance to them primarily (but not exclusively) in motivating the distribution of lex-cons. For example, a conventional metaphor underlying each of the first set of examples above might be expressed as ‘‘an interactive process is a bounded region’’. In addition, each example may appeal to a more specific version of this metaphor. For example, the use of ‘‘enter’’ above might appeal to the metaphor of ‘‘initiating communicating with an interactive process is entering a bounded region’’. In other words, ‘‘enter’’ is presumed to have a distinct lex-con with the sense ‘‘initiate communicating with an interactive process,’’ and, moreover, this sense can described as being the result of applying the motivating metaphor to another (central) sense of the word.

Note that, following Lakoff and Johnson, we use the term ‘‘metaphor’’ to refer to some putative mental representation, and a conventionalized one at that, and not to some linguistic form that is novel. (There have been numerous complaints about this terminology, but we have not attempted to improve it.) However, our position is somewhat more conservative than Lakoff and Johnson’s. In particular, we do not attribute great significance to these metaphors beyond their role in motivating linguistic forms, and, as we will argue below, we do not believe that they are unique in this regard.

As our example above suggests, metaphors are presumed to be related to each other in a manner not unlike the way facts about the world are related to each other. In particular, some metaphors may be more specific instances of other metaphors; some may incorporate others. For example, the sentence ‘‘John gave Mary a cold.’’ is presumed to make recourse to a ‘‘a cold is a possession’’ metaphor. This metaphor might be a more specific verson of ‘‘an illness is a possession’’ metaphor. In addition, the sentence ‘‘John gave Mary a cold’’ is presumed to involve the ‘‘infecting with a cold is transferring possession of the cold’’ metaphor. This metaphor presumes the ‘‘a cold is a possession’’ one (in exactly the same way that giving presumes having). We call metaphors that are presumed by other metaphors, but which do not themselves presume other metaphors, *core* metaphors. A core metaphor is simply the metaphor mapping a core meaning of one meaning complex to that of another. Thus ‘‘a cold is a possession’’ is good candidate for a core metaphor, as possession is the core concept of a meaning complex which gets metaphorically extended; ‘‘an interactive process is a bounded region’’ and ‘‘a process is a (asexual) living thing’’ are core metaphors for other metaphorically extended meaning complexes.

Since all our metaphors do is ‘‘explain’’ the motivation for a word lex-con, we are not attributing any role to them in routine understanding or production. Instead, they are of interest to us primarily for language acquisition.

Previously, we have succeeded in doing automatic lexical acquisition by exploiting conventional metaphors as motivations for linguistic forms. In particular, Martin (1988) implemented the MIDAS system, which can automatically extend the lexicon when a new metaphoric use of a word is encountered. Briefly, his approach involves representing conventional metaphors as full-fledge concepts with relations between them, along the lines described above. Senses of particular lex-cons are likewise represented as being motivated by particular metaphors. For example, the verb ‘‘give’’ would have a lex-con whose sense is the same as that of ‘‘infect’’, this sense being motivated by one of the metaphors alluded to above.

The metaphoric mapping that motivated this sense of ‘‘give’’ is represented and available to the language user. For production and understanding, the word sense is treated just like any other, with no appeal to its metaphoric motivation presumed. However, suppose a new use is encountered by the system, say ‘‘John got the flu from Mary.’’ Suppose further, for the sake of this example, that the system has previously encountered only expressions like ‘‘give a cold’’ and ‘‘have a cold’’, so that its representation of the metaphors underlying these expressions is unduly restricted to colds rather than diseases, and that the use of ‘‘get’’ in this domain is new to the system. Then both the available lex-cons of ‘‘get’’, along with the metaphors involving diseases and possession, are brought to bear to hypothesize the lex-con that might be in play. According to Martin, this hypothesis is generated by two kinds of lexical extension processes, *core extension* and *similarity extension*. Understanding ‘‘get a cold’’ given an appropriate prior metaphoric understanding of ‘‘give a cold’’ involves core extension, as the core metaphor ‘‘a cold is a possession’’ is extended to the ‘‘getting’’ concept (i.e., a new sense for ‘‘get’’ is hypothesized which metaphorically extends the central sense); understanding ‘‘get the flu’’ given an understanding of ‘‘get a cold’’ involves similarity extension, as the generalization about a role in the metaphoric structure must be extended from colds to diseases in general. Understanding ‘‘get the flu’’ given an understanding of ‘‘give a cold’’ involves both kinds of extension.

The MIDAS system has been used in conjunction with UC (Wilensky et al. 1988) to metaphorically extend lex-cons in the computer domain. The following is an example of MIDAS learning a new lex-con of the word ‘‘kill’’, given that it knows some metaphoric extensions of this lex-con outside the computer domain.

```
------------------------------------------------------------
#  How can I kill a process?

No valid interpretations. Attempting to extend existing metaphor.

Searching for related known metaphors.

Metaphors found: Kill-Conversation Kill-Delete-Line Kill-Sports-Defeat

Selecting metaphor Kill-Conversation to extend from.

Attempting a similarity extension inference.

Extending similar metaphor Kill-Conversation with
target concept Terminate-Conversation.

Abstracting Terminate-Conversation to ancestor concept

Creating new metaphor:

  Mapping main source concept Killing to main target concept

      Terminate-Computer-Process
```

```
  Mapping source role killer to target role c-proc-termer.

  Mapping source role kill-victim to target role c-proc-termed.

Calling UC:

You can kill a computer process by typing ^c to the shell.
------------------------------------------------------------
```

In this example, the utterance cannot be interpreted because no known lex-con of ‘‘kill’’ is applicable to processes. Thus, MIDAS is called in an attempted to see if a sense extension is possible. MIDAS first retrieves a number of metaphors related to the input; of these, ‘‘Kill-Conversation’’ is chosen as most applicable. A simple similarity extension is attempted, resulting in a proposed ‘‘Terminate-Computer-Process’’ metaphor for interpretation of the input. The interpretation thus provided is passed along to UC, which can answer this question. Meanwhile, the new, metaphoric lex-con is incorporated into UC’s knowledge base, which allows UC’s language generator to use the same terminology in encoding the answer.

The kinds of generalizations MIDAS has been capable of making is illustrated by the following examples:

Given Known Metaphor	System Can Learn
enter LISP	Exist LISP, Enter mail
kill a conversation	kill a process
kill a process	process died
open a file	close a file
give a cold	have a cold, get a cold, give the flu, give an idea

Note that MIDAS waits until it encounters a new use before attempting to extend a current lex-con. While some automatic extension is probably possible, and psychologically plausible (cf. Gropen et al. 1989), in general, the full extent of the metaphor cannot be computed knowing the core. For example, it seems implausible that the lex-con of ‘‘catch’’ in ‘‘catch a cold’’ could be predicted given the other related metaphorical usages and other lex-cons of ‘‘catch’’. However, the degree to which ‘‘local productivity’’ is possible remains an open question.

It is important to acknowledge that there is a serious question about the claim being made by postulating a conventional metaphor underlying utterances. Specifically, the question is what role such a putative metaphor plays synchronically in the language use of the speaker. We view this largely as an empirical issue. However, we emphasize that our view does not require that such metaphors play any explicit role in understanding or production*. Instead, we attribute relevance to them only in the phase of knowledge

*In particular, we need not worry about making erroneous inferences due to the somewhat inconsistent nature of metaphors. For example, we would not want a system to infer that a person giving a cold to someone no longer has that cold. Such a possibility does not arise in our formulation, since we postulate no ‘‘real-time’’ interpretative use of metaphor at all.

On the other hand, neither do we assert that metaphorical motivations play *no* role in understanding. In particular, we think it likely that a metaphoric motivation for a word sense could play a role in lexical disambiguation. Indeed, such a role occurs quite naturally in an activation-type model. What we want to emphasize is that no *interpretative* role is being posited for an underlying metaphoric motivation once a lex-con has been learned.

acquisition. Of course, this attribution is in need of empirical justification as well.

Implicit in the above discussion is that much of the lexicon is domain-dependent. In some sense, this point is too obvious to belabor; of course there is terminology specific to each domain. However, we claim that much of the vocabulary common to many domains also has a domain-specific interpretation. Thus, the word ''kill'' applies to many domains, but its specific meaning in each of them cannot be computed from some single, abstract meaning of ''kill'' and some domain-specific knowledge. To emphasize the point, note that ''killing an engine'' simply means turning off the engine, and not irreversibly destroying it. This sense of ''kill'' transferred to the computer domain would produce a meaning closer to pausing a process, rather than destroying a process. Thus while the particular meanings in each domain will often have a sensible motivation, they each represent domain-specific knowledge, underscoring the importance of lexical acquisition in the task of retargeting a NLP system to a new domain.

3.1. Why MIDAS Works

MIDAS successfully hypothesizes some new metaphoric word senses. To the degree that this program is successful, one can raise the issue of why it works. As suggested above, MIDAS is exploiting a kind of subregularity in the lexicon. Specifically, the same conventional metaphor provides the motivation for a number of different linguistic facts; but the particular facts cannot be completely predicted even if the metaphor is known to be manifest. The best one can do in general, then, is more or less what MIDAS does: Namely, we can use previous instances of metaphoric extension to motivate new ones when they are encountered (or, as we suggested above, to automatically generalize a metaphorically-based lexical fact to other, very ''similar'' lexical usages, thus anticipating what MIDAS only does reactively).

In effect, MIDAS exploits metaphoric subregularity by a form of analogical reasoning. That is, it finds a metaphorical usage that is closest to the given case according to some conceptual metric; it then exploits the structure of the prior metaphoric usage to construct an analogous one for the case at hand, and proposes this new structure as a hypothetical word sense.

An interesting fact about this explanation is that the fact that the basis for the analogy is metaphoric does not play a crucial role in the extension process. That is, if our analysis of why MIDAS works is correct, then it works precisely because it exploits a subregularity in the lexicon. This subregularity happens to be a metaphoric one, but it is the fact that is a subregularity rather than the fact that it is a metaphor that makes it amenable to analogical exploitation.

3.2. Previous Work on Analogy

Analogy, of course, has played a prominent role in traditional linguistics. Indeed, rather influential linguists (for example, Paul (1891) and Bloomfield (1933)) seemed to attribute all novel language use to analogy. However, today, analogy seems almost entirely relegated to diachronic processes. Some notable exceptions to this trend are the work of Derwing and Skousen (1989), although the primary focus of their work is morphological and phonological, and Ohala (e.g., 1988), whose focus is phonetic. We will return to this work below.

Analogy has also been widely studied in artificial intelligence and cognitive psychology, notably by Carbonell (1982, 1983), Burstein (1983), Gentner (1983), Greiner (1985), Falkenhainer et al. (1986), and Indurkhya (1987). In particular, the work of Carbonell (1982), Burstein (1983), and Greiner (1985) is most relevant to our enterprise, as it explores the role of analogy in knowledge acquisition.

The primary focus of Carbonell, Burstein and Greiner is problem solving and concept formation, respectively. These tasks have a rather different character than the language knowledge acquisition task we are attempting to address. As a result, there are pertinent issues specific to the issue of language acquisition

that this work on analogical learning does not directly address. For example, Carbonell's derivational analogy improves a problem solver's performance by exploiting the structure of the solutions to similar, previously solved problems. Many of the details of the process he proposes are specific to problem solving. However, this characteristic hinders the applicability of the approach to tasks that are not intrinsically problem solving. In particular, Carbonell's analogical problem solver does not assume there is a fact to be learned, as in the case in our language application, but rather, is appropriately satisfied with any solution it arrives at that addresses the problem. In contrast, our language task presumes that a ''correct answer'' exists (i.e., there is a fact about the language that must simply be learned), and that analogy can at best suggest it to us.

On the other hand, Burstein and Greiner's work are both concerned with using analogy to hypothesize facts. The focus of Burstein's work is on formulating constraints on the relations to be considered for analogical mapping. Most of these are related to causal knowledge. While the general approach is appealing, the details of his formulation do not appear to have an obvious application to our problem.

Greiner is primarily concerned with determining, given that some sort of similarity is alleged to exist between two concepts, which of many possible analogical conjectures would be useful for solving a problem. His suggestion is that the process of conjecturing is greatly facilitated by prior generalization, indeed, that generalizations have been created precisely for this reason. He mentions in passing that the kind of metaphoric generalizations discussed by Lakoff and Johnson (1980) are examples of the sort of generalization one could use to aid in the acquisition of linguistic terms, exactly the approach that we have adopted.

The problem we are interested in addressing is in some sense simpler than Greiner's problem, and in some sense harder. Greiner's scenario seems artificial in the sense that it is hard to imagine a situation in which one would know that one concept is usefully analogous to another, but have no idea of why this is the case. In particular, we are not given a particular lexical item to which the new item is to be compared. Thus, upon encountering ''get the flu'' for the first time, we (and MIDAS) are not told that ''get'' is similar to ''give'', or any other lexical item. Instead, we must determine the core metaphor to be extended. In Greiner's terms, this would be equivalent to finding and applying some appropriate generalization without any analogical hint. On the other hand, the kind of knowledge that will help is essentially known in advance. That is, we know that what is relevant to the task at hand is the way in which a central sense has been used as the basis for an extended sense.

Greiner's approach is actually quite similar to Alterman's (1985, 1988) ''adaptive planning'' strategy. The idea of adaptive planning is to use knowledge applicable to one situation to determine what to do in another situation. In effect, the knowledge is used analogically to hypothesize what to do in the new situation. Alterman's approach differs from Carbonell's in that the original knowledge does not have to be a problem that was solved by the planner, but merely some knowledge the planner has acquired about how to function in a situation. Greiner's work resembles it in that Alterman assumes the relevant analog is given. For example, faced with the task of using the New York City subway, Alterman's planner is given a description of BART as a guide. Also, generalization is used to guide the search: For example, if the BART step of buying a ticket from a machine fails in the NYC subway, adaptive planning examines the reason for the step by going up an abstraction hierarchy to find that buying a ticket is a way to obtain permission to enter an area in which service is provided. Unlike Greiner, Alterman assumes that the manner in which the analogical information is to be used is known, and uses abstraction only to modify inappropriate parts of the plan.

As with Greiner, Alterman's work is primarily concerned with details specific to planning, e.g., how to chose the proper component for abstraction if a plan step cannot be executed. Thus, it has no obvious direct application to the task at hand.

Many of the differences between analogical reasoning for problem solving and language knowledge acquisition are discussed in Martin (1988).

Another line of related work is that initiated by Rumelhart and McClelland (1987), and further pursued by researchers such as MacWhinney et al. (1989). This line of research proposes connectionist models, which, when trained with the appropriate input, recapitulate many of the kinds of behavior child language learners are known to display. Moreover, some connectionist advocates claim that such models recapitulate this behavior without embodying the types of explicit rules normally found in more traditional AI-approaches. Primarily, the items to which this sort of learning has been applied are morphological associations, e.g., learning the infinite/past-tense pairs of English verbs. However, similar claims about how such mechanisms can extend their scope to new forms is made for other language issues, such as assigning case roles to constituents of sentences (McClelland and Kawamoto, 1986).

It is interesting to contrast this work with other approaches, such as the various analogic ones just mentioned. In particular, MacWhinney et al. (1989) point out that while analogy is an attractive mechanism for the acquisition of morphology, it is difficult to specify which of several possible analogies apply. Also, in a system in which rote, rules and analogy co-exist, it is difficult to decide which one is implicated in any particular case. Connectionist approaches, they point out, readdress these grievances because a single mechanism is proposed that simultaneously captures the three putative processes; analogy, in particular, falls out of pattern extraction.

While we are sympathetic to this view, and find some aspects of these connectionist approaches attractive, there are several drawbacks. First, we believe it is important to understand what the connectionist models are doing in computational terms. Such an analysis can both explain why the connectionist implementation is effective, as well as suggest alternative theories, explanations and implementations.

Second, the Rumelhart-McClelland model in particular maintains no memory of specific cases, but only a statistical summary of them. Indeed, it is this memory limitation that allows their system to alter its behavior on previously learned cases by experiencing different but related ones. However, this view seems at odds with actual human performance. Humans seem to be able to make familiarity judgments independently of their ability to act productively. I.e., the ability to produce a reasonable morphological form for a novel input should not require that we cannot recognize that input as novel. Indeed, one would conjecture that learners confident about their memory of a particular case would not change their behavior in the manner of child language learners.

Third, the Rumelhart-McClelland model in particular, and many connectionist models in general, cannot flexibly exploit the knowledge they have acquired. For example, a connectionist network that has learned to produce ‘‘sang’’ as the past tense of ‘‘sing’’ will not necessarily be able to *recognize* ‘‘sang’’ as the past tense of ‘‘sing’’. There may be any number of possible inputs that produce a given result, and, in any case, the network will not compute these without some additional machinery, if the information is in fact present at all.

It is not our point here to highlight the defects of connectionist approaches in general, or of the Rumelhart-McClelland model in particular. Indeed, our overall position on polysemy and analogy seems quite close to that expressed by MacWhinney (1989). Rather, we want to point out that the connectionist models proposed so far, while intriguing, do not so obviously address all the issue in which we are interested so as to make them the obvious choice for an implementation. Moreover, these models appear to be deficient in some ways, and the insights captured by these models are not necessarily unique to them, and perhaps cannot be fully understood unless expressed in more abstract terms. It should be possible, for example, to express our theory in terms that are neutral with respect to issues like whether rules or abstractions are in some sense made explicit. Thus, such a theory could be amenable to a connectionist implementation or a more conventional one.

The issue of rules versus analogy versus connectionist models is discussed in Derwing & Skousen (1989). Skousen (1990) has a model of analogy that gives rise to quantitative predictions for morphological variation. In effect, Skousen’s model is a method of statistical interpolation. Evidence from every case is summed in a certain manner to predict a probabilistic outcome for any input. Skousen’s method has the

nice property that it has ‘‘perfect memory’’, i.e., it sums data in such a way that close cases do not influence the output when the input has been seen before, but only when it is an unknown case. This behavior is in contrast to Rumelhart and McClelland’s model, for example, which in effect keeps only statistical summaries of the input. On the other hand, the particular measurement Skousen employs seems hard to justify theoretically. It also has a peculiar instability property, so that the addition of a single data point can dramatically alter the outcome in cases where this behavior is not obviously desirable. Finally, the empirical demonstrations he provides seem only marginally related to the particulars of his theoretical results, justifying some kind of analogical reasoning in general but not the details of the model.

Also competing with analogical and connectionist models is the notion of case-based reasoning applied to NLP (cf. Martin and Riesbeck 1986). Case-based reasoning emphasizes the importance of using previous instances to determine subsequent behavior. However, it seems that there is no distinct notion of case-based reasoning to be applied to the problem at hand. That is, if ‘‘case’’ is to mean something other than ‘‘fact’’, then presumably, the cases here would be individual occurrences of lex-cons. However, the interesting reasoning to be performed here is on the level of lex-cons, and not on the level of lex-con occurrences. On the other hand, if we assume that the cases are really abstractions, (i.e., that ‘‘give’’ being metaphorically extended to mean ‘‘infect’’ is a case), then the distinction between case-based reasoning and analogical reasoning vanishes.

We can summarize this discussion as follows: The general problem is to hypothesize, given a set of prior lex-cons, what the best (or at least, a reasonable) new lex-con for a word in a given context. (The given context presumably rules out known lex-cons of a word.) Previous approaches to exploiting this data fall into one of the following classes:

(1) We can find the single, previous stored instance that seems to best apply to the case at hand. We then exploit this instance to propose a fact for the current case. This method is analogical in the strictest sense: It has the form of finding the A:B that is most relevant to C, and then computing D such that A:B::C:D.

(2) We can combine all the data at once to arrive at a hypothesis for a given case. Connectionist and pattern classification schemes work this way, as does Skousen’s analogical model.

The first strategy has the characteristic that a particular item must be selected as the basis of the extrapolation. As MacWhinney et al. point out, this characteristic seems unsuited for some cases. For example, it seems unreasonable to ask, given that a speaker formed the past tense of a new word by adding ‘‘ed’’, which previously known word the speaker used as the basis for the comparison. On the other hand, there do appear to be cases in which even a single example can serve as the basis for a new case. For example, Marchman (1988) attests a child generating ‘‘hat’’ as the past tense of ‘‘hit’’; the only correct instances of such a paradigm in English are ‘‘sit’’/‘‘sat’’, ‘‘spit’’/‘‘spat’’, and ‘‘shit’’/‘‘shat’’; the second is uncommon in American English, and we are hopeful that the child has not attested the third, which is rare in any case.

The second approach has the characteristic that evidence is combined from multiple sources. In Rumelhart and McClelland’s model, this feature results in inappropriate combinations, such as an occasional utterance of ‘‘camed’’, in apparent correspondence to human performance. On the other hand, the approach maintains no distinct memory of individual cases, and is inflexible, in that the encoded knowledge may not be available for other even closely related tasks. For example, there is no obvious way a conflated form like ‘‘camed’’ might be diagnosed as faulty on the basis of never having been seen before. It is unclear for the problem at hand whether this sort of conflation occurs in human learning, and therefore, is a desirable or undesirable property of a model.

3.3. An Approach to Analogy

We propose a framework for analogical reasoning that attempts to overcome some of the drawbacks of these other approaches. The model is as follows:

(1) We store both individual facts, plus the subregularities abstracted from these facts. Along with each individual fact we maintain a ‘‘confidence parameter’’, a measure of our certainty that the stored data is correct.

(2) When an item is encountered such that we can retrieve an individual fact applicable to the known item and of an acceptable certainty level, then the stored fact is simply used.

(3) When an item is encountered for which applicable knowledge does not exist or has an unacceptable certainty level, we attempt to retrieve applicable subregularities. Some of these regularities (presumably, the single subregularities deemed ‘‘most relevant’’) are applied to produce a hypothesized response. If there is no relevant subregularity, then ‘‘similar’’ previous instances, and similar subregularities, are retrieved, and used analogically to produce a response. The result is used if it is more certain than any stored applicable fact; otherwise the stored fact is used.

(4) If the result of (3) is deemed successful, then the result is stored. Also, if the form was the result of the application of a subregularity, it is stored as such. If it were the result of analogy from another instance or subregularity, a generalization is created and stored to serve as a subsequent subregularity.

Like individual instance-based analogical reasoning, this algorithm schema produces symbolic representations that can be used flexibly. However, it relieves the tension between MacWhinney et al. describe between rote, rules and analogy, and which they turned to connectionist architectures to relieve. (Indeed, it incorporates MacWhinney’s (1989) notion of the degree to which an ‘‘extensional pathway’’ has been exploited.)

This algorithm schema leaves us with many difficult unsolved problems, for example, which of several applicable subregularities or cases to choice from. But it is reasonable to believe that general and illuminating principles may be articulated to govern this problem. For example, above we mentioned a conceptual metric that MIDAS uses to establish similarity. Also, there is the quite general principle of preferring a regularity in portion to its specificity. For example, the ‘‘default’’ rule for the past tense might simply be described as a subregularity having no or few requirements on the verb, whereas the more specific subregularities would have more requirements (e.g., for the past tense to end in ‘‘ang’’, the infinitive must terminate with ‘‘ing’’). Thus, when subregularities of different specificity are candidates for analogical extension, the more specific subregularity would be given preference (all other things being equal).

In addition, some of the interesting properties of Rumelhart and McClelland’s model can be exhibited by versions of this algorithm. If a confidence parameter is also associated with the application of subregularities or analogies, then, in low confidence cases, the subregularities and rote memory are in competition with each other; we look at both predictions, and pick the result in which we are most confident. If we apply more that a single subregularity, we will produce the kind of conflated results their model occasionally produces.

Below we will elaborate some of the issues that this model of analogy poses for the problem of learn new lex-cons.

4. An Analogy-based Model of Lexical Acquisition

The following outline is an application of the general approach to analogy just described to the task of proposing new lex-cons. At this point, our work has been rather preliminary, but we can at least sketch out the basic architecture and comment on the problems we have yet to resolve.

(A) Detect unknown lex-con. For example, suppose the system encountered the following phrase:

''at breakfast''

Suppose further that the centrl use of ''breakfast'' (as in ''I eat a breakfast of cereal and juice'') were known to the system, but the determinerless usage were not. In this case, the system would hypothesize that it is lacking a lex-con because of a failure to parse the sentence. The general questions of what sorts of failures suggest lexical inadequacy, and how to localize the source of the inadequacy, are the issues here.

We note in passing that this issue is part of an important and difficult general problem that has received little attention. Even for the metaphoric examples that MIDAS handles, this problem has only been addressed in the simplest cases. Indeed, most research on non-lexicalized metaphor assumes that metaphoric interpretation is attempted when some semantic anomaly is detected. But is it well-known that semantically well-formed expressions can have preferred metaphoric treatment. Thus, ''no man is an island'' violates no semantic conditions but has an obvious metaphoric intent. The usual analysis is that such sentences violate some sort of ''appropriateness conditions''. In this case, it would be inappropriate to communicate such an obvious truth; thus, another interpretation should be preferred.

(B) Find relevant cases/subregularities. Suppose that the system had other instances of ''function noun/determinerless activity noun'' pairs. Cues from the input would be used to suggest that such pairs would be relevant to hypothesizing a new lex-con here. As is the case for MIDAS, features of the input would probably result in other kinds of cases being retrieved whose utility would subsequently be found wanting. In our example, the retrieved cases might include the following:

bed-1/bed-3, class-1/class-4

Here we have numbered lex-cons so that the first element of each pair designates a lex-con involving a central sense, and the latter a determinerless-activity type of lex-con.

Alternatively, we may have already computed and stored relevant subregularities. If so, then these would be retrieved as well. Some principles come into play here; for example, a case characterized by a retrieved subregularity would not be considered further, certainty factors permitting.

Our assumption is that we can retrieve relevant cases by a conjunction of simple cues, like ''noun'', ''functional meaning'', ''extended determinerless noun lex-con'', etc., and then rely on the next phase to discriminate further among these.

(C) Chose the most pertinent case or subregularity. Again, by analogy to MIDAS, some distance metric is used to pick the best datum to analogize from. In this case, perhaps the correct choice would be the following:

class-1/class-4

One motivation for this selection is that ''class'' is compatible with ''at'', as is the case in point.

Finding the right metric is the primary issue here. Our approach is to start with the MIDAS metric, and attempt to generalize it. The MIDAS metric is a simple sum of two factors: (i) the length of the core-relationship from the input source to the source of the candidate metaphor, and (ii) hierarchical distance between the two concepts. Both factors are measured by the number of links in the representation that must be traversed to get from one concept to the other. For example, suppose the input were ''got the flu'' and the candidate the metaphor underlying ''give a cold''. Intuitively, the similarity between getting and giving is that they both result in having; thus, each is one link away

from ‘‘having’’ in the representation, for a total ‘‘core’’ distance of 2. Similarly, ‘‘cold’’ and ‘‘flu’’ are represented as having a common immediate ancestor (say, ‘‘infectious-disease’’), so the hierarchical distance is 2 as well. Thus the total conceptual distance is 4.

The hierarchical distance factor of the MIDAS metric seems directly relevant to other cases. For the example at hand, using this criterion, the candidate ‘‘lunch’’ would be preferred to the candidate ‘‘class’’ if both were retrieved for the input ‘‘breakfast’’. However, there is no obvious counterpart to the core-relationship component of the MIDAS metric. One possible reason for this is that metaphoric extensions are more complex than most other kinds; if so, then the MIDAS metric may still be applicable to the other subregularities, which are just simpler special cases. Alternatively, a simpler metric may appear to be required in these cases because we have been assuming that the cases retrieved by a conjunction of simple cues will be quite similar to the input. In general, this may not be the case, and we may have to rely on a more complex metric to prune them subsequently.

Note that taking some additional factors into account may not increase the complexity of the closeness metric. For example, one consideration may be the context in which the lex-con is allowed. Above we suggested that ‘‘class’’ might be judged more similar to the input because both can appear as the complement of ‘‘at’’. Thus, a case retrieved that was compatible only with ‘‘in’’ might be less preferable to one compatible with the same preposition. However, including this feature in the measurement just appears to involving adding another hierarchical term.

The MIDAS metric is also unnormalized. In particular, adding another dimension of closeness always increases the conceptual distance, contrary to intuition. This feature does not cause any problems in MIDAS because all comparisons use the same number of dimensions, but we are not confident that it will hold up in the general case. There are several potential solutions to this problem should it in fact arise.

(D) Analogize to a new meaning. Given the best case or subregularity, the system will attempt to hypothesize a new lex-con. For example, in the case at hand, we would like to be produced a representation for the lex-con whose sense is the meaning in quotes:

class-1/class-4 ::
breakfast-1/‘‘situation of eating breakfast’’

In the case of MIDAS, the metaphoric structure of previous examples is assumed to be available. Then, once a best match is established, it is relatively straightforward to generalize or extend this structure to apply to the new input. The same would appear to be true in the general case, provided that the relation between stored polysemous lex-cons is readily available. However, such an explicit relation may not always be available.

(E) Determine the extent of generalization. Supposing that a single new lex-con can be successfully proposed, the question arises as to whether just this particular lex-con is all the system can hypothesize, or whether some local productivity is possible. For example, if this is the first meal term the system has seen as having a determinerless activity lex-con, we suspect that only the single lex-con should be generated. However, if it is the second such meal term, then the first one would have been the likely basis for the analogy, and a generalization to meal terms in general may be attempted. How local productivity works is not well-understood.

We do not believe that the solution to this question will ultimately be separable from that for (D) above, but we do feel that the question is important enough to be posed separately. Also, for both steps, interaction with another agent to validate a hypothesis is possible.

(F) Record a new entry. The new lex-con needs to be stored in the lexicon, and indexed for further reference. Also, the question of whether new subregularities need to be computed and stored is addressed here. This task may interact closely with (E), although generalizing to unattested cases and computing explicit subregularities are logically independent.

There are many additional problems to be addressed beyond the ones alluded to above. In particular, there is the issue of the role of world knowledge in the proposed process. In the example above, the system must know that the activity of eating is the primary one associated with breakfast. A more dramatic example is the role of world knowledge in hypothesizing the meaning of ''treed'' in expressions like ''the dog treed the cat'', assuming that the system is acquainted with the noun ''tree''. All the proposed analogical reasoning mechanism can do is suggest that some specific activity associated with trees is involved; the application of world knowledge would have to do the rest.

5. Other Directions of Investigation

We have also been investigating exploiting subregularities in ''intelligent dictionary reading''. This proposal involves a separate idea, namely, that one could best use a dictionary to gain lexical knowledge by bringing to bear on it a full natural language processing capability. We are just beginning to develop such a system as part of other, related research. One problem we have encountered is that dictionaries are full of inaccuracies about the meaning of words. For example, even relatively good dictionaries have poor entries for the likes of determinerless nouns like ''bed''. E.g., Webster's New World (Second Edition) simply lists ''bedtime'' as a lex-con of ''bed''; Longman's Dictionary of Contemporary English (New Edition) uses ''in bed'' as an example of the ordinary noun ''bed'', then explicitly lists the phrase ''time for bed'' as meaning ''time to go to sleep'', and gives a few other determinerless usages, leaving it to the reader to infer a generalization.* However, a dictionary reader with knowledge of the subregularity mentioned above might be able to correct such deficiencies, and come up with a better meaning that the one the dictionary supplies. Thus, we plan to explore augmenting our intelligent dictionary reader with the ability to use subregularities to compensate for inadequate dictionary entries.

The discussion above has been primarily restricted to nouns and verbs. However, there is at least suggestive evidence for similar analyses for other linguistic objects. For example, Sweetser (1990) argues that ''because'' is polysemous with the central sense projected along epistemic and speech act dimensions to account for ''He's married because I noticed he's wearing a ring'' (epistemic) and ''Give me a minute because I'm pretty busy'' (speech act) (see rule 29 above). Brugman (1981,1984) gives a compelling polysemy account of ''over'' and of ''very''. We expect that attempting to learn such lex-cons by analogy will be difficult because of the lack of a clear understanding of the semantic notions involved.

It seems reasonable to apply the same approach to acquiring the semantics of constructions. In particular, it seems straightforward to move up to items just larger than lexical items, such as verb-particle combinations and nominal compounds. For example, Lindner's (1981) analysis of the verb-particle constructions involving ''up'' and ''out'' would appear to have a useful interpretation in this framework.

It also seems possible to analyze the semantics of basic grammatical complements along these lines. For example, Goldberg (1989, 1990) has argued that much of the data on the ditransitive, and on its acquisition, can be accounted for in terms of subclasses that are conventionally associated with the construction itself, rather than lexical rules and transformations as proposed by Gropen et al. (1989). We believe that this view can be reconciled with the view of language acquisition advocated here. Specifically, this would mean using our analogical reasoning mechanism to propose extended meanings of the ditransitive as new uses are encountered. For example, a system knowing that ''pour Bill a drink'' is acceptable, and that it means ''pour a drink for Bill'', should be able to recognize ''mix Bill a drink'' and ''bake Bill a cake'' as

*Longman's also defines ''make the bed'' as ''make it ready for sleeping in''. We have no idea how to cope with such errors, but they do underscore the problem.

evidence that ‘‘mix’’ and ‘‘bake’’ sanction the ditransitive, and that, in its use with them, it has the same meaning as with ‘‘pour’’. (Indeed, there is empirical evidence (Gropen et al 1989) that people make such generalizations automatically, without any evidence, subject to various phonological and other contraints.) More interesting still is to hypothesize the meaning of the ditransitive in association with verbs like ‘‘pour’’ given an understanding of ‘‘give Bill a book’’ and the evidence such as ‘‘pour Bill a drink’’. While we cannot be certain at this point that the same mechanism can apply for words as for constructions, the parallels seem rather striking.

6. Conclusions

Since our implementation of learning models that exploit subregularities is restricted only to the metaphoric subcase, many difficulties and new issues undoubtably remain to be discovered. We think our theoretical analysis is useful on several accounts, nevertheless. Most generally, it may help us further understand and differentiate between the computational and algorithmic aspects of cognition in general. Analogy, case-based learning, connectionist learning, and statistical pattern classification are all ways of approaching what appears to be the same issue; perhaps the problem of lexical acquisition will serve as the topic to unify what would otherwise remain superficially diverse approaches. On the level of language per se, we believe that, while the details of our analysis may be faulty, the general view of language as a system of motivated but largely unpredictable cases will be borne out. More specifically, we feel that we have illuminated some of the richness of the relation of the lexicon and cognition, which is interesting in its own right, and is not otherwise explicated by conventional linguistic or purely conceptual analyses.

7. References

(1) Alterman, Richard. Adaptive Planning: Refitting Old Plans To New Situations. In the *Proceedings of The Seventh Annual Conference of the Cognitive Science Society*, 1985.

(2) Alterman, Richard. Adaptive Planning. In *Cognitive Science*, vol. 12, pp. 393-421, 1988.

(3) Bloomfield, L. *Language*. New York: Holt, Rinehart & Winston, 1933.

(4) Braverman, M. and Wilensky, R. Toward a Unification of Case-Based Reasoning and Explanation-Based Learning. In the *Proceedings of the AAAI Symposium of Cased-Based Reasoning*, Stanford, CA., March, 1990.

(5) Brugman, Claudia. The Story of *Over*. University of California, Berkeley M.A. thesis, unpublished. Distributed by the Indiana University Linguistics Club. 1981.

(6) Brugman, Claudia. The *Very* Idea: A Case-Study in Polysemy and Cross-Lexical Generalization. In *Papers from the Twentieth Regional Meeting of the Chicago Linguistics Society*. pp. 21-38. 1984.

(7) Burstein, Mark H. Concept Formation by Incremental Analogical Reasoning and Debugging. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (eds.), *Machine Learning: An Artificial Intelligence Approach*, vol. II. Tioga Press, Palo Alto, California, 1982.

(8) Butterworth, B. Lexical representation. In B. Butterworth (ed.)m *Language Production* , vol. 2. Academic Press, New York, 1983.

(9) Carbonell, Jaime. Metaphor: A Key to Extensible Semantic Analysis. In the *Proceedings of the 18th Meeting of the Association for Computational Linguistics*. 1980.

(10) Carbonell, Jaime. Learning by analogy: Formulating and Generalizing Plans from Past Experience. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (eds.) *Machine Learning: An Artificial Intelligence Approach*. Tioga Press, Palo Alto, California, 1982.

(11) Carbonell, Jaime. Derivational analogy: A Theory of Reconstructive Problem Solving and Expertise Acquisition. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (eds.) *Machine Learning: An Artificial Intelligence Approach*, vol. II. Tioga Press, Palo Alto, California, 1982.

(12) Cutler, A. Lexical complexity and sentence processing. In G. B. Flores d'Arcais & and R. J. Jarvella (eds.), *The Process of Language Understanding*, pp. 43-79. Wiley, New York, 1983.

(13) Derwing, Bruce L. and Skousen, Royal. Real-time mrophology: Symbolic Rules or Analogical Networks? *Proceedings of the Fifteenth Annual Meeting of the Berkeley Linguistics Society*. Berkeley, California, February, 1989.

(14) Falkenhainer, B., Forbus, K. D., and Gentner, D. The Structure Mapping Engine. In *Proceedings of AAAI-86*. Philadelphia, PA, 1986.

(15) Fillmore, Charles. The Organization of Semantic Information in the Lexicon. In *Papers from the Parasession on the Lexicon*, Chicago Linguistics Society. pp. 1-11. 1978.

(16) Gentner, D. Structure Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, pp. 155-170, 1983.

(17) Goldberg, Adele. A Unified Account of the Semantics of the Ditransitive Construction. *BLS 15*, 1989.

(18) Goldberg, Adele. The Inherent Semantics of Argument Structure: The Case of the English Ditransitive Construction. Unpublished manuscript, 1990.

(19) Granger, R. H. FOUL-UP: A Program that figures out the meanings of words from context. In the *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Cambridge, MA. 1977.

(20) Gropen, J., Pinker, S., Hollander, M., Goldberg, R., & Wilson, R. The learnability and acquisition of the dative alternation. In *Language*, vol. 65, no. 2, June 1989.

(21) Hinton, G. E. & Sejnowski, T. J. Learning and Relearning in Boltzmann Machines. In J. L. McClelland, D. E. Rumelhart, and the PDP Research Group, *Parallel Distributed Processing*, vol. 1. The MIT Press, Cambridge, Massachusetts, 1986.

(22) Hobbs, J. Metaphor Schemata and Selective Inferencing. SRI Technical Note 204, December 1979.

(23) Indurkhya, B. Approximate Semantic Transference: A Computational Theory of Metaphors and Analogy. In *Cognitive Science*, 11, pp. 445-480, 1987.

(24) Lakoff, G. and Johnson, M. *Metaphors We Live By*. University of Chicago Press, 1980.

(25) Lindner, Susan. A Lexico-Semantic Analysis of Verb-Particle Constructions with UP and OUT. UCSD dissertation. 1981.

(26) MacWhinney, B. Competition and Lexical Categorization. In R. Corrigan, F. Eckman and M. Noonan, *Linguistic Categorization.* John Benjamins Publishing Company, Amsterdam/Philadelphia, 1989.

(27) MacWhinney, B., Leinbach, J., Taraban, R. & McDonald, J. Language Learning: Cues or Rules? In the *Journal of Memory and Language*, 28, pp. 255-277, 1989.

(28) Marchman, V. A. Rules and regularities in the acquisition of the English past tense. In the *Newsletter of the Center for Research in Language*, vol. 2, no. 4. University of San Diego, La Jolla, CA. April, 1988.

(29) Martin, Charles E. and Riesbeck, Christopher K. Uniform Parsing and Inferencing for learning. In the *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, PA. August, 1986.

(30) Martin, James. Knowledge Acquisition through Natural Language Dialogue. In the *Proceedings of the 2nd Conference on Artificial Intelligence Applications*. Miami, Florida. December, 1985.

(31) Martin, James. A Computational Theory of Metaphor. Berkeley Computer Science Technical Report no. UCB/CSD 88/465. November 1988.

(32) McClelland, J. L. and Kawamoto, A. H. Mechanisms of Sentence Processing: Assigning Roles to Constituents. In J. L. McClelland, D. E. Rumelhart, and the PDP Research Group, *Parallel Distributed Processing*, vol. 2. The MIT Press, Cambridge, Massachusetts, 1986.

(33) Norvig, Peter. A Unified Theory of Inference for Text Understanding. Ph. D. Thesis, University of California, Berkeley Technical Report No. UCB/CSD 87/339, 1987.

(34) Norvig, Peter. Building a large lexicon with lexical network theory. In the *Proceedings of the IJCAI Workshop on Lexical Acquisition.* August 1989.

(35) Norvig, Peter and Lakoff, George. Taking: A Study in Lexical Network Theory. In the *Proceedings of the Thirteenth Annual Meeting of the Berkeley Linguistics Society*. Berkeley, CA. February 1987.

(36) Nunberg, G. *The Pragmatics of Reference*. Bloomington, In.: Indiana University Linguistics Club. 1978.

(37) Ohala, J. J. The cognitive cost of phonological analysis. University of California at Berkeley Cognitive Science Report #53, 1988.

(38) Paul, H. *Principles of the History of Language*. Longmans, Green, London, 1891.

(39) Rumelhart, D., & McClelland, J. Learning the past tenses of English verbs: Implicit rules of parallel distributed processes? In B. MacWhinney (ed.), *Mechanisms of Language Acquisition*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1987.

(40) de Saussure, Ferdinand. *A Course in General Linguistics.* McGraw-Hill, New York, 1959.

(41) Selfridge, M. Computer Modeling of Comprehension Development. In W. G. Lehnert & M. H. Ringle, *Strategies for Natural Language Processing*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1982.

(42) Skousen, R. *Analogical Modeling of Language.* Kluwer, Dordrecht, 1990.

(43) Sweetser, E. *From Etymology to Pragmatics*. Cambridge: Cambridge University Press, 1990.

(44) Venneman, T. Words and syllables in natural generative phonology. In *Papers from the parasession on natural phonology*, pp. 346-374. 1974.

(45) Wilensky, R. Primal Content and Actual Content: An Antidote to Literal Meaning. In the **Journal of Pragmatics**, Vol. 13, pp.163-186. Elsevier Science Publishers B.V. (North-Holland), Amsterdam, The Netherlands. 1989.

(46) Wilensky, R. Sentences, Situations and Propositions. To appear in J. Sowa (ed.) *Formal Aspects of Semantic Networks*. Morgan Kaufman Publishers, Los Altos, California, 1990.

(47) Wilensky, R., Mayfield, J., Chin, D., Luria, M., Martin, J. and Wu, D. The Berkeley UNIX Consultant Project. *Computational Linguistics* 14-4, December 1988.

(48) Wittgenstein, Ludwig. *Philosophical Investigations*. Macmillian, New York, 1953.