Revue québécoise de linguistique

WordNet: An Electronic Lexical Reference System Based on Theories of Lexical Memory

George Miller, Christiane Fellbaum, Judy Kegl et Katherine Miller

Citer cet article

Résumé de l'article

Cet article fait la description de WordNet, système de référence électronique, dont le dessin est basé sur des théories psycholinguistiques concernant la mémoire lexicale et l'organisation mentale des mots.

Les noms, les verbes et les adjectifs anglais sont organisés en groupes synonymes (les « synsets »), chacun représentant un concept lexical. Trois relations principales — l'hyponymie, la méronymie et l'antonymie — servent à établir les rapports conceptuels entre les « synsets ». Les présuppositions qui lient les verbes sont indiquées ainsi que leurs contextes syntaxiques et sémantiques.

En tâchant de miroiter l'organisation mentale des concepts lexicaux, WordNet pourrait servir l'utilisateur sans formation en linguistique.

WORDNET: AN ELECTRONIC LEXICAL REFERENCE SYSTEM BASED ON THEORIES OF LEXICAL MEMORY

George Miller, Christiane Fellbaum, Judy Kegl, and Katherine Miller[*]

Introduction

WordNet is an electronic lexical reference system for English, designed in accordance with psycholinguistic theories of the organization of human lexical memory. This novel lexical reference system for English is being developed in the form of an electronic database. Its design derives from psychological and linguistic theories about how lexical information is organized and stored in the memories of people who know English well and speak it fluently. The success of this experimental system would demonstrate the adequacy of the theories from which it derives, but even if those theories must be revised or replaced, the lexical database that is being developed in order to test them will be adaptable to a variety of practical applications. WordNet, supplemented on-line by machine-readable dictionaries and made available via a multi-window workstation, can be profitably incorporated into any task that is facilitated by easy access to lexical information.

Word knowledge is analyzed into: (1) the sound pattern, (2) the concept that the sound pattern can express, and (3) the association of sound and concept. Sounds and concepts are learned differently: as a consequence, different kinds of lexical relations are established: (1) phonological (e.g., rhyme) and morphological relations (e.g., inflection, derivation, compounding) are word-specific, whereas (2) semantic relations (e.g., synonymy, subordination, part-whole) are truth-functional.

Both kinds of relations are incorporated. In WordNet, a concept is represented by a set of synonyms that can be used, in appropriate contexts, to express it; other semantic relations are represented by labeled pointers between the related concepts. WordNet will test the adequacy of current ideas about the structure of the lexicon by testing whether a realistically large sample of the English lexicon can be represented in this way.

The use of synonym sets is both an innovative and an expedient approach to dictionary design. Standard dictionaries develop uniform semantic representations for all the lexical items in English by systematizing the writing of sense definitions or by determining a set of linguistic primitives that constitute the meaning of lexical items. WordNet circumvents the writing and systematizing of sense definitions by representing concepts as relations among words arranged in a "vocabulary matrix," a giant network coding various relations by means of connections between words. It simply looks along a given row of the vocabulary matrix, notes all the words that can be used to express the same concept, and then substitutes that synonym set for the statement of the concept. If one accesses the dictionary via the horizontal word list, one gets a view of the polysemy of a word (all the different concepts that word can express). On the other hand, if one accesses the matrix from the vertical concept list, one gets a row containing all the different synonymous words that express a given concept.

Once the basic matrix is in place, an elaborate system of cross-referencing allows the coding of various relations between synonym sets, including relations of antonymy, superordination, subordination, part-whole, grading, presupposition. Finally, more complex relations termed "theories" can be encoded, including topics, semantic fields, and areas of discourse. WordNet is free from any requirement to encode all the information about a word in the confines of a single entry. Furthermore, the nonlinear nature of this net, together with the freedom afforded by

computer access, captures many important relations obscured by the formatting constraints of hand-held dictionaries.

Psycholinguistic Issues

What a language user must know and how that knowledge is organized are related but separable questions. In order to speak and understand any language, it is necessary to know the sounds and meanings of thousands of different lexical units — some idea of what a language user must know can be gathered from reading an ordinary desk dictionary.

How that lexical knowledge is organized, however, is a much more difficult question. In a printed dictionary it is organized alphabetically. In a person's memory the organization is much more complex. Lexical memory must be so organized that the sounds and the contextually appropriate meanings of thousands of different words can be retrieved from memory at rapid rates. The conversational use of language would scarcely be possible unless the lexical memory system were well organized to support such rapid retrieval. The nature of this organization, and how it comes to be constructed during the process of learning a language, are basic questions for psycholinguistic research. Questions about the organization of lexical memory are easier to consider, however, if one first becomes clear about what a language user must know.

A vocabulary matrix is sufficiently general to represent any lexicon, whether it exists in a person, in a book, or in a computer. It contains a representation of the phonological form of a word and a representation of the conceptual content of a word, along with the associative bond connecting them. The vocabulary matrix is not a complete model of a human language user's lexical knowledge, however. A good model of a person's lexical knowledge would have to include the phonological and morphological features of the words and the semantic and pragmatic relations among lexical concepts.

Lexical Relations

The vocabulary matrix captures the basic structure of lexical memory, but it neglects the complex relations that exist between words.

Phonological relations like rhyme, and morphological relations between derivatives (e.g. *navy* and *naval* , or *high*, *higher*, and *highest*) or compounds (e.g., *ship*, *board*, and *shipboard*, or *pocket*, *pick*, and *pickpocket*) are real and recognizable to anyone who knows English but are not shown in the vocabulary matrix. Judgments of such relations between words depend on familiarity with the spoken patterns; they are rapid and accurate for highly practiced words, but slow and unreliable for infrequently used and unfamiliar words.

Conceptual relations are not shown in the vocabulary matrix, either. A wide variety of such relations have been studied by psycholinguists (Chaffin & Herrmann, 1984). For example, subordination and superordination (e.g., a maple is a tree, and a tree is a plant), which are relations between concepts, not between words, do not appear in a simple listing of lexical concepts. Linguists and lexicographers refer to subordination as *hyponymy*. Hyponymy generates a hierarchical structure, a taxonomy, in the lexicon.

The part-whole relation, called meronymy, is also a relation between concepts (Iris, Litowitz, & Evens, 1985). Simple examples are easily found. For example, a car has an engine, an engine has a carburetor, and a carburetor has a flutter valve; that is, *flutter valve* is a meronym of *carburetor* and *carburetor* is a meronym of *engine*. Like hyponymy, meronymy exhibits a hierarchical organization where, instead of the ISA relation, the HASA relation is exploited.

No adequate theory of the organization of lexical memory can ignore the strong formal relations between the columns or the strong semantic relations between the rows of the vocabulary matrix. Lexical relations must, therefore, be included in any electronic system that hopes to simulate the human memory structure. The vocabulary matrix is merely a skeleton; it must be fleshed out with many formal and conceptual relations.

Sources of Evidence

Any theory must rest on a body of factual data. Two rather different kinds of factual data are available to support claims about the organization of lexical memory. One is linguistic: the data underlying theories of lexical organization are conveniently summarized in printed dictionaries and thesauruses. The second is psychological: a variety of experimental investigations have provided evidence for the psychological reality of the hypothesized mental structures. A few words about each should suffice to indicate the general character of the available data.

First, the linguistic evidence. Dictionaries and thesauruses that summarize the relevant linguistic information derive ultimately from the recorded use of the language by native speakers, and from native speakers' subjective judgments.

How words are strung together in sentences and larger units of discourse provides important information for a person — a child or a foreigner — trying to learn the vocabulary of a language. The ability to produce acceptable sentences is an important indicator that the writer or speaker knows the words they contain. Lexicographers collect such sentences, classify them according to the words they contain, and cite them as the bases for the definitions that they put in their dictionaries. When dealing with dead languages, the written corpus is the only evidence available. The inconvenience of this kind of evidence is that there are many different words, many of them relatively rare, and enormous quantities of text must sometimes be searched in order to turn up a mere handful of examples of sentences using the word that is being studied.

In addition to corpus-based lexicography, some linguists and lexicographers also rely on native speaker intuition. Since native speakers are competent to produce and understand an indefinite variety of sentences containing any particular word, their implicit knowledge of the language provides a basis for subjective judgments that can be used as primary data.

Both linguists and psychologists have developed methods to tap into the linguistic intuitions of others. For example, psychologists sometimes give native speakers a word and ask what other words it suggests, or they may constrain the person's associations by specific instructions, like "What is a kind of plant?" or

"List all the trees you can think of." Judgments that ISA or HASA relations hold take the form of judgments of the truth or falsity of such statements as "A maple is a tree" or "A gasoline engine has a carburetor." General world knowledge is involved in such judgments, of course. Linguists, on the other hand, are more likely to frame questions in terms of sentences, such as "Do S_1 and S_2 have the same meaning?" where S_1 and S_2 are identical sentences except for a pair of words whose meanings are to be compared. Or they may ask for judgments of oddness: "pines and other maples" sounds odd, "trees and other maples" sounds odd, but "pines and other trees" does not.

The experimental evidence gathered by psychologists is of a different nature. By and large, psycholinguistic experiments presuppose the validity of the general structures that linguists and lexicographers have identified and try instead to test hypotheses concerning the way such structures arise or how they contribute to other cognitive processes.

For example, linguists distinguish between open and closed class words. Open class words are nouns, verbs, adjectives, and most adverbs; the language has a great many different open class words, and new ones can easily be added to the vocabulary as needed. Closed class words are articles, prepositions, conjunctions, and some adverbs. English has a limited number of them (around 100); they provide important information about the syntactic structures of the sentences, and new ones are difficult to add to the language. Psychologists have adopted this distinction, calling open class items *content words* and closed class items *function words*.

Psychologists have found a variety of behavioral data correlated with this distinction. For example, hesitations in conversational speech tend to occur before content words, not before function words (Goldman-Eisler, 1968). Or, to take a different example, good readers tend to direct their gaze at content words and to skip over function words. Since there are relatively few function words and they are used in every sentence, they occur much more frequently than do content words; consequently, psychologists translate the content/function distinction into a word-frequency distinction. It is the infrequent and unpredictable words that cause a speaker to hesitate, and the less frequently a word is used, the more time a reader will spend looking at it (Carpenter & Just, 1983).

This word-frequency effect is also found in other experiments. In the lexical decision task, for example, readers are asked to decide as quickly as they can whether a particular string of letters spells an English word (Whaley, 1978). It has been found that the time required to say *Yes* to actual words decreases as the word's frequency of occurrence increases (Gordon, 1985). A wealth of such results strongly suggests that a person's access to lexical information in memory is faster and easier the more often the word has been encountered previously.

The conceptual dimension of lexical memory has also been explored experimentally by psychologists. One of the landmark studies was the work of Collins & Quillian (1969). They reported that it takes people longer to judge the truth of the statement *A canary is an animal* than to judge *A canary is a bird*. They attributed such observations to the fact that *bird* is the immediate superordinate of *canary*, whereas *animal* is a more remote superordinate.

Collins & Quillian's paper stimulated extensive research into the organization of semantic memory. That work need not be summarized here; an excellent review has been written by Smith (1978). It suffices for the present purpose to indicate what kinds of experimental evidence is available to support the claim that words are doubly entered in lexical memory.

Although the work outlined in the present paper is not basic research in the sense that the experimental studies just mentioned clearly are, it can nevertheless contribute to the understanding of the organization of lexical memory. The contribution should follow from the inclusion of a sizeable fraction of the English lexicon, which can act as an antidote against premature enthusiasm. Psychological experiments are almost necessarily conducted with a small number of words and then the results are assumed (often implicitly) to generalize over the entire vocabulary. A failure to look for negative evidence can tempt one into serious mistakes.

This temptation can be strong when lexical properties and relations are at issue. When only bits and pieces of lexical data have been examined, a theorist may begin to see patterns, to formulate hypotheses, and to search for examples to support those hypotheses. Moreover, supporting examples are usually found: it is easy to find words that will fit nicely into almost any pattern a reasonable person

might invent. But the fact that supporting examples can be found does not really test the hypothesis. A list of positive instances — even a long list — offers no assurance that there are no negative instances. Therefore, in order to avoid favoritism (even unconscious favoritism) for words that confirm one's hypothesis, it is advisable to test hypotheses against a large collection of words, a collection assembled in ignorance of the hypotheses in question.

WordNet: An Implementation Of a Model of Lexical Organization

WordNet is an electronic lexical reference system designed in accordance with the theories summarized above. The first step in creating WordNet was to invent an electronic version of the vocabulary matrix.

Synonym Sets

A major problem facing anyone who would construct a vocabulary matrix is how to represent all the various concepts that words can express.

Lexicographers represent lexical concepts by circumlocution. That is to say, they use words to define words. Lexicographers take great pains to distinguish among different senses that a given word can be used to express, but they pay far less attention to establishing a common phrasing for the same sense when it appears in entries for different words. For example, in one widely used dictionary the same lexical concept is phrased as "inferior in quality or value" in the definition of *poor* and as "of little or less importance, value, or merit," in the definition of *inferior*. If WordNet represented the lexical concepts in the vocabulary matrix by definitional phrases borrowed from a conventional dictionary, many, perhaps most, synonymic relations would be overlooked.

Some standard convention for expressing word senses is required. At first glance it might seem that there are many options to choose from. Many different notations for lexical concepts have been proposed (see, for example, Anderson, 1976; Cullingford, 1986; Jackendoff, 1983; Katz, 1972; Miller & Johnson-Laird, 1976; Norman & Rumelhart, 1975; Schank, 1972; Sowa, 1984;

Talmy,1985).   It might be possible to identify one best suited for the present purpose.   But such notations, although easier to standardize than the usual circumlocutions, have been worked out in detail for only small sets of English words, usually for whatever words happen to have been used for demonstration purposes.

How, then, should the list of lexical concepts be constructed? In order to proceed with WordNet, we have used *synonym sets* to represent lexical concepts. That is to say, the identifier for the concept on any given row of the vocabulary matrix is given by the list of words that (in an appropriate context) can be used to express that concept. Actually, since the synonym sets will be numbered, each concept will be represented in the system by a number, but displayed to the user as a set of words having a shared meaning.

It should be noted that synonym sets, unlike dictionary entries, do not have headwords. In a book of synonyms, one entry might have *pipe* as the headword, alphabetized under *P* with "tube" as its contents, and another entry with *tube* as the headword, alphabetized under *T* with "pipe" as its contents. In WordNet, the synonym set { pipe, tube, } stands as an elementary component, and neither word is ahead of the other. This practice has the advantage of symmetry: if *x* is a synonym of *y*, then *y* is a synonym of *x*.

Because synonymy is so central to the design of WordNet, it resembles the electronic thesauruses that are now becoming available commercially (Raskin, 1987). WordNet goes beyond those products, however, by incorporating conceptual relations other than synonymy.

The Master List

Once a satisfactory list of synonym sets is available, it is simple to index it. That is to say, an alphabetical listing of all the words in all the synonym sets can be constructed where each word is followed by the numbers of all the synonym sets of which it is a member. This list, which we have been referring to as the master list, can also contain information that is word-specific and not dependent on the concepts that the word can be used to express. For example, the master list includes information concerning the relative frequencies of use of each word.

Conceptual Relations

As of June 1988, the WordNet files included approximately 11,500 different nouns organized into over 7,000 synonym sets; over 6,000 different verbs organized into over 4,000 synonym sets; and over 13,000 different adjectives organized into more than 10,000 synonym sets. That gave a total of over 29,500 different words organized into approximately 21,000 synonym sets. The next step was to introduce relations between lexical concepts: not only semantic relations (Cruse, 1986; Evens *et al.* , 1983; Lyons, 1977, ch. 9), but others as well. While additional synonym sets continue to be added, we are now introducing cross-references designed to represent conceptual relations.

Conceptual relations are represented in WordNet by cross-references between synonym sets. Each synonym set, therefore, will be followed by a list of the numbers of other synonym sets related to it in particular ways.

Hyponymy, for example, can be introduced in WordNet by appending to a given synonym set one number that points to its superordinate term and other numbers that point to its hyponyms. The relation of meronymy is similar. Since meronymy generates a part-whole hierarchy that is structurally similar to a hyponymic hierarchy, it can be introduced in WordNet in a similar manner, by labeled cross-references.

The Hyponymic Hierarchy

Cognitive psychologists have been interested in lexical hierarchies at least since Collins & Quillian (1969) proposed them as a model of semantic memory. According to the theory, concepts are nodes linked by labeled arcs. Workers in artificial intelligence had observed that a hierarchy of nodes linked by ISA relations is an efficient storage system: since all of the properties attributed to a superordinate node are inherited by its hyponyms, those properties need be stored only once; they need not be stored separately with every hyponym. For example, when you are told that Cuthbert is a cat you know immediately that Cuthbert has four legs, fur, retractable claws, purrs, etc. It is not necessary to learn each of those properties separately.

During the past quarter century, therefore, the hyponymic hierarchy for nominal concepts has been widely exploited. For example, the psychologist Keil (1979) called it an "ontological tree" and used it to organize his observations of vocabulary growth in young children. Other workers have not found the hierarchy as neat and tidy as Keil did: the computer scientist Cullingford (1986) called it a "tangled hierarchy" (e.g., *knife* is a hyponym of both *utensil* and *weapon*) and used it as the basic classification scheme underlying his natural language processing system. Others have proposed other variations. But even those who disagree about the details do agree on the general idea that some kind of semantic hierarchy is required in order to represent lexical knowledge.

It is not difficult to construct demonstrations based on small fragments of the hyponymic hierarchy, but constructing it for a broad sample of the English lexicon is a formidable task. Much of the information required is contained in the defining phrases of standard dictionaries, where a common form of definition is: "x is a y that P," where x is a hyponym of y and P is a relative clause that distinguishes x from the other hyponyms of y. For example, *The Longman's Dictionary of Contemporary English* (Procter, 1978) says that a TREE is "a type of tall PLANT with a wooden trunk and branches, that lives for many years," from which it is obvious that TREE is a hyponym of PLANT.

This kind of information can be extracted from a machine-readable dictionary (Amsler, 1980, 1981; Amsler & White, 1979; Chodorow, Byrd & Heidorn, 1985). The results make it clear that lexicographers work with a fundamentally consistent semantic hierarchy. Unfortunately, however, definitions in standard dictionaries are not written with this analysis in mind, and fortuitous variations in the phraseology of related definitions sometimes obscure their relatedness.

One feature of dictionaries that deserves comment is that it is much easier to identify superordinates from the defining phrases than to identify hyponyms. For example, the definition of *tree* will almost necessarily say that a tree is a plant, but it will not go on to say of *apple, elm, fir, maple, pine* , etc. that they are all trees; for that information a user must consult the individual entries for *apple, elm, fir,* etc., which presupposes that users already have the information that they are searching for. In WordNet moving down the hyponymic hierarchy should be as easy as moving up.

The hyponymic hierarchy is also apparent in standard thesauruses: *Roget's International Thesaurus* has 6-8 tiers of categories, going progressively from highly abstract generic categories to highly concrete specific categories. However, Roget and his successors were not slavishly devoted to the hyponymic relation, and careful judgment is sometimes required in order to extract the hyponymic relation from all the other information in an entry. For example, in Chapman's (1977) version of Roget's thesaurus the path from the root of the hierarchy out to one sense of the word *pipe* goes as follows:

Class Two: Space
  III. Structure; Form
    B. Special Form
      255. Sphericity, Rotundity
      Nouns
        255.4 cylinder, cylindroid, cylindr(o)-, pipe, tube

Although one can agree that a pipe is a cylinder and that a cylinder is a form, the rest of this path introduces other kinds of information. In particular, the more generic concepts seem rather arbitrary. Sedelow & Sedelow (1986) comment that there is much greater descriptive and analytic power, semantically, in the lower tiers of Roget's thesaurus.

In most cases, the judgments required to settle questions about hyponymic relations are not difficult. In order to decide whether x is a hyponym of y, substitute them into a standard frame of the form: x ISA y, then judge, on the basis of general knowledge about such things, whether the resulting proposition is true or false. If it is true, then x can be accepted as a hyponym of y. Uncertainty about the truth value may complicate the judgment when the judge is not knowledgeable about x 's and y 's or when highly abstract concepts are involved, e.g., is VIRTUE a hyponym of IDEA? But the large majority of cases are easily decided.

By using a collection of dictionaries and thesauruses, liberally seasoned with linguistic intuitions, WordNet editors have introduced hyponymic relations into the synonym sets with little trouble. In some cases, a word that seems to have no obvious synonym can be tied into the semantic structure through its superordinate. *Blunderbuss*, for example, has no good synonym in English, but it can be integrated into WordNet as a hyponym of *firearm*. In other instances, an initial

synonym set can be reorganized; coordinate terms — names of trees, for example — that were entered initially as synonyms of *tree*, could, with the introduction of hyponymic relations, be entered more accurately as hyponyms — in this example, as hyponyms of *tree*. In general, the addition of hyponymy has had the effect of sharpening the semantic distinctions that can be drawn and, as a consequence, reducing the average size of the synonym sets. Considerable work is sometimes required to reach a satisfactory solution. In those cases care has been taken not to impose more order on WordNet than a literate speaker of English might find reasonably obvious.

Antonymic Clusters

Psychologists also have an interest in antonymy, since antonyms are so often used to anchor the ends of scales used in subjective judgments: *good-bad, agree-disagree, right-wrong*, and so on. Probably the most extensive use of antonyms for scaling purposes was Osgood, Suci, & Tannenbaum's (1957) attempt to map all concepts into a space whose coordinates were given by pairs of antonymous adjectives.

Not every word has an antonym, of course. This relation is probably clearest between adjectives, although it is by no means limited to adjectives. The adjectival synonym sets were chosen as the most appropriate place to introduce antonymy into WordNet.

The work began with the assumption that antonymy and synonymy are themselves opposites. That is to say, synonyms are words whose meanings are very similar, whereas antonyms are words whose meanings are very dissimilar. That assumption may suffice as long as one does not look too closely, but careful analysis reveals important differences. The long history of disagreement about the nature and definition of antonymy (Egan, 1984) should have been a warning, but the extent of the difference was not recognized until an attempt was made to represent antonymous pairs by symmetrical cross-references between contrasting synonym sets.

The design of WordNet landed it, inadvertently, in the middle of a traditional argument about antonymy. Is an antonym (1) any one of several words that can be

opposed to a group of synonymous terms, or is it (2) a single word, or at most one of two or three words, that can be opposed to a given word? As originally conceived, WordNet incorporated assumption (1). That is to say, relatively large groups of synonyms were first compiled; then attempts were made to cross-reference the antonymous sets. But it proved difficult to carry that program through. When synonym set C_i was put in opposition to synonym set C_j, not every word in C_i was an antonym of every word in C_j, and vice versa, and that fact made it difficult to judge whether the concepts represented by the synonym sets were truly antonymous.

For example, the concept that is represented by the synonym set {*damp, dank, drenched, moist, soaked, waterlogged, wet* } seems to be antonymous to the concept that is represented by the synonym set { *arid, baked, dehydrated, dessicated, parched, sere, withered*} but few people would think of *withered* as an antonym of *waterlogged*, say, or of *baked* as an antonym of *dank*. Assumption (1) defines antonymy as a relation between lexical concepts, whereas assumption (2) defines antonymy as a relation between words. Judgments of antonymy are much easier to make between words than between concepts.

The addition of antonymous relations sharpens considerably the semantic distinctions that are required. That is to say, the adoption of assumption (2) necessarily limits the number of words in many synonym sets to two or three. But the notion that antonymy is a relation between words rather than between concepts, finds support in the frequent use of morphology to signal antonymy: *perfect-imperfect, pleasant-unpleasant, powerful-powerless, superior-inferior, definite-indefinite*, etc., illustrate only a few of the ways in which derivational morphology serves this purpose. Or, to put it differently, prefixing *un-* to adjectives can result in new adjectives (*pleasant-unpleasant*) in much the same way that adding *en-* to adjectives can result in causative verbs (*rich-enrich*). In both cases, the affix does important semantic work, but both dyads reflect formal relations between pairs of words. This fact is consistent with assumption (2), which defines antonymy as a relation between words.

Moreover, if it is assumed that the morphological relations involved in particular antonymous pairs must be learned by repeated exposure and practice, much the way all formal (i.e., phonological and morphological) features of English

are learned, then other observations about antonyms could be explained. For example, although *big-little* and *large-small* are both antonymous pairs, it sounds odd to cross them: *big-small* and *large-little*. The explanation is that we have heard them paired one way much more frequently than the other. Although the cross is conceptually correct, it is morphologically unfamiliar.

How can a conceptual definition of synonymy coexist with a formal conception of antonymy? Or, in more practical terms, how can a loose definition of synonymy be combined with a strict definition of antonymy? Solving this practical problem forced an interesting structure onto the adjective file: antonym pairs must form the basic skeleton of adjectival semantics, and this skeleton is fleshed out by those adjectives that have no obvious antonyms, but are similar to adjectives that do have antonyms. That is to say, another relation, dubbed semantic similarity, is introduced to preserve sets of several synonyms, but without precluding the one:one pairing of antonyms.

Table 1
The antonymic cluster, *wet-dry*
(Antonymic relation, *; similarity relation, &)

1000a { wet, &1001a, &1002a, &1003a, *1005a, }

1001a { damp, dank, moist, &1000a, }

1002a { drenched, saturated, soaked, waterlogged, &1000a, }

1003a { foggy, humid, misty, rainy, &1000a, }

1004a { drunk, slopped, tipsy, wet, *1080a, }

1005a { dry, *1000a, &1006a, &1007a, &1008a, &1009a, &1014a, }

1006a { arid, &1005a, }

1007a { dehydrated, dessicated, sere, withered, &1005a, }

1008a { baked, parched, &1005a, }

1009a { thirsty, &1005a, }

1010a { dry, impassive, matter-of-fact, unemotional, *1020a, }

1011a { barren, dry, sterile, unproductive, *1030a, }

1012a { boring, dry, insipid, wearisome, &1040a, &1090a, }

1013a { bare, dry, plain, unadorned, *1060a, }

1014a { anhydrous, &1005a, }

1015a { dry, &1110a, *1070a }

———

1020a { emotional, *1010a, }

1030a { fruitful, productive, *1011a, }

1040a { dull, &1012a, &1090a, }

1050a { interesting, *1090a, }

1060a { adorned, fancy, *1013a, }

1070a { sweet, *1015a, &1100a, }

1080a { dry, sober, *1004a }

1090a { uninteresting, &1012a, &1040a, *1050a, }

1100a { sugary, *1110a, &1070a, }

1110a { sugarless, &1015a, *1100a, }

The result is illustrated in Table 1 by the cluster of concepts around the antonymous pair *wet-dry*. (The 'a' following each number indicates that it is the name for an adjectival synonym set.) If *dry* in 1005a is consulted in search of an antonym, *wet* will be found in 1000a (and vice versa), whereas if *dry* in 1015a is consulted, the antonym in 1070a will be *sweet*. On the other hand, if 1005a is consulted for near synonyms of *dry* , all the words in 1006a, 1007a, 1008a, 1009a, and 1014a will be found. Thus, a narrow interpretation of antonymy can coexist with a broad interpretation of synonymy. Moreover, this form of representation poses no special problems for polysemous words: the *dry* that is the antonym of *wet* expresses a different concept from the *dry* that is the antonym of *sweet*, and different also from the *dry* that is similar to *dull* and *uninteresting*.

Implicit in the adoption of this structure for WordNet is the hypothesis that native speakers of English have a similar organization of their lexical memory for antonyms. That hypothesis was explored in a series of experiments by Gross, Fischer, and Miller (1988). The first experiment asked native speakers of English to judge relations between different types of contrasting pairs of adjectives: direct antonyms, indirect antonyms, and unrelated adjectives. Direct antonyms are lexically opposed terms such as *wet* and *dry*. An indirect antonymic pair consists of an adjective and a near synonym of its direct antonym that does not have its own lexical antonym: *dry* and *dank*. Examples of unrelated adjectives are *pleasant* and *scarlet,* or *regretful* and *clumsy*. Native speakers of English were expected to judge direct antonymic pairs like *wet-dry* faster than indirect pairs like *wet-parched.* A second experiment asked subjects to distinguish direct antonyms from all other types of adjective pairs. The results of these experiments, although not as robust as might have been expected, were consistent with the hypothesis that semantic memory for adjectives is organized around bipolar attributes and that certain pairs (the direct antonyms) label the poles.

Such experiments serve to illustrate one way that WordNet contributes to our understanding of the organization of lexical memory. As future work incorporates meronymy (work in progress by D. Teibel and G. Miller), association, and verb classification (Fellbaum and Kegl, to appear), further facets of the organization of lexical memory may become apparent.

Meronymy

Meronymy, the part-whole relation, is another basic semantic relation between words and concepts. This relation turns out to play a prominent role in the noun component of the lexicon and is widely exploited in WordNet. Winston, Chaffin, and Herrmann (1988; also Chaffin, Herrmann, and Winston 1987) studied a wide variety of part-whole relations.

The most easily identifiable examples of meronymy are found among words denoting concrete and countable entities. Body parts, for example, lend themselves well to part-whole classification: a *finger* is a part of a *hand*, a *hand* is a part of an *arm*, and an *arm* is a part of a *body*.

Another kind of meronymy is represented by those cases where the concept of the whole exists only by virtue of the existence of a multiple of the parts and is conceptually and linguistically inseparable from them, as in the example *a tree is a part of a forest*. Thus, one can say *a forest is many trees* but not, for example, *a body is many arms*.

In the lexicon of nouns referring to substances, meronymy takes on still another slightly different meaning. As Lyons (1977) points out, *gold* is both a substance and can also be a part of a compound matter. Thus, we can say both *this substance is gold* and *gold is part of this substance*. But the same does not hold for *arm*: Although we can say *The finger is part of an arm*, we cannot say *This arm is a finger.*

Meronymy overlaps with hyponymy in the case of collective nouns such as *furniture*: While *table* is a kind of *furniture*, it is also part of *furniture*, in the sense that the concept *furniture* can be said to prototypically include the concept *table*. The classification of such collectives can, therefore, be problematic.

In the realm of concrete and count nouns, meronymy permits the establishment of hierarchical structures in parallel with, but distinct from, hyponymic structures. Meronymic relations, like hyponymic relations, are also transitive, in that we can say that if *x* is a part of *y*, and *y* is a part of *z*, then *x* is also part of *z*. For example, *foot* is a meronym of *leg* and *leg* is a meronym *body;* therefore, *foot* is a meronym *body*. Bierwisch (1965) discusses redundancies in

these meronymic structures and asks to what extent they should be eliminated by rules. It would be interesting to test whether and how meronymic transitivity is represented in lexical memory: for example, to see whether subjects will easily associate two words that are distantly related by meronymy such as *doorknob* and *house*, and whether such associations require more time than those between less distantly related words like *door* and *house*.

Interesting relations exist between the hyponymic hierarchy and the meronymic hierarchies. For example, it is not necessary to say that *deck* is a meronym of *warship* if it has already been said that *deck* is a meronym of the superordinate *ship*. Tversky & Hemenway (1984) argue that the appropriate level in the hyponymic hierarchy for entering part-whole relations is the level that has been called "basic" by anthropological linguists (Berlin, Breedlove, & Raven, 1966; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976).

Hyponymy, antonymy, and meronymy reflect different aspects of the organization of human lexical memory and they all differ from synonymy. Consequently, the four relations must be represented differently in WordNet. Not until experience had been gained with this task, however, was the extent of their differences and interrelations appreciated. In the final section of this paper, we shall discuss the role of these relations in the verb lexicon, which presents a great challenge to any lexicographer.

Semantic Relations in the Verb Lexicon

At present, over 3,000 synonym sets of verbs have been compiled. They were initially classified into fifteen groups along the lines suggested in Miller and Johnson-Laird (1976). This classification follows very general but intuitively basic semantic criteria; thus, we have verbs of possession, communication, mental state and activity, motion, contact, change, competition, consumption, bodily functions, creation, psychological verbs, existence, social activities, perception, and natural events. The semantic relations of hyponymy, antonymy, and meronymy, that serve naturally to relate nouns and adjectives turn out to be less fitting for verbs.

Superficially, verbs do not seem to be easily represented by a hyponymic taxonomy. Rather than functioning as true hyponyms of a superordinate term, clusters of verbs seem to be related to a core or genus verb via a relation that often specifies the manner in which the subordinate is related to the superordinate. Thus, rather than bearing an ISA relation, a verb's relation to its genus term is expressable by means of a formula such as *to V1* is *to V2 in some way*. For example, *to skulk* is *to walk in a stealthy manner*; *to sew* is *to make by drawing together with a needle and thread.*

Among the nouns that have been entered into WordNet so far, the hyponymic relation has tended to be inherited fairly regularly, so that a hierarchical tree could easily be constructed. However, the corresponding relation among verbs behaves quite differently with respect to inheritance. For example, while *walking* is a kind of *traveling*, and *prowling* and *skulking* are kinds of *walking*, it is rather odd to state that *prowling* and *skulking* are kinds of *traveling*. On the other hand, the hyponymy relation that exists between *walking* and *traveling* and, for example, *marching* and *walking* is inherited, so that we can say that *marching* is a kind of *traveling*. These observations suggest that both *walking* and verbs referring to kinds of *walking* have two principal semantic components, one of displacement (*traveling*) and one of manner. The relation to the hypernym *travel* is only inherited in those verbs where displacement constitutes the salient semantic component. This indicates that the inheritance relations among verbs are less straightforward than they are among nouns.

Meronymy, which was found to play a significant role as a semantic relation among nouns, is not found in the same way among verbs. Its counterpart in the verb lexicon is a hierarchy-building relation that may be called semantic inclusion. Semantic inclusion is related to the logical notion of entailment. Thus, under the literal interpretation of *dream*, (1) entails (2), because when (1) holds, then (2) also holds:

(1)   John is dreaming
(2)   John is sleeping

Note that this relation is different from the one just discussed: *dreaming* is not a kind of *sleeping*. While the kind-of relation always includes the entailment

relation (i.e., you cannot do *x* in a certain way without doing *x* itself), the converse is not true, as examples (1) and (2) show; hyponymy and inclusion are distinct and asymmetric relations. Thus, both the ISA (kind of) and the inclusion relations build hierarchical, partially overlapping structures in the verb lexicon.

Hyponymy in the verb lexicon is complicated. There are two ways of thinking about the architecture of verb relations: (1) in terms of function; and (2) in terms of taxonomies.

Using functional relations such as *in the manner of*, *by means of*, and so forth, yields a nonhyponymic organization that relates a set of verbs to a base verb by mapping the base verb into a set of functions, yielding more specified verbs. In the following examples the base word is related to its more specified relatives by the functions *by means of*, *with a substance*, and *in the manner of*, respectively:

attach   =>   {*nail, tape, paste, glue, pin, sew, button, hook,...*}
cover   =>   {*butter, grease, oil, wax, paint, plaster,...*}
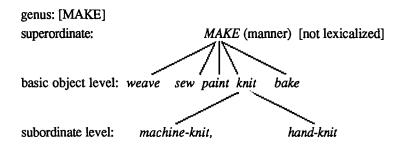*cook*   =>   {*broil, fry, roast, boil, bake, steam,...*}

Notice that the more specified verbs listed in these examples are not merely types of the base verb. The function that relates them to the base verb is additive.

The taxonomic perspective recognizes hyponymic relations among sets of verbs, analogous with the superordinate, basic, and subordinate levels that have been identified in the noun lexicon. A true superordinate of a basic level verb must contain the information contained by the relational functions. Thus, the superordinate of *to nail* is a verb meaning *to attach by means of some instrument*; the superordinate of *to butter* is *to cover with some substance*; and the superordinate of *to broil* is *to cook in some manner*. In contrast to noun taxonomies in English, the architecture of verb taxonomies is confounded by lexical gaps at the superordinate level.[1]
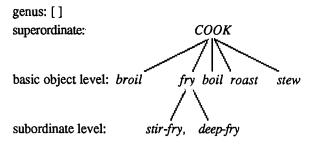
---

1. This may be a language-specific property since languages such as Malayalam and American Sign Language (ASL) have clear lexical gaps in the superordinate level of their taxonomies that are beginning to be instantiated by a word formation process that forms coordinate compounds by stringing together

Consider the taxonomic organization of two standardly recognized verb classes: the CREATION class and the CHANGE-OF-STATE class (see Atkins, Kegl, and Levin, 1988, for a discussion of the semantic and syntactic evidence for putting *bake* into both the CREATION class and the CHANGE-OF-STATE class.)

CREATION CLASS

genus: [MAKE]
superordinate: *MAKE* (manner) [not lexicalized]

basic object level: *weave   sew paint knit   bake*

subordinate level: *machine-knit,   hand-knit*

CHANGE-OF-STATE CLASS

genus: [ ]
superordinate: *COOK*

basic object level: *broil   fry boil roast   stew*

subordinate level: *stir-fry,   deep-fry*

Notice that the verb classes contain both a genus term and a superordinate. The genus term characterizes the class as a whole, picking out a set of sub-classes all of which share certain semantic and syntactic properties in common. Even though the previous schemata do not reflect it, both the CHANGE-OF-STATE and the CREATION classes contain a number of superordinates:

---

prototypical basic level words.   For example, in ASL, the superordinate corresponding to *fruit* is APPLE^ORANGE^BANANA^ETC (Klima and Bellugi, 1979).

CHANGE-OF-STATE:

> *cook* [basic: {*broil, fry, roast, bake, boil,...*}]
> *change (integrity)* [basic: {*break, smash, crack, shatter,...*}]
> *change (shape)* [basic: {*bend, buckle,...*}]
> *change (size)* [basic: {*grow, shrink, expand,...*}]
> *change (consistency)* [basic: {*melt, liquefy, freeze,...*}]
> etc.

CREATION:

> *create (by mental act)* [basic: {*contrive, invent, fabricate,...*}]
> *create (from raw material)* [basic: {*sew, bake, knit, paint,...*}]
> *create (by removal)* [basic: {*dig, drill, tear, bite, break,...*}]

The broad classification of these verbs into CHANGE-OF-STATE and CREATION classes is motivated by certain characteristics that members of each class share in common. For example, all the CHANGE-OF-STATE verbs allow the middle construction, whereas the CREATION verbs do not.

(3)   This glass shatters only under heavy pressure
(4)   This cardboard crushes easily
(5)   Cotton material shrinks when washed hot
(6)   M&Ms melt in your mouth, not in your hand

(7)   *Summer dresses sew easily
(8)   *In Antarctica, wells dig with great difficulty
(9)   *Theories invent with little effort
(10)   *Holes tear easily in this fabric

The CHANGE-OF-STATE class characteristically allows an inchoative form and the CREATION class does not. Contrast the following sets of examples:

CHANGE-OF-STATE:

    (11)   The water boiled
    (12)   The sweater shrank
    (13)   The bottle broke

CREATION:

    (14)   *The dress sewed
    (15)   *The sweater knitted
    (16)   *The portrait was painting

A characteristic of CREATION verbs that is not shared by the CHANGE-OF-STATE verbs involves the cooccurrence of a PP adjunct that indicates a source material:

    (17)   Mary fabricated the story from rumors
    (18)   Elaine sewed the dress from the finest of materials
    (19)   Albert bit a piece from the apple
    (20)   Mary baked a pie from the cherries she had picked

    (21)   *John smashed shards from the pot
    (22)   *Sue froze ice cubes from the water
    (23)   *Mary broke smithereens from the vase
    (24)   *John bent a cylinder from the cardboard

Numerous other tests for class membership are explored in Fellbaum and Kegl (to appear), and an elaborate discussion of verb classes appears in Levin (1985).

Notice that these two classes differ further at the basic level as well with regard to a transitivity alternation involving indefinite object deletion. However, the differences are reflected in only certain subclasses of verbs that can be distinguished at the superordinate level. These differences support a subclassification of verbs at the superordinate level.

One subclass of creation verbs — those whose superordinate is *create (from raw material)* — easily permits indefinite object deletion, whereas others do not.

(25)   John is knitting an afghan/John is knitting

(26)   Sam is painting a portrait/Sam is painting

(27)   Dick is fabricating lies/*Dick is fabricating

(28)   Mary tore a hole in her coat/*Mary tore

Some CREATION and CHANGE-OF-STATE verbs allow deletion, but this is definite object deletion and occurs only with a strong context as in (29) and (30). For discussions of indefinite vs. definite object deletion see Atkins, Kegl, and Levin (1988) and Fillmore (1986).

(29)   a. Elaine is roasting a goose/*Elaine is roasting [Elaine = agent]

         b. What does Jane do in the kitchen? Well, today she is roasting

(30)   a. The man invented a new kind of mousetrap/*The man is inventing

         b. Edison invented the phonograph/Don't bother Edison, he's inventing

Definite object deletion is not an inherent property of the CREATION class and the CHANGE-OF-STATE class per se, but rather is linked to the fact that members of the superordinate level of both these classes can function as activity verbs (like *eat, read, dance, clean*, etc.). In this activity verb realization the indefinite object can be omitted. Notice that at the subordinate level neither the CHANGE-OF-STATE nor CREATION classes allow indefinite object deletion. Allowing object deletion may be a property characteristic of a number of classes that are able to extend to become members of the activity class and to share syntactic and aspectual properties of that class (see Vendler, 1967; Dowty, 1979; Mittwoch, 1982). Nevertheless, the ease with which a verb can extend to another verb class may be related to its most basic class membership.

The distinction between genus and superordinate manifests itself in several ways. First, the genus term is not truly part of the taxonomy in that it does not hold an ISA relation to verbs at the superordinate, basic, and subordinate levels. Second, the genus term does not participate in the syntactic alternations reflected by the superordinate and basic level verbs. Notice that this suggests that *cook* is a

lexicalized superordinate and that *make* and *change* are not. Consider the following examples:

(31)   John is cooking
(32)   *Mary is changing [Mary = agent]
(33)   *Mary is making

The verb *create*, although a likely candidate for a lexicalized superordinate in the CREATION class, actually patterns more like a basic level term. It does not mean "make in some manner". Rather it is specific and refers to more artistic and aesthetic acts. Compare the specificity and infrequency of *Don't bother me, I'm creating* with the more generic *Don't bother me, I'm cooking*.

Representing verb relations with taxonomies versus functions is a matter of choice. As far as we can determine, the two descriptions are equivalent. However, the philosophy behind the architecture of WordNet, which requires that words be represented in terms of their relations to other words, and not by definitions, functions, features, or frames, favors the use of taxonomies. A strict use of taxonomies will, in certain instances, reveal lexical holes, positions in the network where a hypernym is suggested by the overall structure but lacks a lexical instantiation.

Exploring the feasibility of using taxonomies to represent the organization in the verb lexicon is the kind of experiment WordNet encourages. The idea is to see whether certain models or organizations suggested by small fragments of the lexicon remain viable when we consider the lexicon as a whole. With respect to this investigation, we have not yet coded enough of the verb lexicon to know for sure whether hyponymy relations truly deserve a place in the verb lexicon.

So far as antonymy in the verb lexicon is concerned, we strongly suspect that for the most part it is a secondary semantic relation derived from adjectives (of manner, degree, or intensity) or from spatial relations, where it is primary. Whenever an antonymic relation cannot be imported from elsewhere in the lexicon, we might expect a verb pair to lack an antonymic relation.

Members of two verb synonym sets are antonymous if the manner relation by which they differ is antonymic. For example, *nibble* and *gorge* are antonyms

because they are related to *eat* by *little, slow* and by *much, fast*, respectively. Antonymy also shows up systematically among verbs denoting a change from one state to another where each state can be related to a quality (e.g.,*lighten* and *darken* are antonyms by virtue of the antonymic relation that holds between the two adjectives from which they are derived).

The construction of WordNet is based entirely on the conceptual relations that exist between the members of the three major categories noun, verb, and adjective. A psychological model of lexical memory, however, should also account for relations between words belonging to different categories. Therefore, besides an antonymy relation that is imported from other lexical categories, a unidirectional inclusion relation, and verb class hyponymy, WordNet recognizes an additional linking that assigns verbs to a particular semantic domain. For example, the verb *fleece* is linked to the noun *sheep*, and a polysemous verb like *beat* is readily disambiguated when associated with different semantic domains: culinary, musical, contact, competition, and so on.

A speaker's lexical knowledge also includes knowledge of a word's syntactic use. Verbs are special. They form the backbone of the sentence and link in interesting ways to members of all the categories discussed (nouns, adjectives, and other verbs), as well as to function words and whole clauses. Thus, an important part of the lexical entry of a verb is its argument structure and selectional restrictions, that is, the number and kind of nouns it occurs with in a sentence. In this respect we stray from the basic constraint that words be represented only in relation to other words and we include syntactic and semantic frames.

To include this kind of information in WordNet, each verb synonym set is matched with a frame specifying the semantic/syntactic restrictions (a combination of subcategorizations and selectional restrictions) of its members. Since WordNet is intended for use by linguistically unsophisticated users, the codings must be simple and straightforward, drawing upon lexical knowledge the user already possesses. The coding task also presents some interesting theoretical challenges. We do not know at this point how many frames will be needed to account for all the verbs on file, but we are attempting to keep the number small by giving only generic specifications: for example, NP_{human} V $NP_{nonhuman}$. On the other hand, it is hoped that the frames and their relations to the synonym sets can be

connected in some nonrandom fashion to the semantic relations among the verbs. Some of the semantic distinctions made in the relational structures of possession verbs, for example, can be shown to be reflected in a systematic way. The verbs relating to *have* occur in the frame NP $_{human}$ V NP $_{nonhuman}$ (e.g., *John owns a car.*).

The subordinates of *take* and *give* are additionally specified for a prepositional phrase with NP $_{human}$ and *from* and *to* , respectively. Moreover, the frames show the difference between those *give* subordinates that systematically participate in the dative alternation and those that do not (NP V NP NP *vs.* NP V NP to NP).

Conclusion

Significant semantic differences exist between the three major syntactic categories (noun, adjective, and verb). Words from the three categories enter into synonymy relations with other words, yet each category is strongly linked to one additional predominant relation and tends to resist systematic organization by means of other relations. Furthermore, existing models based on fragments of the lexicon need to be, and are being, examined on a large scale. The more we explore the potential implementation of these models, the more complex and rich the problem becomes. Being forced to not only consider those words that fit our theories and therefore come to mind easily, but also to account for the remainder of the lexicon as well, has yielded many insights we would never have stumbled upon otherwise. Our hope is that when considered as a whole, the massive contents of the lexicon will conspire to narrow down the possible descriptions to a more and more general, manageable, and empirically testable set. In addition, the existence in electronic form of a comprehensively coded lexicon should allow other researchers to extend and examine the architecture of the lexicon in new and interesting ways.

George Miller, Christiane Fellbaum, Judy Kegl, and Katherine Miller
Princeton University

Références

AMSLER, R. A. (1980) *The Structure of the Merriam-Webster Pocket Dictionary.* Ph.D. dissertation, The University of Texas at Austin.

AMSLER, R. A. (1981) A taxonomy for English nouns and verbs. In: *Proceedings of the 19th Annual Meeting of the Association for Computational Linguistics.*Stanford, California.133-138.

AMSLER, R. A., and J. S. White (1979) Development of a Computational Methodology for Deriving Natural Language Semantic Structures via Analysis of Machine-Readable Dictionaries. Final Report, NSF Project MCS77-01315. Linguistics Research Center, The University of Texas at Austin.

ANDERSON, J. R. (1976) *Language, Memory, and Thought.* Hillsdale, N.J.: Erlbaum.

ATKINS, B. T., J. Kegl, and B. Levin (1988) Anatomy of a verb entry: From linguistic theory to lexicographic practice. International Journal of Lexicography.1.2. (in press).

BERLIN, B., D. E. Breedlove, and P.H. Raven (1966) Folk taxonomies and biological classification. Science,154. 273-275.

BIERWISCH, M. (1965) Eine Hierarchie syntaktisch-semantischer Merkmale. Studia Grammatica 5.29-86.

CARPENTER, P. A., and M. A. Just (1983) What your eyes do while your mind is reading. In K. Rayner (ed.), *Eye Movements in Reading: Perceptual and Language Processes.* New York: Academic Press.

CHAFFIN, R., and D. J. Herrmann (1984) The similarity and diversity of semantic relations. Memory and Cognition.12.134-141.

CHAFFIN, R., D. J. Herrmann, and M. Winston (1987) An empirical taxonomy of part-whole relations: Effects of part-whole relation type on relation identification. (Unpublished manuscript.)

CHAFFIN, R., and L. Peirce (1987) A taxonomy of semantic relations for the classification of GRE analogy items and an algorithm for the generation of GRE-type analogies. (Unpublished manuscript.)

CHAPMAN, R. L. (1977) *Roget's International Thesaurus.* 4th ed. New York: Crowell.

CHODOROW, M. S., R. J. Byrd, G. E. Heidorn. (1985) Extracting Semantic Hierarchies from a Large On-Line Dictionary. Proceedings of the ACL. 299-304.

COLLINS, A. M., and M. R. Quillian (1969) Retrieval time from semantic memory. Journal of Verbal Learning and Verbal Behavior 8. 240-247.

CRUSE, D. A. (1986). *Lexical Semantics.* Cambridge: Cambridge University Press.

CULLINGFORD, R. E. (1986) *Natural Language Processing Knowledge-An Egineering Approach.* Totowa, N.J.: Rowman & Littlefield.

DOWTY, D. (1979) *Word Meaning and Montague Grammar.* Dordrecht: Reidel.

EVENS, M. W., B. E. Litowitz, J. A. Markowitz, R. N. Smith, and O. Werner (1983) *Lexical-Semantic Relations: A Comparative Survey.* Edmonton, Canada: Linguistic Research.

EGAN, R. F. (1984) Survey of the history of English synonymy. In Gove, P. B. (ed.): Webster's New Dictionary of Synonym. Springfield, Mass.: Merriam-Webster. 5a-31a.

FELLBAUM, C. and J. Kegl (to appear) Taxonomic Hierarchies in the Verb Lexicon. Proceedings of Budalex '88: EURALEX Third International Conference. Budapest, Hungary.

FILLMORE, C. (1986) Pragmatically Controlled Zero-Anaphora. Proceedings of the Berkeley Linguistics Society 12. 95-107.

GOLDMAN-EISLER, F. (1968) *Psycholinguistics.* London: Academic Press.

GORDON, B. (1985) Subjective frequency and the lexical decision latency function: Implications for mechanisms of lexical access. Journal of Memory and Language 24. 631-645.

GROSS, D., U. Fischer, and G. A. Miller (1988) Antonymy and the Representation of Adjectival Meanings. CSL Report 13. Cognitive Science Laboratory, Princeton University, Princeton, NJ.

IRIS, M. A., B. E. Litowitz, and M. W. Evens (1985) The part-whole relation in the lexicon: An investigation of semantic primitives. (Unpublished manuscript.)

JACKENDOFF, R. (1983) *Semantics and Cognition.* Cambridge, Mass.: MIT Press.

KATZ, J. J. (1972) *Semantic Theory.* New York: Harper & Row.

KEIL, F. C. (1979) *Semantic and Conceptual Development: An Ontological Perspective.* Cambridge, Mass.: Harvard University Press.

KLIMA, E. and Bellugi, U. (1979) *Signs of Language*. Cambridge, MA: Harvard University Press.

LEVIN, B. (1985) Lexical Semantics in Review: An Introduction. In B. Levin (ed.): *Lexical Semantics in Review*. Lexicon Project Working Papers 1. 1-62. Cambridge, Mass: MIT.

LYONS, J. (1977) *Semantics*, 2 vols. Cambridge: Cambridge University Press.

MILLER, G. A. (1985) WordNet: A dictionary browser. In: *Information in Data, Proceedings of the First Conference of the UW Centre for the New Oxford English Dictionary*. Waterloo, Canada: University of Waterloo. 25-28.

MILLER, G. A. and P. N. Johnson-Laird (1976) *Language and Perception*. Cambridge, Mass.: Harvard University Press.

MITTWOCH, A. (1982) On the Difference between *Eating* and *Eating Something* : Activities versus Accomplishments. LI 13.113-122.

NORMAN, D. A., and D. E. Rumelhart, (eds.). (1975) *Exploration in Cognition*. San Francisco: Freeman.

OSGOOD, C. E., G. J. Suci, and P. H. Tannenbaum (1957) *The Measurement of Meaning*. Urbana: University of Illinois Press.

PROCTER, P. (Ed.). (1978) *Longman Dictionary of Contemporary English*. London: Longman Group Limited.

RASKIN, R. (1987) Electronic thesauri: Four ways to find the perfect word. PC Magazine. 6.1. 275-283.

ROSCH, E., C. B. Mervis, W. Gray, D. Johnson, and P. Boyes-Braem (1976) Basic objects in natural categories. Cognitive Psychology. 8. 382-439.

SCHANK, R. C. (1972) Conceptual dependency: A theory of natural language understanding. Cognitive Psychology. 3. 552-631.

SEDELOW, S. Y., and W. A. Sedelow, Jr. (1986) Thesaural knowledge representation. In: *Advances in Lexicology, Proceedings of the Second Annual Conference of the UW Centre for the New Oxford English Dictionary*. Waterloo, Canada: University of Waterloo. 29-43.

SMITH, E. E. (1978) Theories of semantic memory. In W. K. Estes (ed.), *Handbook of Learning and Cognitive Processes*. Vol. 6: *Linguistic Functions and Cognitive Theory*. Hillsdale, N.J.: Erlbaum.1-56.

SOWA, J. F. (1984) *Conceptual Structures: Information Processing in Mind and Machine*. Reading, Mass.: Addison-Wesly.

TALMY, L. (1985) Lexicalization Patterns: Semantic Structure in Lexical Form. In T. Shopen (ed.) : *Language Typology and Syntactic Description*, vol.3 : *Grammatical Categories and the Lexicon*. Cambridge: Cambridge University Press.

TVERSKY, B., and K. Hemenway (1984) Objects, parts, and categories. Journal of Experimental Psychology: General, 2. 169-193.

VENDLER, Z. (1967) *Linguistics and Philosophy*. Ithaca, NY: Cornell University Press.

WHALEY, C. P. (1978) Word-nonword classification time. Journal of Verbal Learning and Verbal Behavior.17.143-154.

WINSTON, M.E., R. Chaffin, and D. J. Herrmann, D. J. (1988) A Taxonomy of Part-Whole Relations. Cognitive Science (in press).