Introduction to WordNet: An On-line Lexical Database

George A. Miller, Richard Beckwith, Christiane Fellbaum,
Derek Gross, and Katherine Miller

(Revised August 1993)

WordNet is an on-line lexical reference system whose design is inspired by current psycholinguistic theories of human lexical memory. English nouns, verbs, and adjectives are organized into synonym sets, each representing one underlying lexical concept. Different relations link the synonym sets.

Standard alphabetical procedures for organizing lexical information put together words that are spelled alike and scatter words with similar or related meanings haphazardly through the list. Unfortunately, there is no obvious alternative, no other simple way for lexicographers to keep track of what has been done or for readers to find the word they are looking for. But a frequent objection to this solution is that finding things on an alphabetical list can be tedious and time-consuming. Many people who would like to refer to a dictionary decide not to bother with it because finding the information would interrupt their work and break their train of thought.

In this age of computers, however, there is an answer to that complaint. One obvious reason to resort to on-line dictionaries—lexical databases that can be read by computers—is that computers can search such alphabetical lists much faster than people can. A dictionary entry can be available as soon as the target word is selected or typed into the keyboard. Moreover, since dictionaries are printed from tapes that are read by computers, it is a relatively simple matter to convert those tapes into the appropriate kind of lexical database. Putting conventional dictionaries on line seems a simple and natural marriage of the old and the new.

Once computers are enlisted in the service of dictionary users, however, it quickly becomes apparent that it is grossly inefficient to use these powerful machines as little more than rapid page-turners. The challenge is to think what further use to make of them. WordNet is a proposal for a more effective combination of traditional lexicographic information and modern high-speed computation.

This, and the accompanying four papers, is a detailed report of the state of WordNet as of 1990. In order to reduce unnecessary repetition, the papers are written to be read consecutively.

Psycholexicology

Murray's *Oxford English Dictionary* (1928) was compiled ‘‘on historical principles’’ and no one doubts the value of the *OED* in settling issues of word use or sense priority. By focusing on historical (diachronic) evidence, however, the *OED*, like other standard dictionaries, neglected questions concerning the synchronic organization of lexical knowledge.

It is now possible to envision ways in which that omission might be repaired. The 20th Century has seen the emergence of psycholinguistics, an interdisciplinary field of research concerned with the cognitive bases of linguistic competence. Both linguists and psycholinguists have explored in considerable depth the factors determining the contemporary (synchronic) structure of linguistic knowledge in general, and lexical knowledge in particular—Miller and Johnson-Laird (1976) have proposed that research concerned with the lexical component of language should be called psycholexicology. As linguistic theories evolved in recent decades, linguists became increasingly explicit about the information a lexicon must contain in order for the phonological, syntactic, and lexical components to work together in the everyday production and comprehension of linguistic messages, and those proposals have been incorporated into the work of psycholinguists. Beginning with word association studies at the turn of the century and continuing down to the sophisticated experimental tasks of the past twenty years, psycholinguists have discovered many synchronic properties of the mental lexicon that can be exploited in lexicography.

In 1985 a group of psychologists and linguists at Princeton University undertook to develop a lexical database along lines suggested by these investigations (Miller, 1985). The initial idea was to provide an aid to use in searching dictionaries conceptually, rather than merely alphabetically—it was to be used in close conjunction with an on-line dictionary of the conventional type. As the work proceeded, however, it demanded a more ambitious formulation of its own principles and goals. WordNet is the result. Inasmuch as it instantiates hypotheses based on results of psycholinguistic research, WordNet can be said to be a dictionary based on psycholinguistic principles.

How the leading psycholinguistic theories should be exploited for this project was not always obvious. Unfortunately, most research of interest for psycholexicology has dealt with relatively small samples of the English lexicon, often concentrating on nouns at the expense of other parts of speech. All too often, an interesting hypothesis is put forward, fifty or a hundred words illustrating it are considered, and extension to the rest of the lexicon is left as an exercise for the reader. One motive for developing WordNet was to expose such hypotheses to the full range of the common vocabulary. WordNet presently contains approximately 95,600 different word forms (51,500 simple words and 44,100 collocations) organized into some 70,100 word meanings, or sets of synonyms, and only the most robust hypotheses have survived.

The most obvious difference between WordNet and a standard dictionary is that WordNet divides the lexicon into five categories: nouns, verbs, adjectives, adverbs, and function words. Actually, WordNet contains only nouns, verbs, adjectives, and adverbs.[1] The relatively small set of English function words is omitted on the assumption (supported by observations of the speech of aphasic patients: Garrett, 1982) that they are probably stored separately as part of the syntactic component of language. The realization that syntactic categories differ in subjective organization emerged first from studies of word associations. Fillenbaum and Jones (1965), for example, asked English-

[1] A discussion of adverbs is not included in the present collection of papers.

speaking subjects to give the first word they thought of in response to highly familiar words drawn from different syntactic categories. The modal response category was the same as the category of the probe word: noun probes elicited nouns responses 79% of the time, adjectives elicited adjectives 65% of the time, and verbs elicited verbs 43% of the time. Since grammatical speech requires a speaker to know (at least implicitly) the syntactic privileges of different words, it is not surprising that such information would be readily available. How it is learned, however, is more of a puzzle: it is rare in connected discourse for adjacent words to be from the same syntactic category, so Fillenbaum and Jones's data cannot be explained as association by continguity.

The price of imposing this syntactic categorization on WordNet is a certain amount of redundancy that conventional dictionaries avoid—words like *back*, for example, turn up in more than one category. But the advantage is that fundamental differences in the semantic organization of these syntactic categories can be clearly seen and systematically exploited. As will become clear from the papers following this one, nouns are organized in lexical memory as topical hierarchies, verbs are organized by a variety of entailment relations, and adjectives and adverbs are organized as N-dimensional hyperspaces. Each of these lexical structures reflects a different way of categorizing experience; attempts to impose a single organizing principle on all syntactic categories would badly misrepresent the psychological complexity of lexical knowledge.

The most ambitious feature of WordNet, however, is its attempt to organize lexical information in terms of word meanings, rather than word forms. In that respect, WordNet resembles a thesaurus more than a dictionary, and, in fact, Laurence Urdang's revision of Rodale's *The Synonym Finder* (1978) and Robert L. Chapman's revision of *Roget's International Thesaurus* (1977) have been helpful tools in putting WordNet together. But neither of those excellent works is well suited to the printed form. The problem with an alphabetical thesaurus is redundant entries: if word W_x and word W_y are synonyms, the pair should be entered twice, once alphabetized under W_x and again alphabetized under W_y. The problem with a topical thesaurus is that two look-ups are required, first on an alphabetical list and again in the thesaurus proper, thus doubling a user's search time. These are, of course, precisely the kinds of mechanical chores that a computer can perform rapidly and efficiently.

WordNet is not merely an on-line thesaurus, however. In order to appreciate what more has been attempted in WordNet, it is necessary to understand its basic design (Miller and Fellbaum, 1991).

The Lexical Matrix

Lexical semantics begins with a recognition that a word is a conventional association between a lexicalized concept and an utterance that plays a syntactic role. This definition of ''word'' raises at least three classes of problems for research. First, what kinds of utterances enter into these lexical associations? Second, what is the nature and organization of the lexicalized concepts that words can express? Third, what syntactic roles do different words play? Although it is impossible to ignore any of these questions while considering only one, the emphasis here will be on the second class of problems, those dealing with the semantic structure of the English lexicon.

Since the word ‘‘word’’ is commonly used to refer both to the utterance and to its associated concept, discussions of this lexical association are vulnerable to terminological confusion. In order to reduce ambiguity, therefore, ‘‘word form’’ will be used here to refer to the physical utterance or inscription and ‘‘word meaning’’ to refer to the lexicalized concept that a form can be used to express. Then the starting point for lexical semantics can be said to be the mapping between forms and meanings (Miller, 1986). A conservative initial assumption is that different syntactic categories of words may have different kinds of mappings.

Table 1 is offered simply to make the notion of a lexical matrix concrete. Word forms are imagined to be listed as headings for the columns; word meanings as headings for the rows. An entry in a cell of the matrix implies that the form in that column can be used (in an appropriate context) to express the meaning in that row. Thus, entry $E_{1,1}$ implies that word form F_1 can be used to express word meaning M_1. If there are two entries in the same column, the word form is polysemous; if there are two entries in the same row, the two word forms are synonyms (relative to a context).

Table 1

Illustrating the Concept of a Lexical Matrix:

F_1 and F_2 are synonyms; F_2 is polysemous

Word Meanings	Word Forms F_1	F_2	F_3	. . .	F_n
M_1	$E_{1,1}$	$E_{1,2}$			
M_2		$E_{2,2}$			
M_3			$E_{3,3}$		
⋮				⋱	
M_m					$E_{m,n}$

Mappings between forms and meanings are many:many—some forms have several different meanings, and some meanings can be expressed by several different forms. Two difficult problems of lexicography, polysemy and synonymy, can be viewed as complementary aspects of this mapping. That is to say, polysemy and synonymy are problems that arise in the course of gaining access to information in the mental lexicon: a listener or reader who recognizes a form must cope with its polysemy; a speaker or writer who hopes to express a meaning must decide between synonyms.

As a parenthetical comment, it should be noted that psycholinguists frequently represent their hypotheses about language processing by box-and-arrow diagrams. In that notation, a lexical matrix could be represented by two boxes with arrows going between them in both directions. One box would be labeled ‘Word Meaning’ and the other ‘Word Form’; arrows would indicate that a language user could start with a meaning and look for appropriate forms to express it, or could start with a form and retrieve appropriate meanings. This box-and-arrow representation makes clear the difference between meaning:meaning relations (in the Word Meaning box) and word:word relations (in the Word Form box). In its initial conception, WordNet was concerned solely with the pattern of semantic relations between lexicalized concepts; that is to say, it was to be a theory of the Word Meaning box. As work proceeded, however, it became increasingly clear that lexical relations in the Word Form box could not be ignored. At present, WordNet distinguishes between semantic relations and lexical relations; the emphasis is still on semantic relations between meanings, but relations between words are also included.

Although the box-and-arrow representation respects the difference between these two kinds of relations, it has the disadvantage that the intricate details of the many:many mapping between meanings and forms are slighted, which not only conceals the reciprocity of polysemy and synonymy, but also obscures the major device used in WordNet to represent meanings. For that reason, this description of WordNet has been introduced in terms of a lexical matrix, rather than as a box-and-arrow diagram.

How are word meanings represented in WordNet? In order to simulate a lexical matrix it is necessary to have some way to represent both forms and meanings in a computer. Inscriptions can provide a reasonably satisfactory solution for the forms, but how meanings should be represented poses a critical question for any theory of lexical semantics. Lacking an adequate psychological theory, methods developed by lexicographers can provide an interim solution: definitions can play the same role in a simulation that meanings play in the mind of a language user.

How lexicalized concepts are to be represented by definitions in a theory of lexical semantics depends on whether the theory is intended to be constructive or merely differential. In a constructive theory, the representation should contain sufficient information to support an accurate construction of the concept (by either a person or a machine). The requirements of a constructive theory are not easily met, and there is some reason to believe that the definitions found in most standard dictionaries do not meet them (Gross, Kegl, Gildea, and Miller, 1989; Miller and Gildea, 1987). In a differential theory, on the other hand, meanings can be represented by any symbols that enable a theorist to distinguish among them. The requirements for a differential theory are more modest, yet suffice for the construction of the desired mappings. If the person who reads the definition has already acquired the concept and needs merely to identify it, then a synonym (or near synonym) is often sufficient. In other words, the word meaning M_1 in Table 1 can be represented by simply listing the word forms that can be used to express it: $\{F_1, F_2, \ldots\}$. (Here and later, the curly brackets, '{' and '},' surround the sets of synonyms that serve as identifying definitions of lexicalized concepts.) For example, someone who knows that *board* can signify either a piece of lumber or a group of people assembled for some purpose will be able to pick out the intended sense with no more help than *plank* or *committee*. The synonym sets, {*board, plank*} and {*board, committee*} can serve as unambiguous designators of these two meanings of *board*. These synonym sets (synsets) do not explain what the concepts are; they merely signify that the concepts exist. People who know English are assumed to have already acquired the concepts, and are expected to recognize them from the words listed in the synset.

A lexical matrix, therefore, can be represented for theoretical purposes by a mapping between written words and synsets. Since English is rich in synonyms, synsets are often sufficient for differential purposes. Sometimes, however, an appropriate synonym is not available, in which case the polysemy can be resolved by a short gloss, e.g., {*board*, (a person's meals, provided regularly for money)} can serve to differentiate this sense of *board* from the others; it can be regarded as a synset with a single member. The gloss is not intended for use in constructing a new lexical concept by someone not already familiar with it, and it differs from a synonym in that it is not used to gain access to information stored in the mental lexicon. It fulfills its purpose if it enables the user of WordNet, who is assumed to know English, to differentiate this sense from others with which it could be confused.

Synonymy is, of course, a lexical relation between word forms, but because it is assigned this central role in WordNet, a notational distinction is made between words related by synonymy, which are enclosed in curly brackets, '{' and '}', and other lexical relations, which will be enclosed in square brackets, '[' and ']'. Semantic relations are indicated by pointers.

WordNet is organized by semantic relations. Since a semantic relation is a relation between meanings, and since meanings can be represented by synsets, it is natural to think of semantic relations as pointers between synsets. It is characteristic of semantic relations that they are reciprocated: if there is a semantic relation R between meaning $\{x, x', \ldots\}$ and meaning $\{y, y', \ldots\}$, then there is also a relation R′ between $\{y, y', \ldots\}$ and $\{x, x', \ldots\}$. For the purposes of the present discussion, the names of the semantic relations will serve a dual role: if the relation between the meanings $\{x, x', \ldots\}$ and $\{y, y', \ldots\}$ is called R, then R will also be used to designate the relation between individual word forms belonging to those synsets. It might be logically tidier to introduce separate terms for the relation between meanings and for the relation between forms, but even greater confusion might result from the introduction of so many new technical terms.

The following examples illustrate (but do not exhaust) the kinds of relations used to create WordNet.

Synonymy

From what has already been said, it should be obvious that the most important relation for WordNet is similarity of meaning, since the ability to judge that relation between word forms is a prerequisite for the representation of meanings in a lexical matrix. According to one definition (usually attributed to Leibniz) two expressions are synonymous if the substitution of one for the other never changes the truth value of a sentence in which the substitution is made. By that definition, true synonyms are rare, if they exist at all. A weakened version of this definition would make synonymy relative to a context: two expressions are synonymous in a linguistic context C if the substitution of one for the other in C does not alter the truth value. For example, the substitution of *plank* for *board* will seldom alter truth values in carpentry contexts, although there are other contexts of *board* where that substitution would be totally inappropriate.

Note that the definition of synonymy in terms of substitutability makes it necessary to partition WordNet into nouns, verbs, adjectives, and adverbs. That is to say, if concepts are represented by synsets, and if synonyms must be interchangeable, then words in different syntactic categories cannot be synonyms (cannot form synsets) because they are not interchangeable. Nouns express nominal concepts, verbs express verbal concepts, and modifiers provide ways to qualify those concepts. In other words, the use of synsets to represent word meanings is consistent with psycholinguistic evidence that nouns, verbs, and modifiers are organized independently in semantic memory. An argument might be made in favor of still further partitions: some words in the same syntactic category (particularly verbs) express very similar concepts, yet cannot be interchanged without making the sentence ungrammatical.

The definition of synonymy in terms of truth values seems to make synonymy a discrete matter: two words either are synonyms or they are not. But as some philosophers have argued, and most psychologists accept without considering the alternative, synonymy is best thought of as one end of a continuum along which similarity of meaning can be graded. It is probably the case that semantically similar words can be interchanged in more contexts than can semantically dissimilar words. But the important point here is that theories of lexical semantics do not depend on truth-functional conceptions of meaning; semantic similarity is sufficient. It is convenient to assume that the relation is symmetric: if *x* is similar to *y*, then *y* is equally similar to *x*.

The gradability of semantic similarity is ubiquitous, but it is most important for understanding the organization of adjectival and adverbial meanings.

Antonymy

Another familiar relation is antonymy, which turns out to be surprisingly difficult to define. The antonym of a word *x* is sometimes *not-x*, but not always. For example, *rich* and *poor* are antonyms, but to say that someone is not rich does not imply that they must be poor; many people consider themselves neither rich nor poor. Antonymy, which seems to be a simple symmetric relation, is actually quite complex, yet speakers of English have little difficulty recognizing antonyms when they see them.

Antonymy is a lexical relation between word forms, not a semantic relation between word meanings. For example, the meanings {*rise, ascend*} and {*fall, descend*} may be conceptual opposites, but they are not antonyms; [*rise/fall*] are antonyms and so are [*ascend/descend*], but most people hesitate and look thoughtful when asked if *rise* and *descend*, or *ascend* and *fall*, are antonyms. Such facts make apparent the need to distinguish between semantic relations between word forms and semantic relations between word meanings. Antonymy provides a central organizing principle for the adjectives and adverbs in WordNet, and the complications that arise from the fact that antonymy is a semantic relation between words are better discussed in that context.

Hyponymy

Unlike synonymy and antonymy, which are lexical relations between word forms, hyponymy/hypernymy is a semantic relation between word meanings: e.g., {*maple*} is a hyponym of {*tree*}, and {*tree*} is a hyponym of {*plant*}. Much attention has been devoted to hyponymy/hypernymy (variously called subordination/superordination, subset/superset, or the ISA relation). A concept represented by the synset {*x*, *x*′, . . .} is said to be a hyponym of the concept represented by the synset {*y*, *y*′, . . .} if native speakers of English accept sentences constructed from such frames as *An x is a (kind of) y*. The relation can be represented by including in {*x*, *x*′, . . .} a pointer to its superordinate, and including in {*y*, *y*′, . . .} pointers to its hyponyms.

Hyponymy is transitive and asymmetrical (Lyons, 1977, vol. 1), and, since there is normally a single superordinate, it generates a hierarchical semantic structure, in which a hyponym is said to be below its superordinate. Such hierarchical representations are widely used in the construction of information retrieval systems, where they are called inheritance systems (Touretzky, 1986): a hyponym inherits all the features of the more generic concept and adds at least one feature that distinguishes it from its superordinate and from any other hyponyms of that superordinate. For example, *maple* inherits the features of its superordinate, *tree*, but is distinguished from other trees by the hardness of its wood, the shape of its leaves, the use of its sap for syrup, etc. This convention provides the central organizing principle for the nouns in WordNet.

Meronymy

Synonymy, antonymy, and hyponymy are familiar relations. They apply widely throughout the lexicon and people do not need special training in linguistics in order to appreciate them. Another relation sharing these advantages—a semantic relation—is the part-whole (or HASA) relation, known to lexical semanticists as meronymy/holonymy. A concept represented by the synset {*x*, *x*′, . . .} is a meronym of a concept represented by the synset {*y*, *y*′, . . .} if native speakers of English accept sentences constructed from such frames as *A y has an x (as a part)* or *An x is a part of y*. The meronymic relation is transitive (with qualifications) and asymmetrical (Cruse, 1986), and can be used to construct a part hierarchy (with some reservations, since a meronym can have many holonyms). It will be assumed that the concept of a part of a whole can be a part of a concept of the whole, although it is recognized that the implications of this assumption deserve more discussion than they will receive here.

These and other similar relations serve to organize the mental lexicon. They can be represented in WordNet by parenthetical groupings or by pointers (labeled arcs) from one synset to another. These relations represent associations that form a complex network; knowing where a word is situated in that network is an important part of knowing the word's meaning. It is not profitable to discuss these relations in the abstract, however, because they play different roles in organizing the lexical knowledge associated with different syntactic categories.

Morphological Relations

An important class of lexical relations are the morphological relations between word forms. Initially, interest was limited to semantic relations; no plans were made to include morphological relations in WordNet. As work progressed, however, it became increasingly obvious that if WordNet was to be of any practical use to anyone, it would have to deal with inflectional morphology. For example, if someone put the computer's cursor on the word *trees* and clicked a request for information, WordNet should not reply that the word was not in the database. A program was needed to strip off the plural suffix and then to look up *tree*, which certainly is in the database. This need led to the development of a program for dealing with inflectional morphology.

Although the inflectional morphology of English is relatively simple, writing a computer program to deal with it proved to be a more complex task than had been expected. Verbs are the major problem, of course, since there are four forms and many irregular verbs. But the software has been written and is presently available as part of the interface between the lexical database and the user. In the course of this development it became obvious that programs dealing with derivational morphology would greatly enhance the value of WordNet, but that more ambitious project has not yet been undertaken.

The three papers following this introduction have little to say about lexical relations resulting from inflectional morphology, since those relations are incorporated in the interface to WordNet, not in the central database.

Nouns in WordNet: A Lexical Inheritance System

George A. Miller

(Revised August 1993)

Definitions of common nouns typically give a superordinate term plus distinguishing features; that information provides the basis for organizing noun files in WordNet. The superordinate relation (hyponymy) generates a hierarchical semantic organization that is duplicated in the noun files by the use of labeled pointers between sets of synonyms (synsets). The hierarchy is limited in depth, seldom exceeding more than a dozen levels. Distinguishing features are entered in such a way as to create a lexical inheritance system, a system in which each word inherits the distinguishing features of all its superordinates. Three types of distinguishing features are discussed: attributes (modification), parts (meronymy), and functions (predication), but only meronymy is presently implemented in the noun files. Antonymy is also found between nouns, but it is not a fundamental organizing principle for nouns.

Coverage is partitioned into twenty-five topical files, each of which deals with a different primitive semantic component.

As this is written, WordNet contains approximately 57,000 noun word forms organized into approximately 48,800 word meanings (synsets). The numbers are approximate because WordNet continues to grow—one advantage of an on-line database. Many of these nouns are compounds, of course; a few are artificial collocations invented for the convenience of categorization. No attempt has been made to include proper nouns; on the other hand, since many common nouns once were names, no serious attempt has been made to exclude them. In terms of coverage, WordNet's goals differ little from those of a good standard handheld collegiate-level dictionary. It is in the organization of that information that WordNet aspires to innovation.

If someone asks how to use a conventional dictionary, it is customary to explain the different kinds of information packed into lexical entries: spelling, pronunciation, inflected and derivative forms, etymology, part of speech, definitions and illustrative uses of alternative senses, synonyms and antonyms, special usage notes, occasional line drawings or plates—a good dictionary is a remarkable store of information. But if someone asks how to improve a dictionary, it becomes necessary to consider what is not included. And when, as in the case of WordNet, improvements are intended to reflect psycholinguistic principles, the focal concern becomes what is not included in the definitions.

Examples offer the simplest way to characterize the omissions. Take one meaning of the noun *tree*, the sense having to do with trees as plants. Conventional dictionaries define this sense of *tree* by some such gloss as: *a plant that is large, woody, perennial, and has a distinct trunk*. Of course, the actual wording is usually more felicitous—*a large, woody, perennial plant with a distinct trunk*, for example—but the underlying logic is the same: superordinate plus distinguishers. The point is that the prototypical definition of a noun consists of its immediate superordinate (*plant*, in this example), followed by a relative clause that describes how this instance differs from all other instances.

What is missing from this definition? Anyone educated to expect this kind of thing in a dictionary will not feel that anything is missing. But the definition is woefully incomplete. It does not say, for example, that trees have roots, or that they consist of cells having cellulose walls, or even that they are living organisms. Of course, if you look up the superordinate term, *plant*, you may find that kind of information—unless, of course, you make a mistake and choose the definition of *plant* that says it is a place where some product is manufactured. There is, after all, nothing in the definition of *tree* that specifies which sense of *plant* is the appropriate superordinate. That specification is omitted on the assumption that the reader is not an idiot, a Martian, or a computer. But it is instructive to note that, even though intelligent readers can supply it for themselves, important information about the superordinate term is missing from the definition.

Second, this definition of *tree* contains no information about coordinate terms. The existence of other kinds of plants is a plausible conjecture, but no help is given in finding them. A reader curious about coordinate terms has little alternative but to scan the dictionary from A to Z, noting along the way each occurrence of a definition with the superordinate term *plant*. Even this heroic strategy might not succeed if the lexicographers, not expecting such use of their work, did not maintain strict uniformity in their choice of superordinate terms. *Tree* is probably an unfair example in this respect, since the distinction between trees and bushes is so unclear—the same plant that grows into a tall tree in one location may be little more than a bush in a less favorable climate. Botanists have little use for the lay term *tree*—many trees are gymnosperms, many others angiosperms. Even for well-behaved definitions, however, a conventional dictionary leaves the discovery of coordinate terms as a challenging exercise for the reader.

Third, a similar challenge faces a reader who is interested in knowing the different kinds of trees. In addition to looking through the dictionary for such familiar trees as pine or maple or oak, a reader might wish to know which trees are deciduous, which are hardwoods, or how many different kinds of conifers there are. Dictionaries contain much of this information, but only the most determined reader would try to dig it out. The prototypical definition points upward, to a superordinate term, not sideways to coordinate terms or downward to hyponyms.

Fourth, everyone knows a great deal about trees that lexicographers would not include in a definition of *tree*. For example, trees have bark and twigs, they grow from seeds, adult trees are much taller than human beings, they manufacture their own food by photosynthesis, they provide shade and protection from the wind, they grow wild in forests, their wood is used in construction and for fuel, and so on. Someone who was totally innocent about trees would not be able to construct an accurate concept of them if nothing more were available than the information required to define *tree*. A dictionary definition draws some important distinctions and serves to remind the reader of something that is presumed to be familiar already; it is not intended as a catalogue of general knowledge. There is a place for encyclopedias as well as dictionaries.

Note that much of the missing information is structural, rather than factual. That is to say, lexicographers make an effort to cover all of the factual information about the meanings of each word, but the organization of the conventional dictionary into discrete, alphabetized entries and the economic pressure to minimize redundancy make the reassembly of this scattered information a formidable chore.

Lexical Inheritance Systems

It has often been observed that lexicographers are caught in a web of words. Sometimes it is posed as a conundrum: since words are used to define words, how can lexicography escape circularity? Every dictionary probably contains a few vacuous circles, instances where word W_a is used to define word W_b and W_b is also used to define W_a; in such cases, presumably, the lexicographer inadvertently overlooked the need to define one or the other of these synonyms in terms of something else. Circularity is the exception, not the rule.

The fundamental design that lexicographers try to impose on the semantic memory for nouns is not a circle, but a tree (in the sense of *tree* as a graphical representation). It is a defining property of tree graphs that they branch from a single stem without forming circular loops. The lexical tree can be reconstructed by following trails of superordinate terms: *oak @→ tree @→ plant @→ organism,* for example, where '@→' is the transitive, asymmetric, semantic relation that can be read 'is a' or 'is a kind of.' (By convention, '@→' is said to point upward.) This design creates a sequence of levels, a hierarchy, going from many specific terms at the lower levels to a few generic terms at the top. Hierarchies provide conceptual skeletons for nouns; information about individual nouns is hung on this structure like ornaments on a Christmas tree.

The semantic relation that is represented above by '@→' has been called the ISA relation, or the hypernymic or superordinate relation (since it points to a hypernym or superordinate term); it goes from specific to generic and so is a generalization. Whenever it is the case that a noun W_h @→ a noun W_s, there is always an inverse relation, W_s ~→ W_h. That is to say, if W_s is the superordinate of W_h, then W_h is the subordinate or hyponym of W_s. The inverse semantic relation '~→' goes from generic to specific (from superordinate to hyponym) and so is a specialization.

Since a noun usually has a single superordinate, dictionaries include the superordinate in the definition; since a noun can have many hyponyms, English dictionaries do not list them (the French dictionary *Le Grand Robert* is an exception). Even though the specialization relation is not made explicit in standard dictionaries of English, it is a logical derivative of the generalization relation. In WordNet, lexicographers code the generalization relation '@→' explicitly with a labeled pointer between lexical concepts or senses. When the lexicographers' files are converted automatically into the lexical database, one step in this process is to insert inverse pointers for the specialization relation '~→'. Thus, the lexical database is a hierarchy that can be searched upward or downward with equal speed.

Hierarchies of this sort are widely used by computer programmers to organize large databases (Touretzky, 1986). They have the advantage that information common to many items in the database need not be stored with every item. In other words, database experts and lexicographers both resort to hierarchical structures for the same reason: to save space. Computer scientists call such hierarchies ‘‘inheritance systems,’’ because they think of specific items inheriting information from their generic superordinates. That is to say, all of the properties of the superordinate are assumed to be properties of the subordinate as well; instead of listing those properties redundantly with both items, they are listed only with the superordinate and a pointer from the subordinate to the superordinate is understood to mean ‘‘for additional properties, look here.’’

Inheritance is most easily understood for names. If you hear that your friend has acquired a collie named Rex, you do not need to ask whether Rex is an animal, whether Rex has hair, four legs, and a tail, or whether Rex shares any other properties known to characterize collies. Such questions would be distinctly odd. Since you have been told that Rex is a collie, you are expected to understand that Rex inherits all the properties that define *collie*. And, implicitly, that *collie* inherits the properties of *dog*, which inherits the properties of of *canine*, and so on.

Clearly, an inheritance system is implicit in the prototypical lexicographic definition of a noun. A lexicographer does not store the information that is common to *tree* and *plant* with both entries; the lexicographer stores the redundant information only with *plant*, then writes the definition of *tree* in such a way that a reader will know where to find it. With a printed dictionary, however, a user must look up repeated entries in order to find information that can be instantly retrieved and displayed by a computer.

WordNet is a lexical inheritance system; a systematic effort has been made to connect hyponyms with their superordinates (and vice versa). In the WordNet database, an entry for *tree* contains a reference, or pointer ‘@→,’ to an entry for *plant*; the pointer is labeled ‘‘superordinate’’ by the arbitrary symbol ‘@.’ Thus, the synset for tree would look something like:

{ *tree, plant,@ conifer,~ alder,~* . . . }

where the ‘. . .’ is filled with many more pointers to hyponyms. In the database, the pointer ‘@’ to the superordinate *plant* will be reflected by an inverse pointer ‘~’ to *tree* in the synset for *plant*; that pointer is labeled ‘‘hyponym’’ by the arbitrary symbol ‘~’:

{ *plant, flora, organism,@ tree,~* . . . }

{*tree*} is not the only hyponym of {*plant, flora*}, of course; others have been omitted here in order not to obscure the reciprocity of ‘@’ and ‘~’. The computer is programmed to use these labeled pointers to construct whatever information a user requests; the arbitrary symbols ‘@’ and ‘~’ are suppressed when the requested information is displayed. (There is no need for special tags on *tree* or *plant*, to distinguish which senses are intended because nouns denoting living plants are all in one file, whereas nouns denoting graphical trees or manufacturing plants are elsewhere, as will be explained below.)

It should be noted, at least parenthetically, that WordNet assumes that a distinction can always be drawn between synonymy and hyponymy. In practice, of course, this distinction is not always clear, but in a conventional dictionary that causes no problems. For example, a conventional dictionary can include in its entry for *board* the information that this term can be used to refer to surf boards or to skate boards. That is to say, in addition to the generic meaning of *board*, there are specific meanings of *board* that are hyponyms of the generic meaning. If the information were entered this way in WordNet, however, then a request for information about the superordinates of *board* would elicit the same path twice, the only difference being that one path would be prefaced by {*surf board, board*} @→ *board*. In WordNet, therefore, an effort has been made to avoid entries in which a term is its own hyponym. Thus, for example, *cat* is entered in WordNet as the superordinate of *big cat* and *house cat,* even though to most people the primary sense of *cat*—the meaning that comes first to mind—is {*house cat, tabby, pussy, pussy cat, domesticated cat*}. WordNet does not make explicit the fact that *cat* is frequently used to refer to pet cats, but relies on general linguistic knowledge that a superordinate term can replace a more specific term whenever the context insures that no confusion will result.

What benefits follow from treating lexical knowledge as an inheritance system? In the introduction to this paper, four examples of information missing from conventional definitions were described. Of those four, the first three can be repaired by the judicious use of labeled pointers; with a computer it is as easy to move from superordinate to hyponyms as it is to move from hyponym to superordinate. The fourth omission—of all the associated general knowledge about a referent that is not given in a term's definition—stands uncorrected in WordNet; somewhere a line must be drawn between lexical concepts and general knowledge, and WordNet is designed on the assumption that the standard lexicographic line is probably as distinct as any could be.

Psycholinguistic Assumptions

Since WordNet is supposed to be organized according to principles governing human lexical memory, the decision to organize the nouns as an inheritance system reflects a psycholinguistic judgment about the mental lexicon. What kinds of evidence provide a basis for such decisions?

The isolation of nouns into a separate lexical subsystem receives some support from clinical observations of patients with anomic aphasia. After a left-hemisphere stroke that affects the ability to communicate linguistically, most patients are left with a deficit in naming ability (Caramazza and Berndt, 1978). In anomic aphasia, there is a specific inability to name objects. When confronted with an apple, say, patients may be unable to utter ‘‘apple,’’ even though they will reject such suggestions as *shoe* or *banana,* and will recognize that *apple* is correct when it is provided. They have similar difficulties in naming pictured objects, or in providing a name when given its definition, or in using nouns in spontaneous speech. Nouns that occur frequently in everyday usage tend to be more accessible than are rarely used nouns, but a patient with severe anomia looks for all the world like someone whose semantic memory for nouns has become disconnected from the rest of the lexicon. However, clinical symptoms are characterized by great variability from one patient to the next, so no great weight should be assigned to such observations.

Psycholinguistic evidence that knowledge of nouns is organized hierarchically comes from the ease with which people handle anaphoric nouns and comparative constructions. (1) Superordinate nouns can serve as anaphors referring back to their hyponyms. For example, in such constructions as *He owned a rifle, but the gun had not been fired*, it is immediately understood that *the gun* is an anaphoric noun with *a rifle* as its antecedent. Moreover, (2) superordinates and their hyponyms cannot be compared (Bever and Rosenbaum, 1970). For example, both *A rifle is safer than a gun* and *A gun is safer than a rifle* are immediately recognized as semantically anomalous. Such judgments demand an explanation in terms of hierarchical semantic relations.

More to the point, however, is the question: is there psycholinguistic evidence that people's lexical memory for nouns forms an inheritance system? The first person to make this claim explicit seems to have been Quillian (1967, 1968). Experimental tests of Quillian's proposal were reported in a seminal paper by Collins and Quillian (1969), who assumed that reaction times can be used to indicate the number of hierarchical levels separating two meanings. They observed, for example, that it takes less time to respond True to ‘‘A canary can sing’’ than to ‘‘A canary can fly,’’ and still more time is required to respond True to ‘‘A canary has skin.’’ In this example, it is assumed that *can sing* is stored as a feature of *canary*, *can fly* as a feature of *bird*, and *has skin* as a feature of *animal*. If all three features had been stored directly as features of *canary*, they could all have been retrieved with equal speed. The reaction times are not equal because additional time is required to retrieve *can fly* and *has skin* from the superordinate concepts. Collins and Quillian concluded from such observations that generic information is not stored redundantly, but is retrieved when needed. (In WordNet, the hierarchy is: *canary @→ finch @→ passerine @→ bird @→ vertebrate @→ animal*, but these intervening levels do not affect the general argument that Collins and Quillian were making.)

Most psycholinguists agree that English common nouns are organized hierarchically in semantic memory, but whether generic information is inherited or is stored redundantly is still moot (Smith, 1978). The publication of Collins and Quillian's (1969) experiments stimulated considerable research, in the course of which a number of problems were raised. For example, according to Quillian's theory, *robin* and *ostrich* share the same kind of semantic link to the superordinate *bird*, yet ‘‘A robin is a bird’’ is confirmed more rapidly than is ‘‘An ostrich is a bird’’ (Wilkins, 1971). Or, again, *can move* and *has ears* are both properties that people associate with *animal*, yet ‘‘An animal can move’’ is confirmed more rapidly than is ‘‘An animal has ears’’ (Conrad, 1972). From these and similar results, many psycholinguists concluded that Quillian was wrong, that semantic memory for nouns is not organized as an inheritance system.

An alternative conclusion—the conclusion on which WordNet is based—is that the inheritance assumption is correct, but that reaction times do not measure what Collins and Quillian, and other experimentalists assumed they did. Perhaps reaction times indicate a pragmatic rather than a semantic distance—a difference in word use, rather than a difference in word meaning (Miller and Charles, 1991).

Semantic Components

One way to construe the hierarchical principle is to assume that all nouns are contained in a single hierarchy. If so, the topmost, or most generic level would be semantically empty. In principle, it is possible to put some vague abstraction designated, say, {*entity*}, at the top; to make {*object, thing*} and {*idea*} its immediate hyponyms, and so to continue down to more specific meanings, thus pulling all nouns together into a single hierarchical memory structure. In practice, however, these abstract generic concepts carry little semantic information; it is doubtful that people could even agree on appropriate words to express them.

The alternative is to partition the nouns with a set of semantic primes—to select a (relatively small) number of generic concepts and to treat each one as the unique beginner of a separate hierarchy. These multiple hierarchies correspond to relatively distinct semantic fields, each with its own vocabulary. That is to say, since the features that characterize a unique beginner are inherited by all of its hyponyms, a unique beginner can be regarded as a primitive semantic component of all words in its hierarchically structured semantic field. Partitioning the nouns also has practical advantages: it reduces the size of the files that the lexicographers must work with, and makes it possible to assign the writing and editing of different files to different lexicographers.

Table 1
List of 25 unique beginners for WordNet nouns

{*act, action, activity*}	{*natural object*}
{*animal, fauna*}	{*natural phenomenon*}
{*artifact*}	{*person, human being*}
{*attribute, property*}	{*plant, flora*}
{*body, corpus*}	{*possession*}
{*cognition, knowledge*}	{*process*}
{*communication*}	{*quantity, amount*}
{*event, happening*}	{*relation*}
{*feeling, emotion*}	{*shape*}
{*food*}	{*state, condition*}
{*group, collection*}	{*substance*}
{*location, place*}	{*time*}
{*motive*}	

The problem, of course, is to decide what these primitive semantic components should be. Different workers make different choices; one important criterion is that, collectively, they should provide a place for every English noun. WordNet has adopted the set of twenty-five unique beginners that are listed in Table 1. These hierarchies vary widely in size and are not mutually exclusive—some cross-referencing is required—but on the whole they cover distinct conceptual and lexical domains. They were selected after considering the possible adjective-noun combinations that could be expected to occur (that analysis was carried out by Philip N. Johnson-Laird). The rationale will be discussed below.

Once the primitive semantic components had been chosen, however, some natural groupings among them were observed. Seven of the components, for example, were concerned with living or non-living things; they could be arranged hierarchically as diagrammed in Figure 1. Accordingly, a small 'Tops' file was created in order to include these semantic relations in the system. However, the great bulk of WordNet's nouns are contained in the twenty-five component files.

Figure 1. Diagrammatic representation of hyponymic relations
among seven unique beginners
denoting different kinds of tangible things.

		{*plant, flora*}
	{*living thing, organism*}	{*animal, fauna*}
		{*person, human being*}
{*thing, entity*}		
		{*natural object*}
		{*artifact*}
	{*non-living thing, object*}	{*substance*}
		{*food*}

It is of some interest that these files are relatively shallow. In principle, of course, there is no limit to the number of levels an inheritance system can have. Lexical inheritance systems, however, seldom go more than ten levels deep, and the deepest examples usually contain technical levels that are not part of the everyday vocabulary. For example, a Shetland pony is a pony, a horse, an equid, an odd-toed ungulate, a herbivore, a mammal, a vertebrate, and an animal; pursuing it into the Tops file adds organism and entity: eleven levels, most of them technical. Some hierarchies are deeper than others: man-made artifacts sometimes go six or seven levels deep (*roadster @→ car @→ motor vehicle @→ wheeled vehicle @→ vehicle @→ conveyance @→ artifact*), whereas the hierarchy of persons runs about three or four (one of the deepest is *televangelist @→ evangelist @→ preacher @→ clergyman @→ spiritual leader @→ person*). Advocates of redundant storage of the information associated with these concepts point out that the more generic information would be repeated over and over in a redundant system, so each additional level would put an increasingly severe burden on lexical memory—a possible reason that the number of levels is limited.

Distinguishing Features

These hierarchies of nominal concepts are said to have a level, somewhere in the middle, where most of the distinguishing features are attached. It is referred to as the basic level, and the nominal concepts at this level are called basic-level categories or generic concepts (Berlin, Breedlove, and Raven, 1966, 1973). Rosch (1975; Rosch, Mervis, Gray, Johnson, and Boyes-Braem, 1976) extended this generalization: for concepts at the basic level, people can list many distinguishing features. Above the basic level, descriptions are brief and general. Below the base level, little is added to the features that distinguish basic concepts. These observations have been made largely for the names of concrete, tangible objects, but some psycholinguists have argued that a base or primary level should be a feature of every lexical hierarchy (Hoffman and Ziessler, 1983).

Although the overall structure of noun hierarchies is generated by the hyponymy relation, details are given by the features that distinguish one concept from another. For example, a canary is a bird that is small, colorful, sings, and flies, so not only must *canary* be entered as a hyponym of *bird*, but the attributes of small size and bright color must also be included, as well as the activities of singing and flying. Moreover, *canary* must inherit from *bird* the fact that it has a beak and wings with feathers. In order to make all of this information available when *canary* is activated, it must be possible to associate *canary* appropriately with at least three different kinds of distinguishing features (Miller, in press):

(1) Attributes: *small, yellow*

(2) Parts: *beak, wings*

(3) Functions: *sing, fly*

Each type of distinguishing feature must be treated differently.

Note that attributes are given by adjectives, parts by nouns, and functions by verbs. If the association of *canary* with each of these features is to be represented in WordNet by labeled pointers, then pointers will be required from nouns to adjectives and from nouns to verbs. As this is written, allowance has been made for including such pointers in WordNet, but the possibility has not yet been coded by the lexicographers; only the pointers to parts, which go from nouns to nouns, have been implemented.

When WordNet was first conceived, it was not intended to include information about distinguishing features. It was assumed that WordNet would be used in close conjunction with some on-line dictionary, and that the distinguishing features of a lexical concept would be available from that source. As the coverage of WordNet increased, it became increasingly obvious that alternative senses of a word could not always be identified by the use of synonyms. Rather late in the game, therefore, it was decided to include distinguishing features in the same way that conventional dictionaries do, by including short explanatory glosses as a part of synsets containing polysemous words. These are marked off from the rest of the synset by parentheses. For example, the {*artifact*} hierarchy in WordNet contains eight different senses of the highly polysemous noun *case*:

{*carton, case0, box,@* (a box made of cardboard; opens by flaps on the top)}

{*case1, bag,@* (a portable bag for carrying small objects)}

{*case2, pillowcase, pillowslip, slip2, bed linen,@* (a removable and washable cover for a pillow)} {*bag1, case3, grip, suitcase, traveling bag*,@ (a portable rectangular traveling bag for carrying clothes)}

{*cabinet, case4, console, cupboard*,@ (a cupboard with doors and shelves)}

{*case5, container*,@ (a small portable metal container)}

{*shell, shell plating, case6, casing1, outside surface*,@ (the outer covering or housing of something)}

{*casing, case7, framework*,@ (the enclosing frame around a door or window opening)}

The parenthetical glosses serve to keep the several senses distinct, but a certain redundancy is apparent between the superordinate concepts, indicated by '@,' and the head words of the defining gloss. As more distinguishing features come to be indicated by pointers, these glosses should become even more redundant. An imaginable test of the system would then be to write a computer program that would synthesize glosses from the information provided by the pointers.

At the present time, however, attributive and functional features are not available for many words, and where they are available, it is in the form of defining glosses, not labeled pointers to the appropriate adjectives or verbs. But part-whole relations are available in WordNet; experience with these distinguishing features should provide a basis for the future implementation of cross-part-of-speech pointers.

Attributes and Modification

Values of attributes are expressed by adjectives. For example, size and color are attributes of canaries: the size of canaries can be expressed by the adjective *small*, and the usual color of canaries can be expressed by the adjective *yellow*. There is no semantic relation comparable to synonymy or hyponymy that can serve this function, however. Instead, adjectives are said to modify nouns, or nouns are said to serve as arguments for attributes: Size(*canary*) = *small*.

Although the possibility has not yet been implemented in WordNet, the fact that a canary is small could be represented by a labeled pointer in much the same way as the fact that a canary is a bird is represented. Formally, the difference is that there would be no return pointer from *small* back to *canary*. That is to say, although people will list *small* when asked for the features of canaries, when asked to list small things they are unlikely to group together canaries, pygmies, ponies, and closets. The pointer from *canary* to *small* is interpreted with respect to the immediate superordinate of *canary*, i.e., *small for a bird*, but that anchor to a head noun is lost when *small* is accessed alone.

The semantic structure of adjectival concepts is discussed by Gross and Miller (this volume). Here it is sufficient to point out that the attributes associated with a noun are reflected in the adjectives that can normally modify it. For example, a canary can be hungry or satiated because hunger is a feature of animals and canaries are animals, but *a stingy canary* or *a generous canary* could only be interpreted metaphorically, since generosity is not a feature of animals in general, or of canaries in particular. Keil (1979, 1983) has argued that children learn the hierarchical structure of nominal concepts by observing what can and cannot be predicated at each level. For example, the important semantic distinction between animate and inanimate nouns derives from the fact that the adjectives *dead* and *alive* can be predicated of one class of nouns but not of the other. Although such selectional restrictions on adjectives are not represented explicitly in WordNet, they did motivate the partitioning of the nouns into the twenty-five semantic components listed above.

Parts and Meronymy

The part-whole relation between nouns is generally considered to be a semantic relation, called meronymy (from the Greek *meros*, part; Cruse, 1986), comparable to synonymy, antonymy, and hyponymy. The relation has an inverse: if W_m is a meronym of W_h, then W_h is said to be a holonym of W_m.

Meronyms are distinguishing features that hyponyms can inherit. Consequently, meronymy and hyponymy become intertwined in complex ways. For example, if *beak* and *wing* are meronyms of *bird*, and if *canary* is a hyponym of *bird*, then, by inheritance, *beak* and *wing* must also be meronyms of *canary*. Although the connections may appear complex when dissected in this manner, they are rapidly deployed in language comprehension. For example, most people do not even notice the inferences required to establish a connection between the following sentences: *It was a canary. The beak was injured*. Of course, after *canary* has inherited *beak* often enough, the fact that canaries have beaks may come to be stored redundantly with the other features of *canary*, but that possibility does not mean that the general structure of people’s lexical knowledge is not organized hierarchically.

The connections between meronymy and hyponymy are further complicated by the fact that parts are hyponyms as well as meronyms. For example, {*beak, bill, neb*} is a hyponym of {*mouth, muzzle*}, which in turn is a meronym of {*face, countenance*} and a hyponym of {*orifice, opening*}. A frequent problem in establishing the proper relation between hyponymy and meronymy arises from a general tendency to attach features too high in the hierarchy. For example, if *wheel* is said to be a meronym of *vehicle*, then sleds will inherit wheels they should not have. Indeed, in WordNet a special synset was created for the concept, {*wheeled vehicle*}.

It has been said that distinguishing features are introduced into noun hierarchies primarily at the level of basic concepts; some claims have been made that meronymy is particularly important for defining basic terms (Tversky and Hemenway, 1984). Tests of these claims, however, have been concerned primarily with words denoting physical objects, which is where meronyms tend to occur most frequently. In WordNet, meronymy is found primarily in the {*body, corpus*}, {*artifact*}, and {*quantity, amount*} hierarchies. For concrete objects like bodies and artifacts, meronyms do indeed help to define a basic level. No such level is apparent for terms denoting quantities, however, where small units of measurement are parts of larger units at every level of the hierarchy. Since attributes and functions have not yet been coded, no attempt has been made to see whether a basic level can be defined for the more abstract hierarchies.

The ‘‘part of’’ relation is often compared to the ‘‘kind of’’ relation: both are asymmetric and (with reservations) transitive, and can relate terms hierarchically (Miller and Johnson-Laird, 1976). That is to say, parts can have parts: a finger is a part of a hand, a hand is a part of an arm, an arm is a part of a body: the term *finger* is a meronym of the term *hand*, *hand* is a meronym of *arm*, *arm* is a meronym of *body*. But the ‘‘part of’’ construction is not always a reliable test of meronymy. A basic problem with meronymy is that people will accept the test frame, ‘‘W_m is a part of W_h,’’ for a variety of part-whole relations.

In many instances transitivity seems to be limited. Lyons (1977), for example, notes that *handle* is a meronym of *door* and *door* is a meronym of *house*, yet it sounds odd to say ‘‘The house has a handle’’ or ‘‘The handle is a part of the house.’’ Winston, Chaffin, and Hermann (1987) take such failures of transitivity to indicate that different part-whole relations are involved in the two cases. For example, ‘‘The branch is a part of the tree’’ and ‘‘The tree is a part of a forest’’ do not imply that ‘‘The branch is a part of the forest’’ because the *branch/tree* relation is not the same as the *tree/forest* relation. For Lyons’ example, they suggest, following Cruse (1986), that ‘‘part of’’ is sometimes used where ‘‘attached to’’ would be more appropriate: ‘‘part of’’ should be transitive, whereas ‘‘attached to’’ is clearly not. ‘‘The house has a door handle’’ is acceptable because it negates the implicit inference in ‘‘The house has a handle’’ that the handle is attached to the house.

Such observations raise questions about how many different ‘‘part of’’ relations there are. Winston et al. (1987) differentiate six types of meronyms: component-object (*branch/tree*), member-collection (*tree/forest*), portion-mass (*slice/cake*), stuff-object (*aluminum/airplane*), feature-activity (*paying/shopping*), and place-area (*Princeton/New Jersey*). Chaffin, Hermann, and Winston (1988) add a seventh: phase-process (*adolescence/growing up*). Meronymy is obviously a complex semantic relation—or set of relations. Only three of these types of meronymy are coded in WordNet:

W_m #p→ W_h indicates that W_m is a component part of W_h;

W_m #m→ W_h indicates that W_m is a member of W_h; and

W_m #s→ W_h indicates that W_m is the stuff that W_h is made from.

Of these three, the ‘is a component of’ relation ‘#p’ is by far the most frequent.

The stuff-object relation demonstrates the limits of folk theories of object composition. With the help of modern science it is now possible to analyze ‘‘stuff’’ into smaller and smaller components. At some point, this analysis loses all connection with the object being analyzed. For example, since all concrete objects are composed of atoms, having atoms as a part will not distinguish one category of concrete objects from any other. *Atom* would be a meronym of every term denoting a concrete object. Something has gone wrong here. For commonsense purposes, the dissection of an object terminates at the point where the parts no longer serve to distinguish this object from others with which it might be confused. Knowing where to stop requires commonsense knowledge of the contrasts that need to be drawn.

This problem arises for many parts other than atoms, of course. Some components can serve as parts of many different things: think of all the different objects that have gears. It is sometimes the case that an object can be two kinds of thing at the same time—a piano is both a kind of musical instrument and a kind of furniture, for example—which results in what is sometimes called a tangled hierarchy (Fahlman, 1979). Tangled hierarchies are rare when hyponymy is the semantic relation. In meronymic hierarchies, on the other hand, it is common; *point*, for example, is a meronym of *arrow, awl, dagger, fishhook, harpoon, icepick, knife, needle, pencil, pin, sword, tine*; *handle* has an even greater variety of holonyms. Since the points and handles involved are so different from one holonym to the next, it is remarkable that this situation causes as little confusion as it does.

Functions and Predication

The term ‘function’ has served many purposes, both in psychology and linguistics, so anyone who uses it is obligated to explain what sense they attach to it in this context. A functional feature of a nominal concept is intended to be a description of something that instances of the concept normally do, or that is normally done with or to them. This usage feels more natural in some cases than in others. For example, it seems natural to say that the function of a pencil is to write or the function of a knife is to cut, but to say that the function of a canary is to fly or to sing seems a bit forced. What is really intended here are all the features of nominal concepts that are described by verbs or verb phrases. Nominal concepts can play various semantic roles as arguments of the verbs that they co-occur with in a sentence: instruments (*knife-cut*), materials (*wool-knit*), products (*hole-dig; picture-paint*), containers (*box-hold*), etc.

There does not seem to be an obvious term for this type of distinguishing feature. They resemble the functional utilities or action possibilities that Gibson (1979) called ‘affordances.’ Gardner (1973), borrowing a term from Jean Piaget, spoke of ‘operativity’; operative concepts are acquired by interaction and manipulation, whereas figurative concepts are acquired visually, without interaction. Lacking a better term, function will serve, although the possibility should not be overlooked that a more precise analysis might distinguish several different kinds of functional features.

The need for functional features is most apparent when attempting to characterize a concept like {*ornament, decoration*}. An ornament can be any size or shape or composition; parts and attributes fail to capture the meaning. But the function of an ornament is clear: it is to make something else appear more attractive. At least since Dunker (1945) described functional fixedness, psychologists have been aware that the uses to which a thing is normally put are a central part of a person’s conception of that thing. To call something a box, for example, suggests that it should function as a container, which blocks the thought of using it for anything else.

There are also linguistic reasons to assume that a thing’s function is a feature of its meaning. Consider the problem of defining the adjective *good*. A good pencil is one that writes easily, a good knife is one that cuts well, a good paint job is one that covers completely, a good light is one that illuminates brightly, and so on. As the head noun changes, *good* takes on a sequence of meanings: writes easily, cuts well, covers completely, illuminates brightly, etc. It is unthinkable that all of these different meanings should be listed in a dictionary entry for *good*. How should this problem be handled?

One solution is to define (one sense of) *good* as ‘performs well the function that its head noun is intended to perform’ (Katz, 1964). A good pencil is one that performs well the function that pencils are intended to perform; a good knife is one that performs well the function that knives are supposed to perform; and so on. This solution puts the burden on the head noun. If an object has a normal function, the noun denoting it must contain information about what that function is. Then when the noun is modified by *good*, the functional feature of the noun’s meaning is marked ‘+’; when it is modified by *bad*, the functional feature is marked ‘-’. If an object has no normal function, then it is inappropriate to say it is good or bad: *a good electron* is semantically anomalous. If something serves several functions, a speaker who says it is good or bad can be misunderstood.

A surprising consequence of this formulation is that an object that is not an *X* can be said to be *a good X* if it performs well the function that *X*’s normally perform. For example, calling a box a chair does not make it one, yet a person who sits on a box can say ‘‘This box is a good chair’’ and so indicate that the box is performing well the function that chairs are expected to perform. Such a sentence would be unintelligible if the function that a chair normally serves were not included as part of the meaning of *chair*.

In terms of the present approach to lexical semantics, functional information should be included by pointers to verb concepts, just as attributes are included by pointers to adjective concepts. In many cases, however, there is no single verb that expresses the function. And in cases where there is a single verb, it can be circular. For example, if the noun *hammer* is defined by a pointer to the verb *hammer*, both concepts are left in need of definition. More appropriately, the noun *hammer* should point to the verb *pound*, because it usually plays the semantic role of instrument and is used for pounding; the verb *hammer* is a conflation of its superordinate *hit* and the instrument used to do it. The semantic role of nouns like *hammer, wallpaper*, or *box* tend to be the same wherever they occur in sentences, independent of their grammatical role. That is to say, in both *John hit the mugger with a hammer* and *The hammer hit him on the head*, the semantic role of *hammer* is that of an instrument. Similarly, *wool* is a semantic material in each of the following sentences: *She knitted the wool into a scarf*, *She knitted a scarf out of the wool*, and *This wool knits well*. This consistency in mapping onto the same semantic role independently of syntax is not a feature of all nominal concepts, however: what is the function of *apple* or *cat*?

Although functional pointers from nouns to verbs have not yet been implemented in WordNet, the hyponymic hierarchy itself reflects function strongly. For example, a term like *weapon* demands a functional definition, yet hyponyms of *weapon*—*gun, sword, club*, etc.—are specific kinds of things with familiar structures (Wierzbicka, 1984). Indeed, many tangles in the noun hierarchy result from the competing demands of structure and function. Particularly among the human artifacts there are things that have been created for a purpose; they are defined both by structure and use, and consequently earn double superordinates. For example, {*ribbon, band*} is a strip of cloth on structural grounds, but an adornment on functional grounds; {*balance wheel*} is structurally a wheel, but functionally a regulator; {*cairn*} is a pile of stones that functions as a marker; etc. Functional pointers from these nominal concepts to the verbal concepts {*adorn*}, {*regulate*}, {*mark*}, etc. could eliminate many of these tangles. At this time it is not obvious which representation (if not both) has the greater psycholinguistic validity.

The details are obviously complicated and it is hard to feel that a satisfactory understanding of these functional attributes of nominal concepts has yet been achieved. If support for the continued development of WordNet is forthcoming, the exercise of adding pointers from nouns to the verbs that express their functions should lead to deeper insight into the problem.

Antonymy

The strongest psycholinguistic indication that two words are antonyms is that each is given on a word association test as the most common response to the other. For example, if people are asked for the first word they think of (other than the probe word itself) when they hear ‘‘victory,’’ most will respond ‘‘defeat’’; when they hear ‘‘defeat,’’ most will respond ‘‘victory.’’ Such oppositions are most common for deadjectival nouns: *happiness* and *unhappiness* are noun antonyms because they derive from the antonymous adjectives *happy* and *unhappy*.

Semantic opposition is not a fundamental organizing relation between nouns, but it does exist and so merits its own representation in WordNet. For example, the synsets for *man* and *woman* would contain:

{ [*man, woman*,!], *person*,@ . . . (a male person) }
{ [*woman, man*,!], *person*,@ . . . (a female person) }

where the symmetric relation of antonymy is represented by the ‘!’ pointer, and square brackets indicate that antonymy is a lexical relation between words, rather than a semantic relation between concepts. This particular opposition echoes through the kin terms, being inherited by *husband/wife, father/mother, son/daughter, uncle/aunt, brother/sister, nephew/niece,* and even beyond: *king/queen, duke/duchess, actor/actress,* etc.

When all three kinds of semantic relations—hyponymy, meronymy, and antonymy—are included, the result is a highly interconnected network of nouns. A graphical representation of a fragment of the noun network is shown in Figure 2. There is enough structure to hold each lexical concept in its appropriate place relative to the others, yet there is enough flexibility for the network to grow and change with learning.

Figure 2. Network representation of three semantic relations
among an illustrative variety of lexical concepts

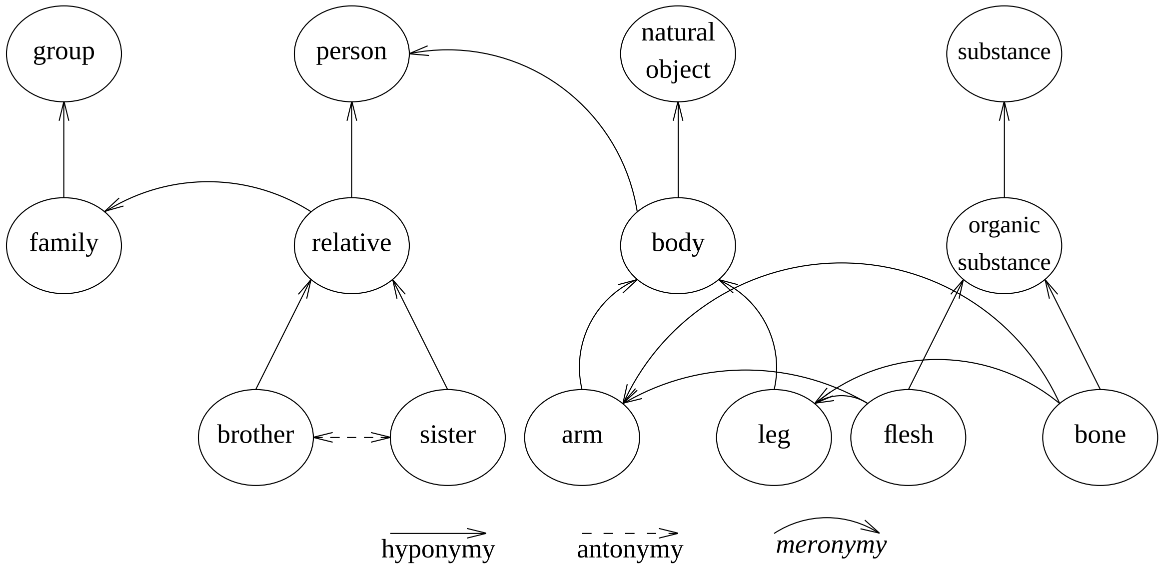

Adjectives in WordNet

Christiane Fellbaum, Derek Gross, and Katherine Miller
(Revised August 1993)

WordNet divides adjectives into two major classes: descriptive and relational. Decriptive adjectives ascribe to their head nouns values of (typically) bipolar attributes and consequently are organized in terms of binary oppositions (antonymy) and similarity of meaning (synonymy). Descriptive adjectives that do not have direct antonyms are said to have indirect antonyms by virtue of their semantic similarity to adjectives that do have direct antonyms. WordNet contains pointers between descriptive adjectives expressing a value of an attribute and the noun by which that attribute is lexicalized. Reference-modifying adjectives have special syntactic properties that distinguish them from other descriptive adjectives. Relational adjectives are assumed to be stylistic variants of modifying nouns and so are cross-referenced to the noun files. Chromatic color adjectives are regarded as a special case.

All languages provide some means of modifying or elaborating the meanings of nouns, although they differ in the syntactic form that such modification can assume.

English syntax allows for a variety of ways to express the qualification of a noun. For example, if *chair* alone is not adequate to select the particular chair a speaker has in mind, a more specific designation can be produced with adjectives like *large* and *comfortable*. Words belonging to other syntactic categories can function as adjectives, such as present and past participles of verbs (*the creaking chair*; *the overstuffed chair*) and nouns (*armchair*, *barber chair*). Phrasal modifiers are prepositional phrases (*chair by the window*, *chair with green upholstery*) and noun phrases (*my grandfather's chair*). Entire clauses can modify nouns, as in *The chair that you bought at the auction*. Prepositional phrases and clausal noun modifiers follow the noun; genitive noun phrases and single word modifiers precede it.

Noun modification is primarily associated with the syntactic category ‘‘adjective.’’ Adjectives have as their sole function the modification of nouns, whereas modification is not the primary function of noun, verb, and prepositional phrases. Adjectives have particular semantic properties that are not shared by other modifiers; some of these are discussed. The lexical organization of adjectives is unique to them, and differs from that of the other major syntactic categories, noun and verb.

The adjective synsets in WordNet contain mostly adjectives, although some nouns and prepositional phrases that function frequently as modifiers have been entered as well. The present discussion will be limited to adjectives.

WordNet presently contains approximately 19,500 adjective word forms, organized into approximately 10,000 word meanings (synsets).

WordNet contains descriptive adjectives (such as *big*, *interesting*, *possible*) and relational adjectives (such as *presidential* and *nuclear*). A relatively small number of adjectives including *former* and *alleged* constitute the closed class of reference-modifying adjectives. Each of these classes is distinguished by the particular semantic and syntactic properties of its adjectives.

Descriptive Adjectives

Descriptive adjectives are what one usually thinks of when adjectives are mentioned. A descriptive adjective is one that ascribes a value of an attribute to a noun. That is to say, *x is Adj* presupposes that there is an attribute A such that A(*x*) = *Adj*. To say *The package is heavy* presupposes that there is an attribute WEIGHT such that WEIGHT(*package*) = *heavy*. Similarly, *low* and *high* are values for the attribute HEIGHT. WordNet contains pointers between descriptive adjectives and the noun synsets that refer to the appropriate attributes.

The semantic organization of descriptive adjectives is entirely different from that of nouns. Nothing like the hyponymic relation that generates nominal hierarchies is available for adjectives: it is not clear what it would mean to say that one adjective ''is a kind of'' some other adjective. The semantic organization of adjectives is more naturally thought of as an abstract hyperspace of *N* dimensions rather than as a hierarchical tree.

Antonymy: The basic semantic relation among descriptive adjectives is antonymy. The importance of antonymy first became obvious from results obtained with word association tests: When the probe is a familiar adjective, the response commonly given by adult speakers is its antonym. For example, to the probe *good*, the common response is *bad*; to *bad*, the response is *good*. This mutuality of association is a salient feature of the data for adjectives (Deese, 1964, 1965). It seems to be acquired as a consequence of these pairs of words being used together in the same phrases and sentences (Charles and Miller, 1989; Justeson and Katz, 1991a, 1991b).

The importance of antonymy in the organization of descriptive adjectives is understandable when it is recognized that the function of these adjectives is to express values of attributes, and that nearly all attributes are bipolar. Antonymous adjectives express opposing values of an attribute. For example, the antonym of *heavy* is *light*, which expresses a value at the opposite pole of the WEIGHT attribute. In WordNet, this binary opposition is represented by reciprocal labeled pointers: *heavy* !→ *light* and *light* !→ *heavy*.

This account suggests two closely related questions, which can serve to organize the following discussion.

(1) When two adjectives have closely similar meanings, why do they not have the same antonym? For example, why do *heavy* and *weighty*, which are closely similar in meaning, have different antonyms, *light* and *weightless*, respectively?

(2) If antonymy is so important, why do many descriptive adjectives seem to have no antonym? For example, continuing with WEIGHT, what is the antonym of *ponderous*? To the suggestion that *light* is the antonym of *ponderous*, the reply must be that the antonym of *light* (in the appropriate sense) is *heavy*. Is some different semantic relation (other than antonymy) involved in the subjective organization of the rest of the adjectives?

The first question caused serious problems for WordNet, which was initially conceived as using labeled pointers between synsets in order to represent semantic relations between lexical concepts. But it is not appropriate to introduce antonymy by labeled pointers between the synsets {*heavy, weighty, ponderous*} and {*light, weightless, airy*}. People who know English judge *heavy/light* to be antonyms, and perhaps *weighty/weightless*, but they pause and are puzzled when asked whether *heavy/weightless* or *ponderous/airy* are antonyms. The concepts are opposed, but the word forms are not familiar antonym pairs.

The problem here is that the antonymy relation between word forms is not the same as the conceptual opposition between word meanings. Except for a handful of frequently used adjectives (most of which are Anglo-Saxon), most antonyms of descriptive adjectives are formed by a morphological rule that changes the polarity of the meaning by adding a negative prefix (usually the Anglo-Saxon *un-* or the Latinate *in-* and its allomorphs *il-*, *im-*, *ir-*). Morphological rules apply to word forms, not to word meanings; they generally have a semantic reflex, of course, and in the case of antonymy the semantic reflex is so striking that it deflects attention away from the underlying morphological process. But the important consequence of the morphological origin of antonyms is that word-form antonymy is not a relation between meanings—which precludes the simple representation of antonymy by pointers between synsets.

If the familiar semantic relation of antonymy holds only between selected pairs of words like *heavy/light* and *weighty/weightless*, then the second question arises: what is to be done with *ponderous, massive,* and *airy*, which seem to have no appropriate antonyms? The simple answer seems to be to introduce a similarity pointer and use it to indicate that the adjectives lacking antonyms are similar in meaning to adjectives that do have antonyms.

Gross, Fischer, and Miller (1989) proposed that adjective synsets be regarded as clusters of adjectives associated by semantic similarity to a focal adjective that relates the cluster to a contrasting cluster at the opposite pole of the attribute. Thus, *ponderous* is similar to *heavy* and *heavy* is the antonym of *light*, so a conceptual opposition of *ponderous/light* is mediated by *heavy*. Gross, Fischer, and Miller distinguish direct antonyms like *heavy/light*, which are conceptual opposites that are also lexical pairs, from indirect antonyms, like *heavy/weightless*, which are conceptual opposites that are not lexically paired. Under this formulation, all descriptive adjectives have antonyms; those lacking direct antonyms have indirect antonyms, i.e., are synonyms of adjectives that have direct antonyms.

In WordNet, direct antonyms are represented by an antonymy pointer, '!→'; indirect antonyms are inherited through similarity, which is indicated by the similarity pointer, '&→.' The configuration that results is illustrated in Figure 1 for the cluster of adjectives around the direct antonyms, *wet/dry*. For example, *moist* does not have a direct antonym, but its indirect antonym can be found via the path, *moist* &→ *wet* !→ *dry*.

This strategy has been successful with the great bulk of English adjectives, but particular adjectives have posed some interesting problems. Among the few adjectives that have no satisfactory antonym, even in an *un-* form, are some of the strongest and

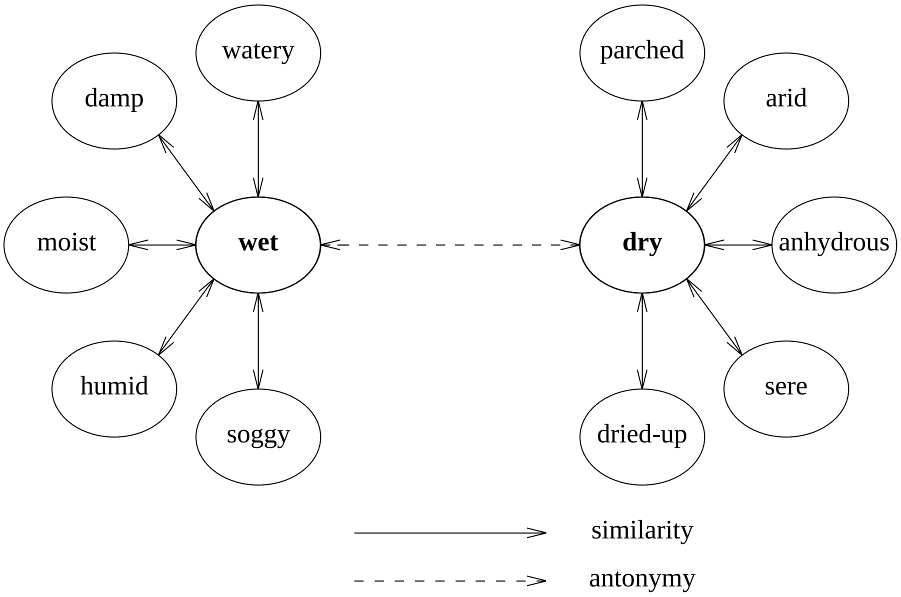

Figure 1. Bipolar Adjective Structure

most colorful. *Angry* is an example. The attribute ANGER is gradable from no anger to extreme fury, but unlike most attributes it does not seem to be bipolar. Many terms are similar in meaning to *angry*: *enraged*, *irate*, *wrathful*, *incensed*, *furious*. But none of them has a direct antonym, either. When adjectives are encountered that do not have direct antonyms, the usual strategy is to search for a related antonym pair and to code the unopposed adjective as similar in meaning to one or the other member of that pair. In the case of *angry*, the best related pair seems to be *pleased/displeased*, but coding *angry* &→ *displeased* seems to miss the essential meaning of *angry*. (And *amicable/hostile* is even worse.) In order to deal with this situation, a special cluster headed *angry/not angry* was created, with *calm* and *placid* (which indicate absence of emotional disturbance) coded as similar in meaning to the synthetic adjective *not angry*. The significance of such exceptions is not obvious, but the recognition that there are exceptions is unavoidable.

The construction of the antonym clusters is discussed in more detail later. We believe that the model presented here—dividing adjectives into two major types, descriptive (which enter into clusters based on antonymy) and relational (which are similar to nouns used as modifiers)—accounts for the majority of English adjectives. We do not claim complete coverage.

Gradation: Most discussions of antonymy distinguish between contradictory and contrary terms. This terminology originated in logic, where two propositions are said to be contradictory if the truth of one implies the falsity of the other and are said to be contrary if only one proposition can be true but both can be false. Thus, *alive* and *dead* are said to be contradictory terms because the truth of *Kennedy is dead* implies the falsity of *Kennedy is alive*, and vice versa. And *fat* and *thin* are said to be contrary terms because *Kennedy is fat* and *Kennedy is thin* cannot both be true, although both can be false if Kennedy is of average weight. However, Lyons (1977, vol. 1) has pointed out that this definition of contrary terms is not limited to opposites, but can be applied so broadly as to be almost meaningless: for example, *Kennedy is a tree* and *Kennedy is a dog* cannot both be true, but both can be false, so *dog* and *tree* must be contraries. Lyons argues that gradability, not truth functions, provides the better explanation of these differences. Contraries are gradable adjectives, contradictories are not.

Gradation, therefore, must also be considered as a semantic relation organizing lexical memory for adjectives (Bierwisch, 1989). For some attributes gradation can be expressed by ordered strings of adjectives, all of which point to the same attribute noun in WordNet. Table 1 illustrates lexicalized gradations for SIZE, WHITENESS, AGE, VIRTUE, VALUE, and WARMTH. (The most difficult grade to find terms for is the neutral middle of each attribute—extremes are extensively lexicalized.)

Table 1
Examples of Some Graded Adjectives

SIZE	WHITENESS	AGE	VIRTUE	VALUE	WARMTH
astronomical	snowy	ancient	saintly	superb	torrid
huge	**white**	**old**	**good**	great	**hot**
large	ash-gray	middle-aged	worthy	**good**	warm
standard	gray	mature	ordinary	mediocre	tepid
small	charcoal	adolescent	unworthy	**bad**	cool
tiny	**black**	**young**	**evil**	awful	**cold**
infinitesimal	pitch-black	infantile	fiendish	atrocious	frigid

But the grading in Table 1 is the exception, not the rule; surprisingly little gradation is lexicalized in English. Most gradation is accomplished in other ways. A gradable adjective can be defined as one whose value can be multiplied by such adverbs of degree as *very, decidedly, intensely, rather, quite, somewhat, pretty, extremely* (Cliff, 1959). And most grading is done by morphological rules for the comparative and superlative degrees, which can be extended if *less* and *least* are used to complement *more* and *most*.

It would not be difficult to represent ordered relations by labeled pointers between synsets, but it was estimated that not more than 2% of the more than 2,500 adjective clusters could be organized in that way. Since the conceptually important relation of gradation does not play a central role in the organization of adjectives, it has not been coded in WordNet.

Markedness: Most attributes have an orientation. It is natural to think of them as dimensions in a hyperspace, where one end of each dimension is anchored at the point of origin of the space. The point of origin is the expected or default value; deviation from it merits comment, and is called the marked value of the attribute.

The antonyms *long/short* illustrate this general linguistic phenomenon known as markedness. In an important paper on German adjectives, Bierwisch (1967) noted that only unmarked spatial adjectives can take measure phrases. For example, *The road is ten miles long* is acceptable; the measure phrase, *ten miles*, describes the LENGTH of the road. But when the antonym is used, as in **The road is ten miles short*, the result is not acceptable (unless the road is short of some goal). Thus, the primary member, *long*, is the unmarked term; the secondary member, *short*, is marked and does not take measure phrases except in special circumstances. Note that the unmarked member, *long*, lends its name to the attribute, LENGTH.

Measure phrases are inappropriate with many attributes, yet markedness is a general phenomenon that characterizes nearly all direct antonyms. In nearly every case, one member of a pair of antonyms is primary: more customary, more frequently used, less remarkable, or morphologically related to the name of the attribute. The primary term is the default value of the attribute, the value that would be assumed in the absence of information to the contrary. Markedness has not been coded in WordNet; it has been assumed that the marked member of the pair is obvious and so needs no explicit indicator. However, the noun that names the attribute—e.g., LENGTH—and all the adjectives expressing values of that attribute (in this case, *long*, *short*, *lengthy*, etc.) are linked in WordNet by a pointer. In a few cases (e.g., *wet/dry, easy/difficult*) it is arguable which term should be regarded as primary, but for the vast majority of pairs the marker is morphologically explicit in the form of a negative prefix: *un+pleasant, in+decent, im+patient, il+legal, ir+resolute*, for example.

Polysemy and Selectional Preferences: Justeson and Katz (1993) find that the different senses of polysemous adjectives like *old*, *right*, and *short* occur with specific nouns (or specific senses of polysemous nouns). For example, the sense of *old* meaning ‘‘not young’’ frequently modifies nouns like *man*, whereas *old* meaning ‘‘not new’’ was found to frequently modify nouns like *house*. Justeson and Katz note that the noun context therefore often serves to disambiguate polysemous adjectives.

An alternative view, put forth by Murphy and Andrew (1993), holds that adjectives are monosemous but that they have different extensions; Murphy and Andrew assert that speakers compute the appropriate meanings in combination with the meanings of the nouns that the adjectives modify. Murphy and Andrew further argue against the claim that antonymy is a relation between two word forms on the basis of the fact that speakers generate different antonyms for an adjective like *fresh* depending on whether it modifies *shirt* or *bread*. WordNet takes the position that these facts point to the polysemy of adjectives like *fresh*; this view is also adopted by Justeson and Katz (1993), who point out that the different antonyms can serve to disambiguate polysemous adjectives.

Adjectives are selective about the nouns they modify. The general rule is that if the referent denoted by a noun does not have the attribute whose value is expressed by the adjective, then that adjective-noun combination requires a figurative or idiomatic interpretation. For example, a building or a person can be tall because buildings and persons have HEIGHT as an attribute, but streets and stories do not have HEIGHT, so *tall street* or *tall story* do not admit literal readings. Nor do antonymy relations hold when nouns lack the pertinent attribute. Compare *short story* with *tall story*, or *short order* with *tall order*. It is really a comment on the semantics of nouns, therefore, when it is said that adjectives vary widely in their breadth of application. Adjectives expressing evaluations (*good/bad, desirable/undesirable*) can modify almost any noun; those expressing activity (*active/passive, fast/slow*) or potency (*strong/weak*, *brave/cowardly*) also have wide ranges of applicability (cf. Osgood, Suci, and Tannenbaum, 1957). Other adjectives are strictly limited with respect to the range of their head nouns (*mown/unmown*; *dehiscent/indehiscent*).

The semantic contribution of adjectives is secondary to, and dependent on, the head nouns that they modify. Edward Sapir (1944) seems to have been the first linguist to point out explicitly that many adjectives take on different meanings when they modify different nouns. Thus, *tall* denotes one range of heights for a building, another for a tree, and still another for a person. It appears that part of the meaning of each of the nouns *building*, *tree*, and *person* is a range of expected values for the attribute HEIGHT. *Tall* is interpreted relative to the expected height of objects of the kind denoted by the head noun: a tall person is someone who is tall *for a person*.

Therefore, in addition to containing a mere list of its attributes, a nominal concept is usually assumed to contain information about the expected values of those attributes: for example, although both buildings and persons have the attribute of HEIGHT, the expected height of a building is much greater than the expected height of a person. The adjective simply modifies those values above or below their default values. The denotation of an adjective-noun combination such as *tall building* cannot be the intersection of two independent sets, the set of tall things and the set of buildings, for then all buildings would be included.

How adjectival information modulates nominal information is not a question to be settled in terms of lexical representations. We assume that the interactions between adjectives and nouns are not prestored but are computed as needed by some on-line interpretative process. As suggested by Miller and Johnson-Laird, ‘‘The nominal information must be given priority; the adjectival information is then evaluated within the range allowed by the nominal information’’ (Miller and Johnson-Laird, 1976, p. 358).

The noun classes of WordNet have been organized in such a way as to make the statement of an adjective’s selectional preferences as simple as possible (Miller, this volume), but as this account is being written those relations have not yet been coded in WordNet.

Syntax: Descriptive adjectives such as *big* and *heavy* are syntactically the freest: they can be used attributively (in prenominal position) or predicatively (after *be*, *become*, *remain*, *stay*, and a few other linking verbs). Some identifying adjectives, such as *best* and *left*, occur mostly attributively: *The best essay* vs. **The essay is best* (but in context, a sentence like *The essay by Mary was best* is acceptable). Reference-modifying and relational adjectives, to be discussed below, are also restricted to attributive use.

Reference-Modifying Adjectives

Bolinger (1967) was the first to note the distinction between reference-modifying and referent-modifying adjectives. He pointed out that in a phrase like *the former president*, what is *former* is the president-hood of the referent, not the referent himself:

the person is former *qua* president. Only the reference to the person as *president* is being qualified by *former*. The nouns modified by adjectives like *former, present, alleged,* and *likely* generally denote a function or social relation. In the phrase *my old friend*, the adjective can be interpreted as reference-modifying, qualifying the friendship between the speaker and the (possibly young) referent of the noun. Under the referent-modifying interpretation of that same adjective, the friend is old (aged), but the friendship need not be. Note that the two senses of this adjective have different antonyms: the reference-modifying sense is opposed to the (reference-modifying) adjectives *recent* or *new*, whereas the referent-modifying adjective has *young* as its antonym. Reference-modifying adjectives are a closed class comprising a few dozen adjectives. Many refer to the temporal status of the noun (*former, present, last, past, late, recent, occasional*); others have an epistemological flavor (*potential, reputed, alleged*); still others are intensifying (*mere, sheer, virtual, actual*). The adjectives express different values of attributes that seem not to be lexicalized, such as ‘‘degree of certainty.’’ Only some of the adjectives have nominalizations (*likelihood, possibility*, and a few others).

The reference-modifying adjectives often function like adverbs: *My former teacher* means *he was formerly my teacher*; *the alleged killer* states that *she is allegedly a killer* or *she allegedly killed*; and *a light eater* is someone that *eats lightly*.

Reference-modifying adjectives can occur only attributively, but not predicatively; compare *The alleged burglar* with **The burglar is alleged*. And the predicative use of *old* as in *My friend is old* disambiguates that adjective, ruling out the *long-standing* reading in favor of the *aged* interpretation. In the current version of WordNet, most reference-modifying adjectives are marked as occurring prenominally only.

Some reference-modifying adjectives resemble descriptive adjectives in that they have direct antonyms: *the possible/impossible task*; *the past/present director*. Those that do not have direct antonyms usually have indirect antonyms.

Color Adjectives

One large and intensively studied class of adjectives is organized differently, and deserves special comment.

English color terms are exceptional in several ways. They can serve as either nouns or adjectives, yet they are not nominal adjectives: they can be graded, nominalized, and conjoined with other descriptive adjectives. But the pattern of direct and indirect antonymy that is observed for other descriptive adjectives does not hold for color adjectives.

Only one color attribute is clearly described by direct antonyms: LIGHTNESS, whose polar values are expressed by *light/dark*. Students of color vision can produce evidence of oppositions between red and green, and between yellow and blue, but those are not treated as direct antonyms in lay speech. The organization of color terms is given by the dimensions of color perception: lightness, hue, and saturation, which define the well-known color solid. In WordNet, however, the opposition *colored/colorless* (cross-referenced to *chromatic/achromatic*) is used to introduce the names of colors. Hues are coded as similar to *colored*, and the shades of gray from white to black are coded as similar to *grey*, which is in a tripartite cluster with *white* and *black*, providing for a graded continuum.

There is some reason to suspect that the elaborate color terminology available in the languages of industrialized countries is a consequence of technological progress and not a natural linguistic development. Speculation about the evolution of color terminology (Berlin and Kay, 1969) suggests that it begins with a single, conventional attribute, LIGHTNESS. Exotic languages are still spoken that have only two color terms to express values of that attribute, and it has been shown that this lexical limitation is not a consequence of perceptual deficits (Heider, 1972; Heider and Olivier, 1972). As technology develops and makes possible the manipulation and control of color, the need for greater terminological precision grows and more color terms appear in the language. They are always added along lines determined by innate mechanisms of color perception rather than by established patterns of linguistic modification.

Relational Adjectives

Another kind of adjective comprises the large and open class of relational adjectives. These, too, can occur only in attributive position, although for some adjectives, this constraint is somewhat relaxed. Relational adjectives, which were first discussed at length by by Levi (1978), mean something like ‘‘of, relating/pertaining to, or associated with’’ some noun, and they play a role similar to that of a modifying noun.

For example, *fraternal*, as in *fraternal twins* relates to *brother*, and *dental*, as in *dental hygiene*, is related to *tooth*. Some head nouns can be modified by both the relational adjective and the noun from which it is derived: both *atomic bomb* and *atom bomb* are admissible.

Some nouns give rise to two homonymous adjectives; one relational adjective restricted to predicative use, the other descriptive. For example, *musical* has a different meaning in *musical instrument* and *musical child*: the first noun phrase does not refer to *an instrument that is musical* but *an instrument used in music*. Similarly, the adjective in *criminal law* is not the same as in *criminal behavior*; this is reflected in the fact that the second adjective, but not the first, is referent-modifying and can be used predicatively. Relational adjectives do not combine well with descriptive adjectives in modifying the same head noun when the two adjectives are linked by a conjunction: *nervous and life-threatening disease* and *musical but not extraordinary talent* sound distinctly odd. (Concatenations like *life-threatening nervous disease* are fine, indicating that the relational adjective acts like a modifying noun.) On the other hand, relational adjectives can easily be conjoined with modifying nouns: *atom and nuclear bombs, the Korean and Vietnam war*.

Relational adjectives are most often derived from Greek or Latin nouns, and less often from the appropriate Anglo-Saxon noun. The English lexicon frequently has several (synonymous) adjectives derived from nouns in different languages that express the same concept: Greek-based *rhinal* and Anglo-Saxon *nasal* both relate to *nose*; the relational adjectives corresponding to *word* are *verbal* (from Latin) and *lexical* (from the Greek). In many cases, these synonyms each pick out their own head nouns and are not substitutable in a given context; compare *nasal/*rhinal passage* and *rhinal/*nasal surgery*.

Conversely, a single relational adjective sometimes points to several nouns: *chemical* has senses corresponding to the two nouns *chemical* (as in *chemical fertilizer*) and *chemistry* (as in *chemical engineer*).

Some homonymous relational adjectives have a common origin but their meanings have drifted apart over time; consequently they point to two distinct noun synsets: one sense of *clerical* points to *clergy* (*clerical leader*); another sense is linked to *clerk* (*clerical work*).

Some relational adjectives do not point to morphologically related English nouns; the Latin or Greek nouns that they are derived from have no exact English equivalents. For example, *fictile* relates to *pottery*, and comes from the Latin word *fictilis*, meaning *made or molded of clay*; there is no corresponding English noun expressing this concept. An adjective like *rural* connects to several related concepts (*country,* as opposed to *city*, and *farming*). In such cases, several senses of the adjective have been entered with pointers to different nouns.

WordNet also has a number of adjectives that are derived from other relational adjectives via some prefix; these adjectives, which include *interstellar*, *extramural*, and *premedical* do not point to any noun but are linked instead to the un-prefixed adjectives (*stellar, mural,* and *medical*, respectively), from which they are derived.

Semantics: Relational adjectives differ from descriptive adjectives in that they do not relate to an attribute: there is no scale of *criminality* or *musicality* on which the adjectives in *criminal law* and *musical training* express a value. The adjective and the related noun refer to the same concept, but they differ formally (morphologically).

Relational adjectives do not refer to a property of their head nouns. This can be seen from the absence of corresponding nominalizations: the descriptive use of *nervous* in *the nervous person* admits such constructions as *the person's nervousness,* but its relational use in *the nervous disorder* does not. Relational adjectives, like nouns and unlike descriptive adjectives, are not gradable: **the extremely atomic bomb,* like **the extremely atom bomb* or **the very baseball game,* are not acceptable. Relational adjectives do not have direct antonyms; although they can often be combined with *non-*, such forms do not express the opposite value of an attribute but something like ‘‘everything else’’; these adjectives have a classifying function. In a few cases, relational adjectives enter into an opposition on the basis of their prefixes: *extracellular* vs. *intracellular*. More frequently, relational adjectives enter into N-way oppositions in combination with a specific head noun (e.g., *civil* opposes *criminal* in combination with *law(yer),* and *mechanical, electrical*, etc., in combination with *engineer(ing)*.)

Since relational adjectives do not have antonyms, they cannot be incorporated into the clusters that characterize descriptive adjectives. And because their syntactic and semantic properties are a mixture of those of adjectives and those of nouns used as noun modifiers, rather than attempting to integrate them into either structure WordNet maintains a separate file of relational adjectives with pointers to the corresponding nouns.

Some 1,700 relational adjective synsets containing over 3,000 individual lexemes are currently included in WordNet. Each synset consists of one or more relational adjectives, followed by a pointer to the appropriate noun. For example, the entry {*stellar, astral, sidereal, noun.object:star*} indicates that *stellar, astral, sidereal* relate to the noun *star*.

Syntax: The semantic relation between a head noun and the noun from which the adjective is derived may differ with different head nouns. For example, *musical evening* means ‘‘an evening with music,’’ whereas *musical instrument* is ‘‘an instrument for (producing) music.’’ Bartning (1980) observes that when the head noun is deverbal, predication is often possible so long as the head noun denotes a state rather than an action. For example, *economic restructuring* refers to an action, and predication is possible: *The restructuring was economic*. By contrast, *economic slump* is a state, and the sentence **the slump is economic* is bad.

Bartning (1980) observed further that if there is a tight, obvious grammatical relation, the adjective cannot be used predicatively; however, when the relation between the adjective and its headnoun is less obvious, predication is possible. Thus, in the noun phrase *presidential election, president* is the object of *elect*; here, the grammatical relation between adjective and head noun is transparent, and predication is not possible: **the election is presidential*. Similarly, *the Pope* is clearly the subject in the phrase *papal visit*, and predication is bad (**The visit was papal*). If, however, the relation is one where the base noun is an adjunct of the head noun, predication is more likely to be acceptable. *Manual labor* is ‘‘labor WITH/BY hand’’, and the phrase ‘‘This labor is (mostly) manual’’ is fine. The syntactic behavior of some relational adjectives that differs with the semantic relation to the particular head noun cannot presently be accounted for in WordNet. The relational adjectives are not provided with syntactic codes.

Predication is also possible when the relation between the noun and the base noun of the adjective is ‘‘like’’: *Nixonian politics* are politics reminiscent of those of a former president; a *presidential speech* is a speech that is like that of a president. Both allow predication: *These politics are truly Nixonian*; *His speech was rather presidential*. Arguably, the meaning of *presidential* is not the same in *presidential speech* and in *presidential election* (where the adjective cannot be used predicatively). In WordNet, such distinctions have generally not been made, because there are too many semantic relations between a relational adjective and its different head nouns to classify the adjectives into distinct senses.

Virtually all relational adjectives can be used predicatively in contrastive contexts: *These weapons are not chemical or biological, but nuclear; They hired a criminal not a corporate lawyer.* However, these cases arguably involve ellipsis of the head noun.

Coding

The semantic organization of descriptive adjectives illustrated in Figure is 1 coded by organizing them into bipolar clusters. There are over 2,500 of these clusters, one for each pair of antonyms; they can be likened to the subject files for nouns and verbs. Each bipolar cluster stands alone, and coding is restricted to within-cluster relations.

The cluster for *wet*/*dry*, which define the attribute WETNESS or MOISTNESS, illustrates the basic coding devices used, and shows the variety and range of senses that can be represented within a cluster.

```
[{ [WET1, DRY1,!] bedewed,& boggy,& clammy,& damp,& drenched,&
   drizzling,& hydrated,& muggy,& perspiring,& saturated2,&
   showery,& tacky,& tearful,& watery2,& WET2,& }
{ bedewed, dewy, wet1,& }
{ boggy, marshy, miry, mucky, muddy, quaggy, swampy, wet1,& }
{ clammy, dank, humid1, wet1,& }
{ damp, moist, wet1,& }
{ drenched, saturated1, soaked, soaking, soppy, soused, wet1,& }
{ drizzling, drizzly, misting, misty, wet1,& }
{ hydrated, hydrous, wet1,& ((chem) combined with water molecules) }
{ muggy, humid2, steamy, sticky1, sultry, wet1,& }
{ perspiring, sweaty, wet1,& }
{ saturated2, sodden, soggy, waterlogged, wet1,& }
{ showery, rainy, wet1,& }
{ sticky2, tacky, undried, wet1,& ("wet varnish") }
{ tearful, teary, watery1, wet1,&  }
{ watery2, wet1,& (filled with water; "watery soil") }
-
{ [DRY1, WET1,!] anhydrous,& arid,& dehydrated,& dried,& dried-up1,&
  dried-up2,& DRY2,& rainless,& thirsty,& }
{ anhydrous, dry1,& ((chem) with all water removed) }
{ arid, waterless, dry1,& }
{ dehydrated, desiccated, parched, dry1,& }
{ dried, dry1,& ("the ink is dry") }
{ dried-up1, dry1,& ("a dry water hole") }
{ dried-up2, sere, shriveled, withered, wizened, dry1,&
  (used of vegetation) }
{ rainless, dry1,& }
{ thirsty, dry1,& }]
```

Each half of the cluster is headed by what is called a head synset. The first two items in each head synset are the antonymous pair that define the attribute represented; head words are coded as such by being capitalized. These head words are followed by &-pointers, one to each synset in the half cluster, each of which has a reciprocal pointer back to the head word. The numerals following certain items distinguish different subsenses or different privileges of occurrence—for example, the *dried-up1* of a water hole in one synset and the *dried-up2* of autumn leaves or fruit in another. Each of these cases, furthermore, contains parenthetical information designed to help distinguish these particular senses or indicate acceptable contexts.

As already mentioned, many adjectives are limited as to the syntactic positions they can occupy, and that limitation is usually coded in WordNet. Because it is a word-form limitation, it is coded for individual adjectives rather than for synsets. Consider the cluster *awake*/*asleep*, both of which are limited to predicate position. Although these are the head words of the cluster, the limitation does not hold for all of the synonyms in the cluster. Therefore, the individual words so limited are all coded with `(p)`.

```
[{ [AWAKE(p), ASLEEP,!] ALERT,& astir(p),& AWARE(p),& CONSCIOUS,&
   insomniac,& unsleeping,& }
{ astir(p), out_of_bed(p), up(p), awake,& }
{ insomniac, sleepless, wakeful, awake,& }
{ unsleeping, wide-awake, awake,& }
-
{ [ASLEEP(p), AWAKE,!] at_rest(p),& benumbed,& DEAD,& dormant,& drowsing,&
  drowsy,& unconscious,& UNAWARE,& UNCONSCIOUS,& }
{ at_rest(p), resting, asleep,& }
{ benumbed, insensible, numb, unfeeling, asleep,& ("my foot is asleep") }
{ dormant, inactive, hibernating, torpid, asleep,& }
{ drowsing, dozing, napping, asleep,& }
{ drowsy, nodding, sleepy, slumberous, slumbrous, somnolent, asleep,& }
{ unconscious, asleep,& }]
```

For adjectives limited to prenominal (attributive) position, the code is `(a)`: for example, {*putative(a), reputed(a), supposed(a),*} as in *the putative father* but not *the father is putative*, and {*bare(a), mere(a),*} as in *the bare minimum* but not *the minimum is bare*. As already mentioned, *former* is used only prenominally, and so are several of its synonyms, *{preceding(a), previous(a), prior(a),}*. When, however, *previous* is used predicatively the sense becomes {*premature, too soon(p),*}, as in *our condemnation of him was a bit previous*.

And, finally, for those few adjectives that can appear only immediately following a noun, the code is `(ip)` for ‘‘immediately postnominal’’: *galore* as in *gore galore*, *elect* as in *president elect*, and *aforethought* as in *malice aforethought*. In many cases the adjectives constitute part of what is essentially a frozen construction.

In addition to the lowercase within-cluster pointers, many head synsets contain pointers to other, related clusters. In this AWAKE/ASLEEP cluster, the capitalized pointer ALERT,& points to the head word of the ALERT/UNALERT cluster. These capitalized pointers are planned to serve as ‘‘see also’’ cross-references to related clusters, even though the present system software is not yet able to make use of them, being tightly restricted to within-cluster coding.

The restricted within-cluster coding leads to a problem when closely related attributes are expressed by more than one pair of antonyms. In such cases, exactly the same set of synsets can be related to two different antonymous pairs, some of which are presently in different clusters. Consider *large*/*small* and *big*/*little*. *Big/little* and *large/small* are equally salient as antonyms: many synsets could just as well be coded as similar to *big* as to *large*. Therefore, a single cluster has been created headed by both pairs, thus avoiding unnecessary redundancy. In addition, a particular synset can be coded with two pointers, one to its own cluster head, the other to the head of an outside cluster.

A final word about *large*/*small* and *big*/*little*: although *large* is clearly opposed to *little*, the pair *large* and *little* are simply not accepted as antonyms. Overwhelmingly, association data and co-occurrence data indicate that *big* and *little* are considered a pair and so are *large* and *small*. These two pairs constitute a prime demonstration that antonymy is as a semantic relation between words rather than between lexicalized concepts.

English Verbs as a Semantic Net

Christiane Fellbaum

This paper describes the semantic network of English verbs in WordNet. The semantic relations used to build networks of nouns and adjectives cannot be applied without modification, but have to be adapted to fit the semantics of verbs, which differ substantially from those of the other lexical categories. The nature of these relations is discussed, as is their distribution throughout different semantic groups of verbs, which determines certain idiosyncratic patterns of lexicalization. In addition, four variants of lexical entailment are distinguished, which interact in systematic ways with the semantic relations. Finally, the lexical properties of the different verb groups are outlined.

Verbs are arguably the most important lexical and syntactic category of a language. All English sentences must contain at least one verb, but, as grammatical sentences with ‘‘dummy’’ subjects like *It is snowing* show, they need not contain a (referential) noun. Many linguists have argued for a model of sentence meaning in which verbs occupy the core position and function as the central organizers of sentences (Chafe, 1970; Fillmore, 1968; and others). The verb provides the relational and semantic framework for its sentence. Its predicate-argument structure (or subcategorization frame) specifies the possible syntactic structures of the sentences in which it can occur. The linking of noun arguments with thematic roles or cases, such as INSTRUMENT, determines the different meanings of the events or states denoted by the sentence, and the selectional restrictions specify the semantic properties of the noun classes that can flesh out the frame. This syntactic and semantic information is generally thought to be part of the verb’s lexical entry, that is to say, part of the information about the verb that is stored in a speaker’s mental lexicon. Because of the complexity of this information, verbs are probably the lexical category that is most difficult to study.

Polysemy

Even though grammatical English sentences require a verb though not necessarily a noun, the language has far fewer verbs than nouns. For example, the *Collins English Dictionary* lists 43,636 different nouns and 14,190 different verbs. Verbs are more polysemous than nouns: the nouns in *Collins* have on the average 1.74 senses, whereas verbs average 2.11 senses.[2]

The higher polysemy of verbs suggests that verb meanings are more flexible than noun meanings. Verbs can change their meanings depending on the kinds of noun arguments with which they co-occur, whereas the meanings of nouns tend to be more stable in the presence of different verbs. Gentner and France (1988) have demonstrated what they call the high mutability of verbs. They presented subjects with sentences

[2] We are indebted to Richard Beckwith for computing these figures.

containing verbs in conjunction with nouns that violated the verbs' selectional restrictions. When asked to paraphrase the sentences, subjects assigned novel interpretations to the verbs, but did not modify the literal meanings of the nouns. Gentner and France concluded that verb meanings are more easily altered because they are less cohesive than those of nouns—a flexibility that makes a semantic analyis of verbs an even more challenging task.

The most frequently used verbs (*have, be, run, make, set, go, take*, and others) are also the most polysemous, and their meanings often depend heavily on the nouns with which they co-occur. For example, dictionaries differentiate between the senses of *have* in sentences like *I have a Mercedes* and *I have a headache*. The difference is less due to the polysemy of *have*, however, than to the concrete or abstract nature of its objects.

In the case of such polysemous verbs as *beat*, meaning differences are determined less by the semantics of the verb's arguments than by different elaborations of one or two common core components shared by most senses of *beat*. Different senses of *beat* occur in very different semantic domains: {*beat, strike, hit*} is a contact verb; {*beat, flatten*} is a verb of change; {*beat, throb, pulse*} is a motion verb; {*beat, defeat*} is a competition verb; {*beat, flog, punish*} and {*beat, circumvent (the system)*} are verbs in the domain of social interaction; {*beat, shape, do metalwork*} is a creation verb; {*beat, baffle*} is a cognition verb; {*beat, stir, whisk*} belongs to the domain of cooking verbs; and {*beat, mark*} is a kind of motion performed to indicate the counts in music. Although most of these verbs seem to share a semantic component of CONTACT or IMPACT, the differences illustrate how flexible these core meanings can be.

In order to reduce ambiguity in WordNet, verb synsets could contain cross-reference pointers to the noun synsets that contain nouns selected for by the verbs. For example, one sense of the verb *throw* (*throw on a wheel*) always selects the noun *pottery* or its hyponyms as its object; that selectional restriction could be represented by a labeled pointer. At the present time, however, this possibility has not been implemented in WordNet.

The Organization of Verbs in WordNet

Currently, WordNet contains over 21,000 verb word forms (of which over 13,000 are unique strings) and approximately 8,400 word meanings (synsets). Included are phrasal verbs like *look up* and *fall back*.

Verbs are divided into 15 files, largely on the basis of semantic criteria. All but one of these files correspond to what linguists have called semantic domains: verbs of bodily care and functions, change, cognition, communication, competition, consumption, contact, creation, emotion, motion, perception, possession, social interaction, and weather verbs. Virtually all the verbs in these files denote events or actions. Another file contains verbs referring to states, such as *suffice, belong,* and *resemble*, that could not be integrated into the other files. The verbs in this latter group do not constitute a semantic domain, and share no semantic properties other than that they refer to states. This file, whose organization resembles that of the adjectives in WordNet, consists of small semantic clusters. The division of verbs into 14 files corresponding to different semantic domains, each containing event and action verbs, and one file containing semantically diverse stative verbs reflects the division between the major conceptual categories EVENT and STATE found in Jackendoff's (1983:170) and Dowty's (1979) analyses. The boundaries between the files are not rigid, and the particular classification was chosen largely because it permits us to get a grasp on the organization of the verbs; it has no further theoretical or psychological implications.

Many of the files derive their names from the topmost verbs, or ''unique beginners,'' which head these semantically coherent lexical groups. These topmost verbs resemble the ''core components'' of Miller and Johnson-Laird (1976). They are the unelaborated concepts from which the verbs constituting the semantic field are derived via semantic relations.

Synonymy

Few truly synonymous verbs, such as *shut* and *close*, can be found in the lexicon—the number depends on how loose a definition of synonymy one adopts. The best examples are probably verb concepts that are represented by both an Anglo-Saxon and a Greco-Latinate word: *begin-commence*, *end-terminate*, *rise-ascend*, *blink-nictate*, *behead-decapitate*, *spit-expectorate*. In general, the Greco-Latinate verbs are used in more formal or technical speech registers: *buy* vs. *purchase*, *sweat* vs. *perspire*, or *shave* vs. *epilate*. Cruse (1986:268) points out that frequently only one member of such a synonym pair tends to be felicitous in a given context: *Where have you hidden Dad's slippers?* sounds more natural than *Where have you concealed Dad's slippers?* And subtle meaning differences can show up in different selectional restrictions. For example, *rise* and *fall* can select as an argument such abstract entities as *the temperature*, but their close synonyms *ascend* or *descend* cannot. Because many apparently synonymous verbs exhibit such differences, verb synsets in WordNet often contain periphrastic expressions, rather than lexicalized synonyms.

These periphrases break down a synonymous verb into an entire verb phrase and thereby often reflect the way in which the verb has become lexicalized by showing constituents that have been conflated in the verb. For example, a denominal verb such as *hammer* is listed with the parenthetical gloss *hit with a hammer*; the conflated verb denotes the function of the noun. The periphrases indicate the basic action and the role of the noun (material and instrument) with which the action is performed. The synonymous expressions of deadjectival verbs often have the form *make* or *become* + some adjective: {*whiten, become white*}, {*enrich, make rich*}, etc. Thus they reflect the fact that these verbs are for the most part verbs of change. The synonymous expressions of many verbs show that they are manner elaborations of a more basic verb: {*swim, travel through water*}, {*mumble, talk indistinctly*}, {*saute, fry briefly*}, etc.

Representing the meanings of verbs is difficult for any theory of lexical semantics, but especially so for WordNet, which differs from previous approaches in avoiding a semantic decomposition in favor of a relational analysis.

Decompositional vs. Relational Semantic Analysis

Most approaches to verb semantics have been attempts at decomposition in one form or another. Early proponents of semantic decomposition (Katz & Fodor, 1963; Katz, 1972; Gruber, 1976; Lakoff, 1970; Jackendoff, 1972; Schank, 1972; Miller and Johnson-Laird, 1976; and others), whether in a generative or an interpretative framework, argued for the existence of a finite set of universal semantic-conceptual components (or primes, or primitives, or atomic predicates, or, in the case of nouns, markers) into which all lexical items could be exhaustively decomposed. Among the examples one finds discussed in the literature, most have been English verbs. The best known example of the decomposition of a verb is probably McCawley's (1968) analysis of *kill* into CAUSE TO BECOME NOT ALIVE, which has been much discussed and criticized (Fodor, 1970; Shibatani, 1972; and others).

Although semantic decomposition has been judged by some to be an inadequate theory of semantic representation (Chomsky, 1972, and others), more recent approaches have taken a similar path to semantic analysis. Both Jackendoff (1983) and Talmy (1985) have proposed an anlysis of verbs in terms of such conceptual categories as EVENT, STATE, ACTION, PATH, MANNER, PLACE, etc. For example, Talmy analyzes the verb *roll* as being a lexicalized conflation of MOVE and MANNER. Both analyses share the assumption of a limited inventory including components or categories that can be expressed not only by verbs, but also by nouns and by operators like NEG.

Relational semantic analysis differs from semantic decomposition primarily in taking lexical items, rather than hypothetically irreducible meaning atoms, as the smallest unit of analysis. Thus, relational analysis has the advantage that its units can be thought of as entries in speakers' mental dictionaries. However, the relational analysis adopted in WordNet shares some aspects of decomposition. Although WordNet does not explicitly recognize conceptual components, some components are reflected in the nature of the semantic relations linking verbs to each other. For example, one of generative semantics' important subpredicates, CAUSE, has the status of a semantic relation in WordNet—a relation that links such verb pairs as *teach*-*learn* and *show*-*see*. This relation also distinguishes systematically between the causative (transitive) and the anticausative (intransitive) senses of certain verb classes, including *break*, *rot* and *move*.

Components like NEG and PATH do not play an overt role in WordNet. But NEG is clearly implicit in the opposition relation that holds between contradictory verbs like *live* and *die*, or *succeed* and *fail*, and between gradables like *like* and *dislike*. And a semantic PATH component is clearly part of the manner relation that links verb concepts like {*move*, *travel*} and *soar*. Other features, such as Talmy's MANNER, are also part of the semantics of WordNet's manner relation (troponymy) linking basic verbs like *eat* and *communicate* to other verbs denoting particular elaborations of the base verb, such as *gobble* and *telex*.

Proponents of semantic decomposition have also argued for the existence of subpredicates, corresponding to abstract verbal concepts, in many lexicalized verbs. Some abstract semantic predicates such as MOVE or GO are argued to be the basic components of most verbs from a wide variety of different semantic fields (Gruber, 1976;

Jackendoff, 1972, 1976). Gruber's Location Hypothesis argues that all events and states can be analyzed as more or less abstract spatial locations and motions. Similarly, Dowty (1979) analyzes all English verbs with the exception of statives as verbs of change, and posits as part of their meaning the semantic subpredicate CHANGE. Thus, Gruber analyzes verbs of giving as an abstract motion undergone by the object that is given; Dowty sees giving as a change of possession. Such decompositions of verb semantics are defensible on an abstract plane, but it is not clear that they reflect the way in which people's memory of verbs is structured; there is no evidence that speakers store verbs of giving together with the concept of (abstract) motion or change of location. Pinker (1989:101) claims that speakers of English decompose verbs into such semantic subpredicates as CAUSE, GO, BE, and PATH, which enables them to predict the verbs' idiosyncratic syntactic behavior. Such an analysis may well be part of speakers' linguistic competence, but there is no indication that it serves to organize the mental lexicon. And there is no evidence that verbs with a complex composition take longer to use or understand.

The subpredicates of lexical decomposition are accorded the same status as other verbs in WordNet. Because verbs like {*move, go*} and *change* refer to very basic concepts, they constitute the root verbs, or topmost ‘‘unique beginners,’’ heading two semantic field. The verbs making up these fields are linked to the root verbs via semantic relations. In other words, WordNet characterizes the verbs of change not by the presence of a semantic component CHANGE, but rather by a semantic relation to the verb *change.* This distinction may appear subtle; it hinges on the formulation of the semantic relations that are coded in WordNet.

Componential analysis could be viewed in terms of entailment, in that a verb V_1 that is a component of another verb V_2 must be entailed by V_1. This idea forms the basis of Carnap's (1947) theory of meaning postulates. Rather than attempting to provide an exhaustive breakdown of words into components, meaning postulates state inferential rules between sentences based on the semantic composition of the words in these sentences. A well-known example is the relation between the sentence *John is a bachelor* and the sentences *John is a male* / *John is an adult* / *John is unmarried*. These entailment relations reflect the fact that the lexical composition of *bachelor* includes the components MALE, ADULT, and UNMARRIED. The meaning of *bachelor* can be expressed in terms of its position in a taxonomic hierarchy, that is, a *bachelor* is a kind of a *man,* which is a kind of a *person*. *Bachelor*, therefore, could be said to inherit all the semantic components of its superordinates, such as HUMAN. Consequently, the sentence *John is a bachelor* also entails the sentence *John is human*. (Katz, 1972, expresses a similar idea by postulating lexical redundancy rules concerning the co-occurrence of semantic components.) Carnap's meaning postulates also predict an entailment relation between such verbs as *kill* and *die*. The entailment here reflects the fact that *die* and *kill*, which is decomposed into CAUSE TO DIE, share the semantic component DIE.

Although WordNet avoids semantic decomposition in favor of a relational analysis, the semantic relations among verbs in WordNet all interact with entailment.

Lexical Entailment

The principle of lexical inheritance can be said to underlie the semantic relations between nouns, and bipolar oppositions serve to organize the adjectives. Similarly, the different relations that organize the verbs can be cast in terms of one overarching principle, lexical entailment.

In logic, entailment, or strict implication, is properly defined for propositions; a proposition *P* entails a proposition *Q* if and only if there is no conceivable state of affairs that could make *P* true and *Q* false. Entailment is a semantic relation because it involves reference to the states of affairs that *P* and *Q* represent. The term will be generalized here to refer to the relation between two verbs V_1 and V_2 that holds when the sentence *Someone* V_1 logically entails the sentence *Someone* V_2; this use of entailment can be called lexical entailment. Thus, for example, *snore* lexically entails *sleep* because the sentence *He is snoring* entails *He is sleeping*; the second sentence necessarily holds if the the first one does.

Lexical entailment is a unilateral relation: if a verb V_1 entails another verb V_2, then it cannot be that case that V_2 entails V_1. The exception is that where two verbs can be said to be mutually entailing, they must also be synonyms, that is, they must have the same sense. For example, one might say both that *The Germans beat the Argentinians* entails *The Germans defeated the Argentinians*, and that *The Germans defeated the Argentinians* entails *The Germans beat the Argentinians*. However, we find such statements rather unnatural. Negation reverses the direction of entailment: *not sleeping* entails *not snoring*, but *not snoring* does not entail *not sleeping*. The converse of entailment is contradiction: If the sentence *He is snoring* entails *He is sleeping*, then *He is snoring* also contradicts the sentence *He is not sleeping* (Kempson, 1977).

The entailment relation between verbs resembles meronymy between nouns, but meronymy is better suited to nouns than to verbs. To begin with, in order for sentences based on the formula *An x is a part of a y* to be acceptable, both *x* and *y* must be nouns. It might seem that using the nominalizing gerundive form of the verbs would convert them into nouns, and as nouns the HASA relation should apply. For example, Rips and Conrad (1989) obtained consistent results when they asked subjects to judge questions like *Is thinking a part of planning?* vs. *Is planning a part of thinking?* But this change in syntactic category does not overcome fundamental meaning differences between nouns and verbs. Fellbaum and Miller (1990) argue that, first, verbs cannot be taken apart in the same way as nouns, because the parts of verbs are not analogous to the parts of nouns. Most nouns and noun parts have distinct, delimited referents. The referents of verbs, on the other hand, do not have the kind of distinct parts that characterize objects, groups, or substances. Componential analyses have shown that verbs cannot be broken up into referents denoted solely by verbs. And, second, the relations among parts of verbs differ from those found among noun parts. Any acceptable statement about part-relations among verbs always involves the temporal relation between the activities that the two verbs denote. One activity or event is part of another activity or event only when it is part of, or a stage in, its temporal realization.

It is true that some activities can be broken down into sequentially ordered subactivities. For the most part, these are complex activities that are said to be mentally represented as scripts (Schank and Abelson, 1977). They tend not to be lexicalized in English: *eat at a restaurant*, *clean an engine*, *get a medical check-up*, etc. The analysis into lexicalized sub-activities that is possible for these verb phrases is, however, not available for the majority of simple verbs in English. Yet people will also accept part-whole statements involving verb pairs like *drive-ride* and *snore-sleep*. The reason lies in the kinds of entailment that hold between the verbs.

Consider the relation between the verbs *ride* and *drive*. Although neither activity is a discrete part of the other, the two are connected in that when you drive a vehicle, you necessarily also ride in it. The relation between the two activities denoted by these verbs is different from that holding between an activity like *get a medical check-up* and its temporally sequential sub-activities, like *visit (the doctor)* and *undress*. *Riding* and *driving* are carried on simultaneously. Yet most people accept *Riding is a part of driving* and reject *Driving is a part of riding*, even though neither activity can be considered a subactivity of the other.

Consider also the relations among the activities denoted by the verbs *snore*, *dream*, and *sleep*. Snoring or dreaming can be a part of sleeping, in the sense that the two activities are, at least partially, temporally co-extensive: the time that you spend snoring or dreaming is a proper part of the time you spend sleeping. And it is true that when you stop sleeping you also necessarily stop snoring or dreaming.

The differences between pairs like *drive* and *ride* and *snore* and *sleep* are due to the temporal relations between the members of each pair. The activities can be simultaneous (as with *drive* and *ride*) or one can include the other (as with *snore* and *sleep*). For both pairs, engaging in one activity necessitates engaging in the other activity. Therefore, the first activity in each pair entails the second.

The two semantic relations subsumed under lexical entailment that we have considered so far share the feature of temporal inclusion. That is to say, the sets of verbs related by entailment have in common that one member temporally includes the other. A verb V_1 will be said to include a verb V_2 if there is some stretch of time during which the activities denoted by the two verbs co-occur, but no time during which V_2 occurs and V_1 does not. If there is a time during which V_1 occurs but V_2 does not, V_1 will be said to properly include V_2.

Temporal inclusion may go in either direction. Verb pairs like *buy* and *pay* differ from those like *snore* and *sleep* in that whereas *snore* entails *sleep* and is properly included by it, *buy* entails *pay* but properly includes it. That is to say, either the entailing or the entailed verb may properly include the other. In the case of *snore-sleep*, it is the entailed verb *sleep* that properly includes the entailing verb *snore*, whereas in the pair *buy-pay*, the entailing verb *buy* properly includes the entailed verb *pay*.

Our analysis so far yields a simple generalization: if V_1 entails V_2, and if a temporal inclusion relation holds between them, then people will accept a part-whole statement relating V_2 and V_1.

Hyponymy Among Verbs

The sentence frame used to test hyponymy between nouns, *An x is a y*, is not suitable for verbs, because it requires that *x* and *y* be nouns: *to amble is a kind of to walk* is not a felicitous sentence. Even when this formula is used with verbs in the gerundive form, there is a noticeable difference between nouns and verbs. Although people are quite comfortable with statements like *A horse is an animal* or *A spade is a garden tool*, they are likely to reject such statements as *Ambling is walking* or *Mumbling is talking*, where the superordinate is not accompanied by some qualification. The semantic distinction between two verbs is different from the features that distinguish two nouns in a hyponymic relation.

An examination of ‘verb hyponyms’ and their superordinates shows that lexicalization involves many kinds of semantic elaborations across different semantic fields. For example, Talmy’s (1985) analysis of motion verbs treats them as conflations of *move* and such semantic components as MANNER and CAUSE, exemplified by *slide* and *pull*, respectively. To these components could be added SPEED (encoded in *run, stroll*) or the CONVEYANCE of displacement (*bus, truck, bike*). Similarly, English verbs denoting different kinds of *hitting* express the DEGREE OF FORCE used by the agent (*chop, slam, whack, swat, rap, tap, peck*, etc.). Some verbs refer to different degrees of INTENSITY of the action or state (*drowse, doze, sleep*).

Since the aim is to study relations between verbs, rather than between the building blocks that make them up, the many different kinds of elaborations that distinguish a ‘verb hyponym’ from its superordinate have been merged into a manner relation that Fellbaum and Miller (1990) have dubbed troponymy (from the Greek *tropos*, manner or fashion). The troponymy relation between two verbs can be expressed by the formula *To* V_1 *is to* V_2 *in some particular manner*. ‘Manner’ is interpreted here more loosely than in Talmy’s work, for example, and troponyms can be related to their superordinates along many semantic dimensions. Subsets of particular kinds of manners tend to cluster within a given semantic field. Among competition verbs, for example, many troponyms are conflations of the basic verb *fight* with nouns denoting the occasion for, or form of, the fight: *battle, war, tourney, joust, duel, feud,* etc. Troponyms of communication verbs often encode the speaker’s INTENTION or motivation for communicating, as in *examine, confess,* or *preach*, or the MEDIUM of communication: *fax, e-mail, phone, telex*.

Troponymy and Entailment

Troponymy is a particular kind of entailment, in that every troponym V_1 of a more general verb V_2 also entails V_2. Consider the pair *limp-walk*. The verbs in this pair are related by troponymy: *to limp* is also *to walk in a certain manner*; *limp* is a troponym of *walk*. The verbs are also in an entailment relation: the statement *He is limping* entails *He is walking*, and walking can be said to be a part of limping. Unlike the activitites denoted by *snore* and *sleep*, or *buy* and *pay*, the activities referred to by a troponym and its more general superordinate are always temporally co-extensive, in that one must necessarily be walking every instant that one is limping. Troponymy therefore represents a special case of entailment: pairs that are always temporally co-extensive and are related by entailment.

In contrast with pairs like *limp-walk*, a verb like *snore* entails and is included in *sleep*, but is not a troponym of *sleep*; *get a medical check-up* entails and includes *visit the doctor*, but is not a troponym of *visit the doctor*; and *buy* entails *pay*, but is not a troponym of *pay*. The verbs in these pairs are related only by entailment and proper temporal inclusion. The important generalization here is that verbs related by entailment and proper temporal inclusion cannot be related by troponymy. For two verbs to be related by troponymy, the activities they denote must be temporally co-extensive. One can sleep before or after snoring, buying includes activities other than paying, and visiting the doctor is not temporally co-extensive with getting a medical check-up, so none of these pairs are related by troponymy.

The two categories of lexical entailment that have been distinguished so far are related diagrammatically in Figure 1.

Entailment

+Troponymy	**-Troponymy**
(Co-extensiveness)	**(Proper Inclusion)**
limp-walk	*snore-sleep*
lisp-talk	*buy-pay*

Figure 1. Two kinds of entailment with temporal inclusion

Verb Taxonomies

In trying to construct verb taxonomies using the troponymy relation, it became apparent that verbs cannot easily be arranged into the kind of tree structures onto which nouns are mapped. First, within a single semantic field it is frequently the case that not all verbs can be grouped under a single unique beginner; some semantic fields must be represented by several independent trees. Motion verbs, for example, have two top nodes, {*move, make a movement*}, and {*move, travel*}. Verbs of possession can be traced up to the three verbs {*give, transfer*}, {*take, receive*}, and {*have, hold*}; for the most part, their troponyms encode ways in which society has ritualized the transfer of possessions: *bequeath, donate, inherit, usurp, own, stock,* etc. The semantic field containing verbs of bodily care and functions consists of a number of independent hierarchies that form a coherent semantic field by virtue of the fact that most of the verbs (*wash, comb, shampoo, make up; ache, atrophy*) select for the same kinds of noun arguments (body parts). The communication verbs are headed by the verb *communicate* but immediately divide into verbs of verbal and non-verbal (gestural) communication; the former divide further into actions denoting the communication of spoken vs. written language.

Verb hierarchies tend to have a more shallow, bushy structure than nouns; in few cases does the number of hierarchical levels exceed four. Moreover, virtually every verb taxonomy shows what might be called a bulge, that is to say, a level far more richly lexicalized than the other levels in the same hierarchy. Call this layer *L0*, the layer above it *L+1*, and the layer below *L-1*. Certain parallels can be drawn between *L0* and what has been called the basic level in noun hierarchies (Rosch et al., 1976). Not only do most of the verbs in a hierarchy cluster at *L0*, but the troponymy relation between these verbs and their superordinate is semantically richer than between verbs on other levels. Consider, for example, the taxonomy arising from the *L+1* verb *walk*: the superordinate of *walk*, on level *L+2*, is {*move, travel*}; troponyms of *walk*, on level *L0*, are *march, strut, traipse, amble, mosey, slouch*, etc. Although a statement relating *L+1* to *L+2*—*To walk is to move in some manner*—is perfectly acceptable, statements relating *L0* to *L+1*—*To march/strut/traipse/amble . . . is to walk in some manner*—seem more felicitous; these verbs elaborate the concept of *walking* in distinct ways, yet the features of *walking* are still clearly present. *Walk*, on the other hand, seems semantically more remote from its superordinate, *move*.

An alternative way to think about verb taxonomies—one that reflects the prominence of the two levels, *L+1* and *L0*—is in terms of a radial structure, or cluster, with unelaborated *L+1* verbs like *walk, talk, hit*, and *fight* in the center and their troponyms *march, strut; lisp, babble; tap, slam; battle, joust*, etc., clustered around them.

In most hierarchies, *L-1*, the level below the most richly lexicalized one, has few members. For the most part, they tend not to be independently lexicalized, but are compounded from their superordinate verb and a noun or noun phrase. Examples are *goose-step*, a troponym of *march* from the *walk* hierarchy; and *spoonfeed, forcefeed, bottle-feed, breastfeed*, troponyms of {*feed, cause to eat*} from the *ingest* hierarchy among the consumption verbs.

As one descends in a verb hierarchy, the variety of nouns that the verbs on a given level can take as potential arguments decreases. This seems to be a function of the increasing elaboration and meaning specificity of the verb. Thus, *walk* can take a subject referring either to a person or an animal; most troponyms of *walk*, however, are restricted to human subjects. And *goose-stepping* is usually, though not necessarily, done by soldiers; this verb rarely takes *children* or *old people* as arguments. On the other hand, {*move, travel*} can take not only person and animal subjects, but also vehicles, or objects moved by external forces. Similarly, figures or pictures can *communicate* and *talk*; they can even *deceive* or *lie*, but they cannot *fib* or *perjure themselves*, as only human speakers can. A piece of news may *hit*, *touch* or even *grab* you, but it cannot *punch*, *stroke* or *collar* you; only people can be agents of these verbs.

Opposition Relations between Verbs

There is evidence that opposition relations are psychologically salient not only for adjectives, but also for verbs. For example, in the course of teaching foreign languages the author has experienced that students, when given only one member of an antonymous or opposed verb pair, will insist on being told the other member; they believe that it is easier to learn semantically opposed words together. Fellbaum and Chaffin (1990), in an analogy task involving different semantic relations between verbs, asked subjects to generate verbs whose relation to the stimulus matched that of a given pair; they found that subjects were most successful in completing analogies that involved an opposition relation. Moreover, analogies based on opposition relations took the least time to complete. In building the database for verbs, it was found that after synonymy and troponymy, opposition is the most frequently coded semantic relation.

The semantics of opposition relations among verbs is complex. As in the case of adjectives, much of the opposition among verbs is based on the morphological markedness of one member of an opposed pair, as in the pairs *tie/untie* and *appear/disappear/fR.*

Like the semantically similar adjective pairs *weighty/weightless* and *heavy/light*, there are pairs like *fall/rise* and *ascend/descend* that seem identical in meaning, yet are distinguished by the way their members pick out their direct antonyms: *rise/descend* and *ascend/fall* are conceptually opposed, but are not direct antonyms.

Many deadjectival verbs formed with a suffix such as *-en* or *-ify* inherit opposition relations from their root adjectives: *lengthen/shorten*, *strengthen/weaken*, *prettify/uglify*, for example. These are, for the most part, verbs of change and would decompose into BECOME + adjective or MAKE + adjective. As in the case of the corresponding adjectives, these are direct antonyms. Synonyms of these verbs, when they exist, are generally of Latin or Greek origin and tend to be more constrained in the range of their potential arguments, that is to say, they are usually reserved for more specialized uses. Thus, *fortify* is a synonym of *strengthen*, but its opposition relation to *weaken* is conceptual: *fortify/weaken* are indirect antonyms. In short, deadjectival verbs can be represented by the same configuration that Gross and Miller (this volume) describe for adjectives.

As in the case of adjectives, a variety of negative morphological markers attach to verbs to form their respective opposing members. Examples are *tie/untie*, *approve/disapprove*, and *bone/debone*. The semantics of these morphological oppositions is not simple negation. *To untie* is a kind of *undoing*, and the sense of this verb is one of reversing an action; it does not mean *to not tie*. A pair of verbs like *approve/disapprove* are gradables: the two lexicalized terms are points on a scale (of approval, in this case). Gradable verbs, like gradable adjectives, can be modified by degree adverbs, such as *quite*, *rather*, etc. Perhaps the most striking example illustrating that negative morphology is not simple negation is seen in such pairs as *bone/debone* where, despite the lexical opposition induced by the presence of a negative prefix, there is no semantic opposition at all—both verbs refer to the same activity of removing the bones of an animal (Horn, 1989). In some pairs, the marked member cannot be inferred simply from the morphological marker because the opposition derives from the prefixes themselves: *emigrate/immigrate*, *exhale/inhale*, *predate/postdate*.

Other pairs whose members seem to be direct antonyms are *rise/fall* and *walk/run*. Members of these pairs are associated with each other rather than with verbs that are synonyms of their respective opposites and that express the same concept as that opposite. These pairs are illustrative of an opposition relation that is found quite systematically between co-troponyms (troponyms of the same superordinate verb). For example, the motion verbs *rise* and *fall* both conflate the superordinate *move* with a semantic component denoting the direction of motion; they constitute an opposing pair because the direction of motion, upward or downward, is opposed. Similarly, the opposition between *walk* and *run,* two co-troponyms of {*move, travel*}, is due to the opposing manners (slow or fast, respectively) that distinguish each troponym from its superordinate. And the opposition of *nibble* and *gorge,* co-troponyms of *eat,* derives from the opposition between the quantities eaten. Similarly, *breastfeed* and *bottlefeed* are opposites because they refer to two opposite manners of feeding an infant.

Still other pairs illustrate the variety and complexity of verb opposition. Some pairs, called converses, are opposites that are associated with no common superordinate or entailed verb: *give/take*, *buy/sell*, *lend/borrow*, *teach/learn*, etc. They have in common that they occur within the same semantic field: they refer to the same activity, but from the viewpoint of different participants. This fact would lead one to surmise that their strong lexical association is probably due to their frequent co-occurrence in usage.

Most antonymous verbs are stative or change-of-state verbs that can be expressed in terms of attributes. There are many opposition relations among stative verbs: *live/die*, *exclude/include*, *differ/equal*, *wake/sleep*. Opposition relations are also frequent among change verbs. Virtually no other relation (other than synonymy) holds these verbs together. Thus, the organization of this suburb of the lexicon is flat rather than hierarchical—there are no superordinates (except the generic *change* and *be* or *have*), and virtually no troponyms. Change verbs and stative verbs thus have a structure resembling that of the adjectives, with only synonymy and opposition relations.

Opposition and Entailment

Many verb pairs in an opposition relation also share an entailed verb. For example, both *hit* and *miss* entail *aim*, because one must necessarily aim in order to hit or miss. In contrast to the kinds of entailment discussed earlier, these verbs are not related by temporal inclusion. The activities denoted by *hit* (or *miss*) and *aim* occur in a sequential order: in order to either hit or miss, one must have aimed first; aiming is a precondition for both hitting and missing. The relation between the entailing and the entailed verbs here is one of backward presupposition, where the activity denoted by the entailed verb always precedes the activity denoted by the entailing verb in time. Other examples are *fail* and *succeed*, which both entail *try*; and *win* and *lose*, both entailing *play* or *gamble*.

Entailment via backward presupposition also holds between certain verb pairs related by a result or purpose relation, such as *fatten-feed*.

A verb V_1 that is entailed by another verb V_2 via backward presupposition cannot be said to be a part of V_2. Part-whole statements between verbs are possible only when a temporal inclusion relation holds between these verbs.

The set of verbs related by entailment that we have considered so far can be classified exhaustively into two mutually exclusive categories on the basis of temporal inclusion (see Fig.2).

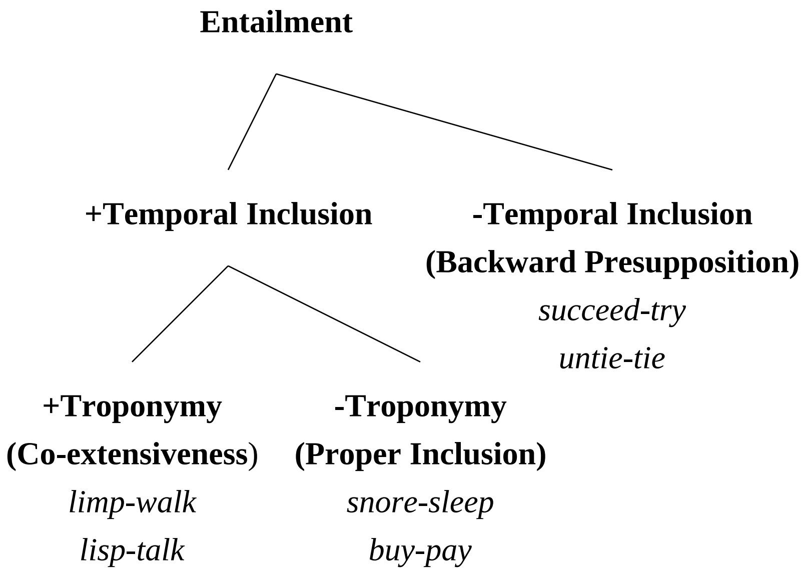

Figure 2. Three kinds of entailment

Opposing verbs like *fail* and *succeed* tend to be contradictories, and, like contradictory adjectives, they do not tolerate degree adverbs. Some opposition relations interact with the entailment relation in a systematic way. Cruse (1986) distinguishes an opposition relation that holds between verb pairs like *damage* and *repair*, and *remove* and *replace*. One member of these pairs, Cruse states, constitutes a ‘‘restitutive.’’ This kind of opposition also always includes entailment, in that the restitutive verb always presupposes what one might call the ‘‘deconstructive’’ one. Many reversive *un-* or *de-* verbs also presuppose their unprefixed, opposed member: in order to *untie* or *unscrew* something, someone must have *tied* or *screwed* it first. Again, these accompanying entailment relations are not marked in WordNet; only the more salient opposition relation is entered.

The Causal Relation

The causative relation picks out two verb concepts, one causative (like *give*), the other what might be called the ‘‘resultative’’ (like *have*). In contrast to the other relations coded in WordNet, the subject of the causative verb usually has a referent that is distinct from the subject of the resultative; the subject of the resultative must be an object of the causative verb, which is therefore necessarily transitive. The causative member of the pair may have its own lexicalization, distinct from the resultative, as in the pair *show* and *see*; sometimes, the members of such a pair differ only by a small variation in their common stem, as in the case of *fell-fall* and *raise-rise*. Although many languages have a means to express causation, not all languages lexicalize the causative member independently; causation is often marked by a morpheme reserved for this function. English does not have many lexicalized causative-resultative pairs, such as *show-see*; it has an analytic, or periphrastic, causative, formed with *cause to/make/let/have/get to*, that is used productively.

It has frequently been pointed out that a periphrastic causative is not semantically equivalent to a lexicalized causative (Fodor, 1970; Shibatani, 1972; and others), but refers to a more indirect kind of causation than the direct, lexicalized form. *Kill* and *cause to die* usually cannot be used interchangeably to refer to the same action, and so are not strictly speaking synonymous expressions of the same concept. For example, Chomsky (1977) notes that you can cause someone to die by having him drive across the country with a pathological murderer, but your action could not properly be called killing. For the purposes of WordNet, such pragmatic considerations have been disregarded.

WordNet recognizes only lexicalized causative-resultative pairs. The synonyms of the members of such a pair inherit the Cause relation, indicating that this relation holds between the entire concept rather than between individual word forms only: the synonyms {*teach, instruct, educate*}, for example, are all causatives of the concept {*learn, acquire knowledge*}. However, unlike entailment, the causation relation is not inherited by the troponyms of each of these concepts: *spoonfeed, indoctrinate,* and *tutor* do not necessarily cause the student to *cram, stuff,* or *memorize*.

Causative verbs have the sense of *cause to be/become/happen/have* or *cause to do*. That is to say, they relate transitive verbs to either states or actions. For example, *give* and *teach* are related via causation to the statives *have* and *know*; *raise* and *feed* are related to the events or actions referred to by *rise* and *eat*. In both cases, causation can be seen as a kind of change. Many verbs clearly have the semantics of such a causative change, but they do not have lexicalized resultatives. The *amuse* and *annoy* subgroup of the psych verbs all refer to causing the experiencer to have an emotion, but only one such causative concept, {*frighten, scare*}, has a lexicalized resultative, {*fear, dread*}.

There are many verbs in English that have both a causative and an anticausative usage. Most of them cluster in the file containing the verbs of change, where many verbs alternate between a transitive causative form and an intransitive anticausative (or unaccusative, or inchoative) form. Here, the surface form of the causative and the anticausative verbs are identical. Examples are the verbs *whiten, grow, break,* and *shrink*. Most anticausative verbs imply either an animate agent or an inanimate cause (*The glass door broke-The storm/The children broke the glass door*). A few verbs are compatible only with an inanimate cause: *Johnnie's teeth rotted—All that candy rotted Johnnie's teeth*, is acceptable, but *His mother rotted Johnnie's teeth* is not.

The causative relation also shows up systematically among the motion verbs: *bounce, roll, blow*, etc., alternate between a causative and an anticausative usage (*She blew a soap bubble in his face vs. The soap bubble blew in his face*). While the causative variants of these verbs usually require an inanimate object, some unergative verbs like *run, jump, gallop, walk, race,* which select for an animate agent, can also have a causative reading, as in the sentences *He raced the horse past the barn* and *The father* *walked his son to school*. (See Levin, 1985, and Pinker, 1989, for a conceptual analysis of these verbs.)

Causation and Entailment

Carter (1976) notes that causation is a specific kind of entailment: if V_1 necessarily causes V_2, then V_1 also entails V_2. He cites the entailment relation between verb pairs like *expel* and *leave*, or *bequeath* and *own*, where the entailing verb denotes the causation of the state or activity referred to by the entailed verb. Like the backward presupposition relation that holds between verbs like *fail/succeed* and *try*, the entailment between verbs like *bequeath* and *own* is characterized by the absence of temporal inclusion.

The causation relation is unidirectional: although *giving* something to somebody causes the recipient to *have* it, for someone to *have* something does not entail that he was *given* it. Similarly, *feeding* somebody causes that person to *eat*, but somebody's *eating* does not entail that someone *feeds* the eater. Except that when the subject of *eat* is not a potentially independent agent, such as a baby or a confined animal, then *eating* does entail the causative act of *feeding*. The direction of the entailment may therefore be reversed in such cases, depending on the semantic features of the verb's subject. But because the entailment depends on specific features of the subject, it can no longer be said to be lexical, that is, it is no longer a relation between two verbs only.

We have now distinguished four different kinds of lexical entailment that systematically interact with the semantic relations we code in WordNet. These four kinds of entailment are related in Figure 3.

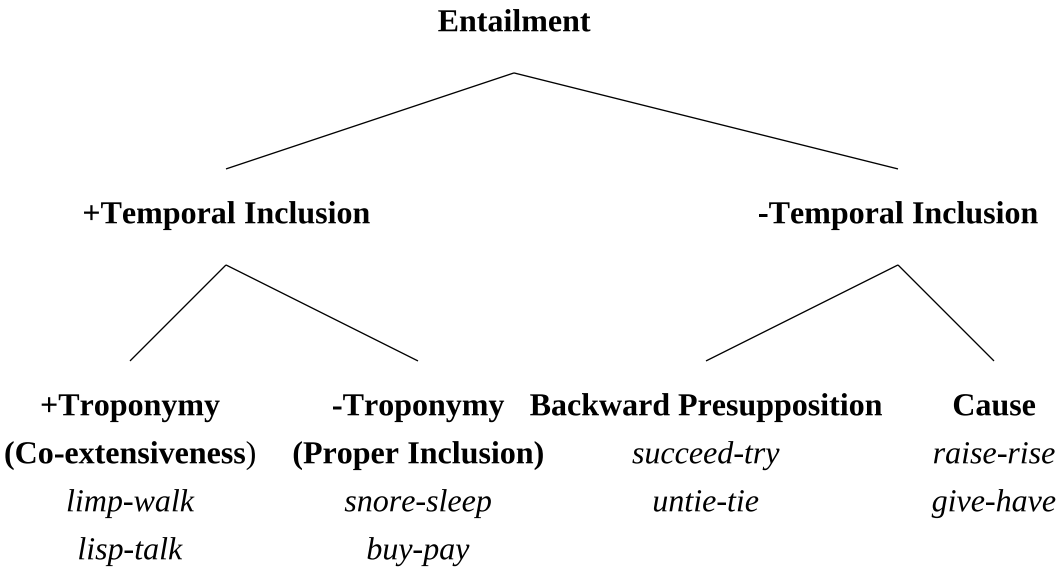

Figure 3. Four kinds of entailment relations among verbs

Syntactic Properties and Semantic Relations

In recent years the lexicon has gained increasing attention from linguists. Verbs in particular have been the subject of much research in pursuit of a theory of lexical knowledge. The work of Levin (1985, 1989) and others has focused on properties of verbs as lexical items that combine with noun arguments to form sentences. This research analyzes the constraints on verbs' argument-taking properties in terms of their semantic makeup, based on the assumption that the distinctive syntactic behavior of verbs and verb classes arises from their semantic components. Pinker (1989) undertakes a fine-grained semantic analysis of verbs that participate in syntactic alternations, claiming that children can discern subtle differences that enable them to distinguish semantically-based verb classes with certain syntactic properties. Gleitman (1990) asserts that children exploit the syntactic-semantic regularities of verbs to infer their meanings on the basis of their syntactic properties.

WordNet was designed to model lexical memory rather than represent lexical knowledge, so it excludes much of a speaker's knowledge about both semantic and syntactic properties of verbs. There is no evidence that the syntactic behavior of verbs (or any other lexical category) serves to organize lexical memory. But there is a substantial body of research (cited in Levin, 1989) showing undeniable correlations between a verb's semantic make-up and its syntax, and the possible implications for children's acquisition of lexical knowledge (Pinker, 1989; Gleitman, 1990).

To cover at least the most important syntactic aspects of verbs, therefore, WordNet includes for each verb synset one or several sentence frames, which specify the subcategorization features of the verbs in the synset by indicating the kinds of sentences they can occur in. This information permits one quickly to search among the verbs for the kinds of semantic-syntactic regularities studied by Levin and others. One can either search for all the synsets that share one or more sentence frames in common and compare their semantic properties; or one can start with a number of semantically similar verb synsets and see whether they exhibit the same syntactic properties. An exploration of the syntactic properties of co-troponyms occasionally provides the basis for distinguishing semantic subgroups of troponyms.

As a case in point, consider verbs like *fabricate* and *compose*, which are members of the creation verb class. Many creation verbs participate in a syntactic alternation that Levin (1989) terms the Material/Product alternation, illustrated by the following examples:

> She wove a rug from the black sheep's wool
> She wove the black sheep's wool into a rug
>
> They molded a head from the clay
> They molded the clay into a head

Some verbs, like *fabricate* and *compose*, which also share membership in the class of creation verbs, do not participate in this syntactic alternation, despite their semantic similarity to verbs like *weave* and *mold*:

The reporter fabricated a story out of the girl's account
*The reporter fabricated the girl's account into a story

She composed a quartet out of the old folk song
*She composed the old folk song into a quartet

In discussing these verbs, Fellbaum and Kegl (1988) point out that the data suggest a need for a fine-grained sub-classification of creation verbs that distinguishes a class of verbs referring to acts of mental creation (such as as *fabricate* and *compose*) from verbs denoting the creation from raw materials (such as *weave* and *mold*). Such a distinction would account for the systematic difference among the verbs in most cases. Thus, Levin (1989) distinguishes these verbs in terms of membership in one of two classes: the BUILD class, which comprises verbs like *bake*, and the CREATE class constituted by such verbs as *compose* and *fabricate*. However, English does not have a lexicalized generic verb denoting the concepts of *create from raw material* and *create mentally*, which would make it possible to capture this generalization by means of different superordinates whose troponyms differ syntactically. But the observation can be formulated in terms of differences in the manner relations that link verbs like *mold* on the one hand, and verbs like *compose* on the other hand, to their common superordinate *create*. This example demonstrates how syntactic differences between apparently similar verbs can be cast in terms of the particular way that the meanings of words are represented in WordNet.

Viewing verbs in terms of semantic relations can also provide clues to an understanding of the syntactic behavior of verbs. Fellbaum and Kegl (1989) studied a class of English verbs that participate in the following transitive-intransitive alternation:

Mary ate a bag of pretzels
Mary ate

Previous analyses of these verbs have explained the alternation in terms of discourse control (Fillmore, 1986) or aspect (Mittwoch, 1982). However, an analysis of the troponyms of the verb *eat* showed that they fall into two syntactic classes: those that must always be used transitively, and those that are always intransitive. The first class includes the verbs *gobble*, *guzzle*, *gulp*, and *devour*; the second class includes verbs like *dine*, *graze*, *nosh*, and *snack*. Fellbaum and Kegl suggest that this syntactic difference is not just a transitivity alternation characteristic of a single verb, but is semantically motivated. They show that English has two verbs *eat*, and that each verb occupies a different position in the network, that is to say, each verb is part of a different taxonomy. Intransitive *eat* has the sense of *eat a meal*. In some troponyms of this verb, such as the denominals *dine*, *breakfast*, *picnic*, and *feast*, the verb *eat* has become conflated with hyponyms of the noun *meal*. These verbs are intransitive because they are all lexicalizations of the verb *eat* that means *eat a meal*. Other intransitive troponyms of this verb are *munch*, *nosh*, and *graze*. Although these verbs are not conflations of *eat* and a noun, they are semantically related in that they refer to eating informal kinds of meals or repasts. By contrast, the transitive verb *eat* has the sense of *ingest in some manner*, and its troponyms all refer to a specific manner of eating: *gobble*, *gulp*, *devour*, etc.[3] Thus,

[3] Kenneth Hale (personal communication) informs us that two *eat* verbs having this semantic distinction are found cross-linguistically.

the semantics of the troponyms in each case provide a classification in terms of two distinct hierarchies matching the syntactic distinction between the two verb groups.

Summary

Different semantic groups of verbs have distinct structures. Some parts can be cast into a taxonomic framework by means of the troponymy relation; this is generally true for verbs of creation, communication, competition, contact, motion, and consumption. The troponymy relation covers a number of different manner relations, one or more of which tend to cluster in specific semantic domains. In some semantic domains, distinct patterns of lexicalization can be discerned; often a base verb conflates with nouns from the corresponding semantic domain. In such verb hierarchies, which tend to be much flatter than noun hierarchies, one level can be distinguished that is more richly lexicalized than the other levels. Stative verbs and verbs of change exhibit an entirely different structure: they tend to be organized in terms of opposition and synonymy relations and they can be mapped into the bipolar clusters that characterize the adjectives.

Comments on Specific Verb Clusters

This section lists particular properties of the several verb files.

1. *Verbs of Bodily Functions and Care*. This relatively small file of approximately 275 synsets contains many verbs referring to actions or events that are not under the control of the argument that functions as their subject (*sweat, shiver, faint, burp, ache, tire, sleep, freeze,* etc.). For this reason these verbs, which are mostly intransitives, have been called unaccusatives, that is, their subjects are thought to be underlyingly objects. Some verbs, like *snort* and *wink*, that have a basic sense referring to an involuntary action, acquire a sense of communication when the action is intended and controlled by the agent. A number of body verbs referring to grooming activities such as *wash*, and *dress* have a reflexive reading in their intransitive form, and are transitive when the action is performed on someone other than the agent, or on a specific body part of the agent.

2. *Verbs of Change*. The verbs of change constitute one of the largest verb files in WordNet (about 750 synsets), owing in part to the fact that the concept of *change* is flexible enough to accommodate verbs whose semantic description makes them unfit for any other semantically coherent group. Dowty (1979) analyzes all verbs as ultimately being composable as stative predicates, which, by means of aspectual connectives and operators such as DO, yield other verb classes. With the exception of the statives, Dowty additionally assigns to all verbs the operator BECOME as part of their lexical make-up. His analysis shows that all (non-stative) verbs can be classified as verbs of change, either as intransitives (stative predicates plus the operator BECOME) or as transitives (with the additional operator DO). An analysis like Dowty's is based on the decomposition of verbs, rather than on a relational analysis of the kind WordNet has undertaken, but it shows that, given the abstract concept of CHANGE, all verbs can be derived from it. In WordNet, this concept has been broken down into several superordinate verbs of change: {*change, alter, vary, modify*}, {*change1, change state*}, {*change2, change by reversal,*

turn, revert}, {*change integrity*}, {*change shape*}, and {*change3, adjust, conform, adapt*}. Most verbs in the change file are derived from these verbs via troponymy. Thus, WordNet characterizes the verbs of change not by the presence of a semantic component CHANGE but by a semantic relation to one of these more specific concepts of change.

The fact that the change file is so large is also due, in part, to the fact that English has productive morphological rules deriving change verbs from adjectives or nouns via affixes such as -*ify* and -*ize*. The derived verbs, such as *humidify* and *magnetize*, generally refer to a state or attribute, and there is usually an opposite state or attribute value expressed by a verb that is identical except for a negative prefix (*dehumidify* and *demagnetize*). Many change verbs are also derived from adjectives via the -*en* suffix; they tend to have an antonymic verb derived from the base adjective's antonym: *weaken-strengthen*, *shorten-lengthen*, etc.

3. *Verbs of Communication*. The communication verbs, which comprise over 710 synsets, include verbs of verbal and nonverbal communication (gesturing); the former are further divided into verbs of speaking and verbs of writing. Most communication verbs involving language, such as *petition* and *hail*, can be classified as troponyms of speech act verbs (Austin, 1962). Verbs of verbal communciation are richly lexicalized in English. They are elaborated in terms of manner of speaking (*lisp, stammer, babble*) or the speaker's intention (*beg, order, thank*). The communication verb file is also rich in denominals (*mandate, appeal, quiz*, and many more). Many sub-areas of lexicalization show where society values communication: politics (*veto, inaugurate*), law (*libel, plead, pardon*), religion (*preach, proselytize, catechize*), education (*teach, examine*), telecommunciations (denominals derived from the nouns denoting the medium of communication, like *fax, telex*, and *e-mail*), to name just a few. This file also contains verbs referring to animal noises (*neigh, moo*, etc.) and verbs of noise production and uttering that have an inanimate source and lack a communicative function (*creak, screech*).

4. *Competition Verbs*. These verbs, grouped into over 200 synsets, cover the semantic areas of sports, games, and warfare. In this file, the *L-1* layer is relatively well lexicalized for many hierarchies, because there are many verbs referring to actions that are specific to games or sports, and that are troponyms of verbs with more general meanings. We find many composite troponyms (*face-off, run-off, counterstrike*) and denominals (*referee, handicap, arm, team, campaign, chicken-fight, duel*). Many of the verbs can be used reciprocally, i.e., they can take a plural subject, and either a surface or an implicit ‘‘each other’’ referring to the two opposing sides (*fight, race*, etc.).

5. *Consumption Verbs*. The consumption file includes, besides verbs of ingesting, verbs of using, exploiting, spending, and sharing, organized into approximately 130 synsets. Many verbs are syntactic unergatives—that is to say, they are either strictly intransitive or they select only cognate objects or their subordinates. For example, *drink*, when it is not used intransitively, takes as its object the cognate noun *drink* (usually with a modifier) or one of its subordinates (*water, beer*, etc.). The analysis of transitive and intransitive verbs like *eat* (Fellbaum & Kegl, 1989) described above shows how transitivity patterns can serve to tease out semantic subgroups of verbs. The consumption

verb file contains other unergatives like *eat*, which can be analyzed in a similar fashion.

6. *Contact Verbs*. This file is the largest verb file, consisting of over 820 synsets. Because most of its verbs are troponyms of very few base verbs, the largest tree structures are to be found in this file. The central verb concepts are {*fasten, attach*}, *cover*, *cut*, and *touch*. Many troponyms are derived via a manner relation encoding the force, intensity, or iteration of the action. For example, the troponyms of *rub* denote different degrees of force (*scrub, wipe, fray, chafe, scour, abrade,* and others). There are many verbs of holding (*grasp, squeeze, grab, pinch, grip*, and others), and touching (*paw, finger, stroke, hit, jab, poke*). Some of the base verbs require an instrument or material; Carter (1976) and Jackendoff (1983) refer to these as ‘‘entailed’’ or ‘‘open’’ arguments, respectively. These arguments are frequently conflated in troponyms together with the base verb characterizing the action. Troponyms of *cut* encode the instrument: (*knife, saw*) or the resultant shape (*cube, slice*); troponyms of *cover* express the material with which an object is covered (*paint, tar, feather*); troponyms of *enclose* encode the container (*box, bag, crate, shroud*). Verbs of removal refer to the removed stuff (*skin, bark, fleece*) or the resultant empty space (*furrow, hole, groove*). Body part denominals indicate what kind of contact action the body part is typically used for: *shoulder* (*support, carry*); *elbow* (*push*); *finger, thumb* (*touch, manipulate*), and so on.

7. *Cognition Verbs*. This file contains verbs denoting various cognitive actions and states, such as *reasoning, judging, learning, memorizing, understanding*, and *concluding*. The organizing relation in this file is troponymy. Some troponyms express kinds of reasoning (*deduce, induce*) or degrees of certainty (*infer, guess, assume, suppose*). The cognition verbs overlap to a large degree with the communication verbs; one verb can refer both to the mental activity of, say, reasoning and judging and to the action of articulating one’s reasoning and judging.

8. *Creation Verbs*. The creation verbs, organized into about 250 synsets, fall into several subgroups that are both semantically and syntactically motivated, but whose superordinates, referring to manners of creation, are not lexicalized: *create by mental act* (*invent, conceive,* etc.); *create by artistic means* (*engrave, illuminate, print*); *create from raw material* (*weave, sew, bake*). Many of these verbs can appear either transitively, where the direct object refers to the creation, or intransitively, where the verb no longer necessarily has the *create* sense but refers only to a manipulation of some material (compare *He sewed* and *He sewed a shirt*) without implying the accomplishment of a creation. Other troponyms of creation verbs are such denominals as *lithograph, fresco,* and *silkscreen*, where the created object has been conflated with the verb.

9. *Motion Verbs*. The motion verbs, grouped into over 500 synsets, derive from two roots: {*move, make a movement*}, and {*move, travel*}. The first sense, exemplified by *shake* and *twist*, is what Miller and Johnson-Laird (1976, p. 529) call "motion-in-place"(1976, p.529) and Pinker calls ‘‘contained’’ motion (1989); the second is the concept of locomotion, as in *run* and *crawl*. Both senses of *move* can also have a transitive causative meaning (this is not true for all of their troponyms, though). Troponyms encode the speed of locomotion (*gallop, race*), the medium of transportation (*canoe, taxi*), and the medium in which the travel takes place (*fly, swim*), as well as other

elaborations.

10. *Emotion or Psych Verbs*. These verbs fall into two grammatically distinct classes: those whose subject is the animate Experiencer and whose object (if there is one) is the Source (*fear, miss, adore, love, despise*); and those whose object is the animate Experiencer and whose subject is the Source (*amuse, charm, encourage, anger*). In both cases, the Source may be either animate or inanimate. If the Source is animate, it may be either intentionally causing the emotion, i.e., it may an Agent, or it may be the unintentional Source of the emotion. This distinction is shown by the ambiguity of *The teacher frightened the children*. Inanimates are of course always unintentional sources: *The skeleton/The cry of the owl scared the children*. Most of the verbs have been structured along the analysis given in Johnson-Laird and Oatley (1989) for nouns, where five basic emotions (*happiness, sadness, fear, anger, disgust*) are posited along with their subordinates; most of these nouns have corresponding verbs. Besides being linked by troponymy, some emotions enter into opposition relations (*love-hate, hope-despair*).

11. *Stative Verbs*. Stative verbs (some 200 synsets) are for the most part verbs of *being* and *having*. Many stative verbs also have non-stative senses that have been placed into other files. For example, verbs like *surround*, *cross*, and *reach* have both a stative sense referring to spatial relations, and a non-stative sense denoting verbs of motion.

Many verbs have the meaning of *be* ADJECTIVE: *equal, suffice, necessitate*, to name just a few. Like adjectives, they usually have opposite terms (*differ, lack, obviate*) and synonyms (*match, cover, require*), but they rarely have superordinates, other than *be* or {*have, feature, have as a feature*}. The stative verb file, therefore, consists of small, semantically independent clusters, and resembles the adjective file.

12. *Perception Verbs*. The top nodes in this file of about 200 synsets are verbs referring to perception by means of the five senses. Their troponyms encode different elaborations. For example, the troponyms of the base verb *see* conflate the intention on the part of the perceiver (*watch, spy, survey*), the circumstances (*witness, discover*), and manner (*gaze, stare, ogle, glance*). Verbs of smelling, for example, take as their subject either experiencer (*snuff, sniff, whiff*) or the source of the perception (*reek, stink*); *smell* and *scent* can be used in either frame. Verbs denoting both the causation and the perception of cutaneous irritation are particularly richly lexicalized: *ache, hurt, prickle, sting, prick, tingle, tickle, scratch, itch, bite*, and others.

13. *Verbs of Possession*. The verbs of possession, organized into almost 300 synsets, are mostly derived from three basic concepts: {*have, hold, own*}, {*give, transfer*}, and {*take, receive*}, which denote the change of possession and its prior or resultant state. Troponyms of these verbs refer to ways in which the transfer of possession takes place in our society: by legal or illegal means, such as inheritance (*bequeath, will, inherit*) vs. theft (*rob, plagiarize, loot*); by formal or informal gifts (*bestow, confer, grant* vs. *beg, bribe, extort*); and by various business transactions (*peddle, scalp, auction, retail*), to name a few of the possible elaborations.

14. *Verbs of Social Interaction*. This file contains well over 400 synsets with verbs from different areas of social life: law, politics, economy, education, family, religion, etc. Many have a specialized meaning, restricted to a particular domain of social life, and

they tend to be monosemous (*impeach, court-martial, moonlight, franchise, gerrymander*, *excommunicate*, to name a few). Frequently, the verbs are denominals conflating a basic verb—often a communication or a cognition verb—and a noun from one of the areas named above, for example, *petition*, *quarrel, charm*, or *veto*.

15. *Weather Verbs*. Weather verbs constitute the smallest verb file (66 synsets), but they are semantically and syntactically distinct. They include mostly verbs like *rain* and *thunder*, which are all intransitives (except for such idiomatic expressions as *It is raining cats and dogs*). They do not select for any arguments; their subject is the semantically empty expletive, *it*. Many of these verbs are derived from their homonymous nouns (*rain, thunder, snow, hail,* etc.).

Conclusion

A relational analysis of English verbs has revealed some of the striking ways in which verbs differ from nouns and adjectives. The relations between verbs are distinct from those between words of other parts of speech; in general, their semantics are considerably more complex. The predominance of different relations and different lexicalization patterns in various semantic domains has been discussed.

Design and Implementation of the WordNet Lexical Database and Searching Software†

Richard Beckwith, George A. Miller, and Randee Tengi

Lexicographers must be concerned with the presentation as well as the content of their work, and this concern is heightened when presentation moves from the printed page to the computer monitor. Printed dictionaries have become relatively standardized through many years of publishing (Vizetelly, 1915); expectations for electronic lexicons are still up for grabs. Indeed, computer technology itself is evolving rapidly; an indefinite variety of ways to present lexical information is possible with this new technology, and the advantages and disadvantages of many possible alternatives are still matters for experimentation and debate. Given this degree of uncertainty, manner of presentation must be a central concern for the electronic lexicographer.

WordNet is a pioneering excursion into this new medium. Considerable attention has been devoted to making it useful and convenient, but the solutions described here are unlikely to be the final word on these matters. It is hoped that readers will not merely note the shortcomings of this work, but will also be inspired to make improvements on it.

One's first impression of WordNet is likely to be that it is an on-line thesaurus. It is true that sets of synonyms are basic building blocks, and with nothing more than these synonym sets the system would have all the power of a thesaurus. When short glosses are added to the synonym sets, it resembles an on-line dictionary that has been supplemented with synonyms for cross referencing (Calzolari, 1988). But WordNet includes much more information than that. In an attempt to model the lexical knowledge of a native speaker of English, WordNet has been given detailed information about relations between word forms and synonym sets. How this relational structure should be presented to a user raises questions that outrun the experience of conventional lexicography.

In developing this on-line lexical database, it has been convenient to divide the work into two interdependent tasks which bear a vague similarity to the traditional tasks of writing and printing a dictionary. One task was to write the source files that contain the basic lexical data — the contents of those files are the lexical substance of WordNet. The second task was to create a set of computer programs that would accept the source

† This is a revised version of "Implementing a Lexical Network" in CSL Report #43, prepared by Randee Tengi. UNIX is a registered trademark of UNIX System Laboratories, Inc. Sun, Sun 3 and Sun 4 are trademarks of Sun Microsystems, Inc. Macintosh is a trademark of Macintosh Laboratory, Inc. licensed to Apple Computer, Inc. NeXT is a trademark of NeXT. Microsoft Windows is a trademark of Microsoft Corporation. IBM is a registered trademark of International Business Machines Corporation. X Windows is a trademark of the Massachusetts Institute of Technology. DECstation is a trademark of Digital Equipment Corporation.

files and do all the work leading ultimately to the generation of a display for the user.

The WordNet system falls naturally into four parts: the WordNet lexicographers' source files; the software to convert these files into the WordNet lexical database; the WordNet lexical database; and the suite of software tools used to access the database. The WordNet system is developed on a network of Sun-4 workstations. The software programs and tools are written using the C programming language, Unix utilities, and shell scripts. To date, WordNet has been ported to the following computer systems: Sun-3; DECstation; NeXT; IBM PC and PC clones; Macintosh.

The remainder of this paper discusses general features of the design and implementation of WordNet. The ''WordNet Reference Manual'' is a set of manual pages that describe aspects of the WordNet system in detail, particularly the user interfaces and file formats. Together the two provide a fairly comprehensive view of the WordNet system.

Index of Familiarity

One of the best known and most important psycholinguistic facts about the mental lexicon is that some words are much more familiar than others. The familiarity of a word is known to influence a wide range of performance variables: speed of reading, speed of comprehension, ease of recall, probability of use. The effects are so ubiquitous that experimenters who hope to study anything else must take great pains to equate the words they use for familiarity. To ignore this variable in a lexical database that is supposed to reflect psycholinguistic principles would be unthinkable.

In order to incorporate differences in familiarity into WordNet, a syntactically tagged index of familiarity is associated with each word form. This index does not reflect all of the consequences of differences of familiarity — some theorists would ask for strength indices associated with each relation — but accurate information on all of the consequences is not easily obtained. The present index is a first step.

Frequency of use is usually assumed to be the best indicator of familiarity. The closed class words that play an important syntactic role are the most frequently used, of course, but even within the open classes of words there are large differences in frequency of occurrence that are assumed to correlate with — or to explain — the large differences in familiarity. The frequency data that are readily available in the technical literature, however, are inadequate for a database as extensive as WordNet. Thorndike and Lorge (1944) published data based on a count of some 5,000,000 running words of text, but they reported their results only for the 30,000 most frequent words. Moreover, they defined a ''word'' as any string of letters between successive spaces, so their counts for homographs are untrustworthy; there is no way to tell, for example, how often *lead* occurred as a noun and how often as a verb. Francis and Kučera (1982) tag words for their syntactic category, but they report results for only 1,014,000 running words of text — or 50,400 word types, including many proper names — which is not a large enough sample to yield reliable counts for infrequently used words. (A comfortable rate of speaking is about 120 words/minute, so that 1,000,000 words corresponds to 140 hours, or about two weeks of normal exposure to language.)

Fortunately, an alternative indicator of familiarity is available. It has been known at least since Zipf (1945) that frequency of occurrence and polysemy are correlated. That is to say, on the average, the more frequently a word is used the more different meanings it will have in a dictionary. An intriguing finding in psycholinguistics (Jastrezembski, 1981) is that polysemy seems to predict lexical access times as well as frequency does. Indeed, if the effect of frequency is controlled by choosing words of equivalent frequencies, polysemy is still a significant predictor of lexical decision times.

Instead of using frequency of occurrence as an index of familiarity, therefore, WordNet uses polysemy. This measure can be determined from an on-line dictionary. If an index value of 0 is assigned to words that do not appear in the dictionary, and if values of 1 or more are assigned according to the number of senses the word has, then an index value can be made available for every word in every syntactic category. Associated with every word form in WordNet, therefore, there is an integer that represents a count (of the Collins *Dictionary of the English Language*) of the number of senses that word form has when it is used as a noun, verb, adjective, or adverb.

A simple example of how the familiarity index might be used is shown in Table 1. If, say, the superordinates of *bronco* are requested, WordNet can respond with the sequence of hypernyms shown in Table 1. Now, if all the terms with a familiarity index (polysemy count) of 0 or 1 are omitted, which are primarily technical terms, the hypernyms of *bronco* include simply: *bronco @→ pony @→ horse @→ animal @→ organism @→ entity*. This shortened chain is much closer to what a layman would expect. The index of familiarity should be useful, therefore, when making suggestions for changes in wording. A user can search for a more familiar word by inspecting the polysemy in the WordNet hierarchy.

WordNet would be a better simulation of human semantic memory if a familiarity index could be assigned to word-meaning pairs rather than to word forms. The noun *tie*, for example, is used far more often with the meaning {*tie, necktie*} than with the meaning {*tie, tie beam*}, yet both are presently assigned the same index, 13.

Lexicographers' Source Files

WordNet's source files are written by lexicographers. They are the product of a detailed relational analysis of lexical semantics: a variety of lexical and semantic relations are used to represent the organization of lexical knowledge. Two kinds of building blocks are distinguished in the source files: word forms and word meanings. Word forms are represented in their familiar orthography; word meanings are represented by synonym sets — lists of synonymous word forms that are interchangeable in some syntax. Two kinds of relations are recognized: lexical and semantic. Lexical relations hold between word forms; semantic relations hold between word meanings.

WordNet organizes nouns, verbs, adjectives and adverbs into synonym sets (*synsets*), which are further arranged into a set of lexicographers' source files by syntactic category and other organizational criteria. Adverbs are maintained in one file, while nouns and verbs are grouped according to semantic fields. Adjectives are divided between two files: one for descriptive adjectives and one for relational adjectives.

Hypernyms of *bronco* and their index values

Word	**Polysemy**
bronco	1
@→ mustang	1
@→ pony	5
@→ horse	14
@→ equine	0
@→ odd-toed ungulate	0
@→ placental mammal	0
@→ mammal	1
@→ vertebrate	1
@→ chordate	1
@→ animal	4
@→ organism	2
@→ entity	3

Table 1

Appendix A lists the names of the lexicographers' source files.

Each source file contains a list of synsets for one part of speech. Each synset consists of synonymous word forms, relational pointers, and other information. The relations represented by these pointers include (but are not limited to): hypernymy/hyponymy, antonymy, entailment, and meronymy/holonymy. Polysemous word forms are those that appear in more than one synset, therefore representing more than one concept. A lexicographer often enters a textual gloss in a synset, usually to provide some insight into the semantics intended by the synonymous word forms and their usage. If present, the textual gloss is included in the database and can be displayed by retrieval software. Comments can be entered, outside of a synset, by enclosing the text of the comment in parentheses, and are not included in the database.

Descriptive adjectives are organized into clusters that represent the values, from one extreme to the other, of some attribute. Thus each adjective cluster has two (occasionally three) parts, each part headed by an antonymous pair of word forms called a head synset. Most head synsets are followed by one or more satellite synsets, each representing a concept that is similar in meaning to the concept represented by the head synset. One way to think of the cluster organization is to visualize a wheel, with each head synset as a hub and its satellite synsets as the spokes. Two or more wheels are logically connected via antonymy, which can be thought of as an axle between wheels.

The Grinder utility compiles the lexicographers' files. It verifies the syntax of the files, resolves the relational pointers, then generates the WordNet database that is used with the retrieval software and other research tools.

Word Forms

In WordNet, a word form is represented as the orthographic representation of an individual word or a string of individual words joined with underscore characters. A string of words so joined is referred to as a collocation and represents a single concept, such as the noun collocation *fountain_pen*.

In the lexicographers' files a word form may be augmented with additional information, necessary for the correct processing and interpretation of the data. An integer sense number is added for sense disambiguation if the same word form appears more than once in a lexicographer file. A syntactic marker, enclosed in parentheses, is added to any adjectival word form whose use is limited to a specific syntactic position in relation to the noun that it modifies. Each word form in WordNet is known by its orthographic representation, syntactic category, semantic field, and sense number. Together, these data make a ‘‘key’’ which uniquely identifies each word form in the database.

Relational Pointers

Relational pointers represent the relations between the word forms in a synset and other synsets, and are either lexical or semantic. Lexical relations exists between relational adjectives and the nouns that they relate to, and between adverbs and the adjectives from which they are derived. The semantic relation between adjectives and the nouns for which they express values are encoded as attributes. The semantic relation between noun attributes and the adjectives expressing their values are also encoded. Presently these are the only pointers that cross from one syntactic category to another. Antonyms are also lexically related. Synonymy of word forms is implicit by inclusion in the same synset. Table 2 summarizes the relational pointers by syntactic category. Meronymy is further specified by appending one of the following characters to the meronymy pointer: **p** to indicate a part of something; **s** to indicate the substance of something; **m** to indicate a member of some group. Holonymy is specified in the same manner, each pointer representing the semantic relation opposite to the corresponding meronymy relation.

Many pointers are reflexive, meaning that if a synset contains a pointer to another synset, the other synset should contain a corresponding reflexive pointer back to the original synset. The Grinder automatically generates the relations for missing reflexive pointers of the types listed in Table 3.

A relational pointer can be entered by the lexicographer in one of two ways. If a pointer is to represent a relation between synsets — a semantic relation — it is entered following the list of word forms in the synset. Hypernymy always relates one synset to another, and is an example of a semantic relation. The lexicographer can also enclose a word form and a list of pointers within square brackets (**[...]**) to define a lexical relation between word forms. Relational adjectives are entered in this manner, showing the lexical relation between the adjective and the noun that it pertains to.

WordNet Relational Pointers

Noun		Verb		Adjective		Adverb	
Antonym	!	Antonym	!	Antonym	!	Antonym	!
Hyponym	~	Troponym	~	Similar	&	Derived from	\
Hypernym	@	Hypernym	@	Relational Adj.	\		
Meronym	#	Entailment	*	Also See	ˆ		
Holonym	%	Cause	>	Attribute	=		
Attribute	=	Also See	ˆ				

Table 2

Reflexive Pointers

Pointer	Reflect
Antonym	Antonym
Hyponym	Hypernym
Hypernym	Hyponym
Holonym	Meronym
Meronym	Holonym
Similar to	Similar to
Attribute	Attribute

Table 3

Verb Sentence Frames

Each verb synset contains a list of verb frames illustrating the types of simple sentences in which the verbs in the synset can be used. A list of verb frames can be restricted to a word form by using the square bracket syntax described above. See Appendix B for a list of the verb sentence frames.

Synset Syntax

Strings in the source files that conform to the following syntactic rules are treated as synsets. Note that this is a brief description of the general synset syntax and is not a formal description of the source file format. A formal specification is found in the manual page **wninput(5)** of the ‘‘WordNet Reference Manual’’.

[1] Each synset begins with a left curly bracket (**{**).

[2] Each synset is terminated with a right curly bracket (**}**).

[3] Each synset contains a list of one or more word forms, each followed by a comma.

[4] To code semantic relations, the list of word forms is followed by a list of relational pointers using the following syntax: a word form (optionally preceded by *"filename***:***"* to indicate a word form in a different lexicographer file) followed by a comma, followed by a relational pointer symbol.

[5] For verb synsets, *"***frames:***"* is followed by a comma separated list of applicable verb frames. The verb frames follow all relational pointers.

[6] To code lexical relations, a word form is followed by a list of elements from [4] and/or [5] inside square brackets (**[...]**).

[7] To code adjective clusters, each part of a cluster (a head synset, optionally followed by satellite synsets) is separated from other parts of a cluster by a line containing only hyphens. Each entire cluster is enclosed in square brackets.

Archive System

The lexicographers' source files are maintained in an archive system based on the Unix Revision Control System (RCS) for managing multiple revisions of text files. The archive system has been established for several reasons — to allow the reconstruction of any version of the WordNet database, to keep a history of all the changes to lexicographers' files, to prevent people from making conflicting changes to the same file, and to ensure that it is always possible to produce an up-to-date version of the WordNet database. The programs in the archive system are Unix shell scripts which envelop RCS commands in a manner that maintains the desired control over the lexicographers' source files and provides a user-friendly interface for the lexicographers.

The **reserve** command extracts from the archive the most recent revision of a given file or files and locks the file for as long as a user is working on it. The **review** command extracts from the archive the most recent revision of a given file or files for the purpose of examination only, therefore the file is not locked. To discourage making changes, review files do not have write permission since any such changes could not be incorporated into the archive. The **restore** command verifies the integrity of a reserved file and returns it to the archive system. The **release** command is used to break a lock placed on a file with the **reserve** command. This is generally used if the lexicographer decides that changes should not be returned to the archive. The **whose** command is used to find out whether files are currently reserved, and if so, by whom.

Grinder Utility

The Grinder is a versatile utility with the primary purpose of compiling the lexicographers' files into a database format that facilitates machine retrieval of the information in WordNet. The Grinder has several options that control its operation on a set of input files. To build a complete WordNet database, all of the lexicographers' files must be processed at the same time. The Grinder is also used as a verification tool to ensure the syntactic integrity of the lexicographers' files when they are returned to the archive system with the **restore** command.

Implementation

The Grinder is a multi-pass compiler that is coded in C. The first pass uses a parser, written in **yacc** and **lex**, to verify that the syntax of the input files conforms to the specification of the input grammar and lexical items, and builds an internal representation of the parsed synsets. Additional passes refer only to this internal representation of the lexicographic data. Pass one attempts to find as many syntactic and structural errors as possible. Syntactic errors are those in which the input file fails to conform to the input grammar's specification, and structural errors refer to relational pointers that cannot be resolved for some reason. Usually these errors occur because the lexicographer has made a typographical error, such as constructing a pointer to a non-existent file, or fails to specify a sense number when referring to an ambiguous word form. Pass one cannot determine structural errors in pointers to files that are not processed together. When used as a verification tool, as from the **restore** command, only pass one is run.

In its second pass, the Grinder resolves all of the semantic and lexical pointers. To do this, the pointers that were specified in each synset are examined in turn, and the target of each pointer (either a synset or a word form in a synset) is found. The source pointer is then resolved by adding an entry to the internal data structure which notes the ''location'' of the target. In the case of reflexive pointers, the target pointer's synset is then searched for a corresponding reflexive pointer. If found, the data structure representing the reflexive pointer is modified to note the ''location'' of its target, the original source pointer. If a reflexive pointer is not found, the Grinder automatically creates one with all the pertinent information.

A subsequent pass through the list of word forms assigns a polysemy index value, or sense count, to each word form found in the on-line dictionary. There is a separate sense count for each syntactic category that the word form is found in. The Grinder's final pass generates the WordNet database.

Internal Representation

The internal representation of the lexicographic data is a network of interrelated linked lists. A hash table of word forms is created as the lexicographers' files are parsed. Lower-case strings are used as keys; the original orthographic word form, if not in lower-case, is retained as part of the data structure for inclusion in the database files. As the parser processes an input file, it calls functions which create data structures for the word forms, pointers, and verb frames in a synset. Once an entire synset had been parsed, a data structure is created for it which includes pointers to the various structures representing the word forms, pointers, and verb frames. All of the synsets from the input files are maintained as a single linked list. The Grinder's different passes access the structures either through the linked list of synsets or the hash table of word forms. A list of synsets that specify each word form is maintained for the purposes of resolving pointers and generating the database's index files.

WordNet Database

For each syntactic category, two files represent the WordNet database — **index.***pos* and **data.***pos*, where *pos* is either **noun**, **verb**, **adj** or **adv** (the actual file names may be different on platforms other than Sun-4). The database is in an ASCII format that is human- and machine-readable, and is easily accessible to those who wish to use it with their own applications. Each index file is an alphabetized list of all of the word forms in WordNet for the corresponding syntactic category. Each data file contains all of the lexicographic data gathered from the lexicographers' files for the corresponding syntactic category, with relational pointers resolved to addresses in data files.

The index and data files are interrelated. Part of each entry in an index file is a list of one or more byte offsets, each indicating the starting address of a synset in a data file. The first step to the retrieval of synsets or other information is typically a search for a word form in one or more index files to obtain all data file addresses of the synsets containing the word form. Each address is the byte offset (in the data file corresponding to the syntactic category of the index file) at which the synset's information begins. The information pertaining to a single synset is encoded as described in the **Data Files** section below.

One shortcoming of the database's structure is that although all the files are in ASCII, and are therefore editable, and in theory extensible, in practice this is almost impossible. One of the Grinder's primary functions is the calculation of addresses for the synsets in the data files. Editing any of the database files would (most likely) create incorrect byte offsets, and would thus derail many searching strategies. At the present time, building a WordNet database requires the use of the Grinder and the processing of all lexicographers' source files at the same time.

The descriptions of the Index and Data files that follow are brief and are intended to provide only a glimpse into the structure, syntax, and organization of the database. More detailed descriptions can be found in the manual page **wndb(5)** included in the ''WordNet Reference Manual''.

Index Files

Word forms in an index file are in lower case regardless of how they were entered in the lexicographers' files. The files are sorted according to the ASCII character set collating sequence and can be searched quickly with a binary search.

Each index file begins with several lines containing a copyright notice, version number and license agreement, followed by the data lines. Each line of data contains the following information: the sense count from the on-line dictionary; a list of the relational pointer types used in all synsets containing the word (this is used by the retrieval software to indicate to a user which searches are applicable); a list of indices which are byte offsets into the corresponding data file, one for each occurrence of the word form in a synset. Each data line is terminated with an end-of-line character.

Data Files

A data file contains information corresponding to the synsets that were defined in the lexicographers' files with pointers resolved to byte offsets in **data.***pos* files.

Each data file begins with several lines containing a copyright notice, version number and license agreement. This is followed by a list of the names of all the input files that were specified to the Grinder, in the order that they were given on the command line, followed by the data lines. Each line of data contains an encoding of the information entered by the lexicographer for a synset, as well as additional information provided by the Grinder which is useful to the retrieval software and other programs. Each data line is terminated with an end-of-line character. In the data files, word forms in a synset match the orthographic representation entered in the lexicographers' files.

The first piece of information on each line is the byte offset, or address, of the synset. This is slightly redundant, since almost any computer program that reads a synset from a data file knows the byte offset that it read it from; however this piece of information is useful when using UNIX utilities like **grep** to trace synsets and pointers without the use of sophisticated software. It also provides a unique ''key'' for a synset, if a user's application requires one. An integer, corresponding to the location in the list of file names of the file from which the synset originated, follows. This can be used by retrieval software to annotate the display of a synset with the name of the originating file, and can be helpful for distinguishing senses. A list of word forms, relational pointers, and verb frames follows. An optional textual gloss is the final component of a data line.

Relational pointers are represented by several pieces of information. The symbol for the pointer comes first, followed by the address of the target synset and its syntactic category (necessary for pointers that cross over into a different syntactic category), followed by a field which differentiates lexical and semantic pointers. If a lexical pointer is being represented, this field indicates which word forms in the source and target synsets the pointer pertains to. For a semantic pointer, this field is 0.

Retrieving Lexical Information

In order to give a user access to information in the database, an interface is required. Interfaces enable end users to retrieve the lexical data and display it via a window-based tool or the command line. When considering the role of the interface, it is important to recognize the difference between a printed dictionary and a lexical database. WordNet's interface software creates its responses to a user's requests on the fly. Unlike an on-line version of a printed dictionary, where information is stored in a fixed format and displayed on demand, WordNet's information is stored in a format that would be meaningless to an ordinary reader. The interface provides a user with a variety of ways to retrieve and display lexical information. Different interfaces can be created to serve the purposes of different users, but all of them will draw on the same underlying lexical database, and may use the same software functions that interface to the database files.

User interfaces to WordNet can take on many forms. The standard interface is an X Windows application, which has been ported to several computer platforms. Microsoft Windows and Macintosh interfaces have also been written. An alternative command line interface allows the user to retrieve the same data, with exactly the same output as the window-based interfaces, although the specification of the retrieval criteria is more cumbersome, and the whole effect is less impressive. Nevertheless, the command line interface is useful because some users do not have access to windowing environments. Shell scripts and other programs can also be written around the command line interface.

The search process is the same regardless of the type of search requested. The first step is to retrieve the index entry located in the appropriate index file. This will contain a list of addresses of the synsets in the data file in which the word appears. Then each of the synsets in the data file is searched for the requested information, which is retrieved and formatted for output. Searching is complicated by the fact that each synset containing the search word also contains pointers to other synsets in the data file that may need to be retrieved and displayed, depending on the search type. For example, each synset in the hypernymic pathway points to the next synset in the hierarchy. If a user requests a recursive search on hypernyms a recursive retrieval process is repeated until a synset is encountered that contains no further pointers.

The user interfaces to WordNet and other software tools rely upon a library of functions that interface to the database files. A fairly comprehensive set of functions is provided: they perform searches and retrievals, morphology, and various other utility functions. Appendix C contains a brief description of these functions. The structured, flexible design of the library provides a simple programming interface to the WordNet database. Low-level, complex, and utility functions are included. The user interface software depends upon the more complex functions to perform the actual data retrieval and formatting of the search results for display to the user. Low-level functions provide basic access to the lexical data in the index and data files, while shielding the programmer from the details of opening files, reading files, and parsing a line of data. These functions return the requested information in a data structure that can be interpreted and used as required by the application. Utility functions allow simple manipulations of the search strings.

The basic searching function, *findtheinfo()*, receives as its input arguments a word form, syntactic category, and search type; *findtheinfo()* calls a low-level function to find the corresponding entry in the index file, and for each sense calls the appropriate function to trace the pointer corresponding to the search type. Most traces are done with the function *traceptrs()*, but specialized functions exist for search types which do not conform to the standard hierarchical search. As a synset is retrieved from the database, it is formatted as required by the search type into a large output buffer. The resulting buffer, containing all of the formatted synsets for all of the senses of the search word, is returned to the caller. The calling function simply has to print the buffer returned from *findtheinfo()*.

This general search and retrieval algorithm is used in several different ways to implement the user interfaces to WordNet. Search types vary by syntactic category but correspond to the relational pointers listed in Table 2. Hierarchical searches may be performed on all relational pointers except for antonyms and ‘‘also see’’. In addition, a call to *findtheinfo()* may retrieve polysemy information, verb sentence frames, or noun coordinate terms (those with the same hypernym as the search string).

The searching function does not perform morphological operations; therefore calls to *findtheinfo()* are made from within a loop that calls *morphstr()* to translate the search string into one or more base forms before calling the searching function.

X Windows Interface

An attempt is made here to give the reader an idea of the look and feel of the X Windows interface to the WordNet database. The Microsoft Windows and Macintosh interfaces are very similar. The command line interface provides the same functions, but the user must specify the search string and search type, as well as other options, on the command line. The command line interface allows multiple searches on a search string with a single command, but a separate command line must be constructed for each search word.

The command **xwn** runs the **xwordnet** program in the background, freeing up the window from which it was started for other tasks. The *xwordnet* window provides full access to the WordNet database. The standard X Windows mouse functions are used to open and close the *xwordnet* window, move the window, and change its size. Help on the general operation of **xwordnet** can be obtained by pressing the middle mouse button with the cursor in the top part of the window.

Searching the Database

The top part of the *xwordnet* window provides a buffer for entering a search string and buttons corresponding to syntactic categories and options. Below this area a status line indicates which type of search is being displayed in the large buffer below.

To search the WordNet database, a user moves the cursor into the large, horizontal box below **"Enter Search Word:"** and enters a search string, followed by a carriage return. A single word, hyphenated word, or collocation may be entered. A highlighted button indicates each syntactic category in the WordNet database that contains the search string. If the search string is not present exactly as typed (except for case, which is ignored), a morphological process is applied to the search string in an attempt to automatically generate a form that is present in WordNet. See the section on **Morphy** for a discussion of this process.

Holding any mouse button on a highlighted part-of-speech button reveals a pull-down menu of searches specific to that syntactic category. All of the searches available for the search string are highlighted. The user selects a search by scrolling down with the mouse until the desired search type is in reverse video, then releasing the mouse button. The retrieval is then performed and the formatted results are displayed in the lower window. The status line shows the type of search that was selected.

Although most searches return very quickly, the WordNet hierarchies can be quite deep and broad, and some retrievals can take a long time. While a search is running, the mouse pointer displays as a watchface when the mouse is in the upper part of the window (above the output buffer), and the message **Searching...** is displayed in the output buffer.

By default, all of the senses found in WordNet that match the selected search are displayed. The search may be restricted to one or more specific senses by entering a comma-separated list of sense numbers in the **"Sense Number:"** box. These numbers are used for one search only, and the box is cleared after the search is completed.

Options

The **Options** menu displays a list of options that are not directly associated with WordNet searches. The *Help*, *Textual Gloss*, and *Log* options are toggles. *Help* and *Log* are initially **Off**, and *Textual Gloss* is initially **On**. An option is toggled by highlighting the option and releasing the mouse button. The following options are available:

[1] The *Help* option is used to display information that is helpful in understanding the search results. The help information is displayed in the output buffer before the search results.

[2] Many WordNet synsets have a textual gloss which often provides an explanation of what the synset represents. The *Textual Gloss* option controls this display.

[3] In addition to being viewed in the output buffer, search results may be appended to a file. When the *Log* option is **On**, search results are appended to the file named when the option is displayed. By default this file is **wnoutput.log**. If the **WNLOG** environment variable is set, the filename is the value of the variable with **.log** appended.

[4] *Display license* allows a user to view the WordNet copyright notice, version number, and license.

[5] Selecting *Quit* exits **xwordnet**.

Output

The output of a WordNet search is intended to be self-explanatory, given that the user knows what type of search was requested. Visual cues, such as indentation to represent levels in retrieved hierarchies, are relied upon to aide a user in interpreting the formatted search results. The complex nature of the adjective structure, unfortunately, makes for less straightforward output of retrieved adjective synsets. In an attempt to clarify the display of adjectival information, direct antonyms, which are generally represented only by head synsets, are always displayed together. This allows a user to distinguish head synsets from satellite synsets, as well as different senses of a head synset.

The output of a search is displayed in the large buffer below the status line. Both horizontal and vertical scroll bars are used to view data that exceeds the window’s borders. The output consists of an ordinal sense number (simply indicating position in the list of senses), followed by a line with the synset that the search string is in, followed by the search results. Each line of search output is preceded by a marker and the synset containing the requested information. If a search traverses more than one level of the tree, then successive lines are indented by spaces corresponding to its level in the hierarchy. If a search doesn't apply to all senses of the search string, the search results are headed by a string such as:

2 of 5 senses of table

When ‘‘**Sample sentences for verb _____**’’ is selected, verb frames that are acceptable for all words in a synset are preceded by the string ‘‘*>’’. If a frame is acceptable for the search string only, it is preceded by the string ‘‘=>’’.

When an adjective is printed, its direct antonym, if it has one, is also printed in parentheses. Since adjectives can be in either head synsets, satellite synsets, or both, any head synsets that the word appears in are printed first, followed by all of the satellite synsets that the word appears in, with an indication of the head synset that the adjective is a satellite of. When the search string is in a head synset, all of the head synset's satellites are also displayed. The position of an adjective in relation to the noun may be restricted to the prenominal, postnominal, or predicative position. Where present, these restrictions are noted in parentheses.

When an adverb is derived from an adjective, the specific adjectival sense on which it is based is printed, along with the relevant adjective synset. If the adjective synset indicated is a satellite synset, then the pertinent head synset is printed following the satellite synset.

Morphy

Many dictionaries hang their information on uninflected headwords without separate listings for inflectional (or many derivational) forms of the word. In a printed dictionary, that practice causes little trouble; with a few highly irregular exceptions, morphologically related words are generally similar enough in spelling to the reference form that the eye, aided by boldface type, quickly picks them up. In an electronic dictionary, on the other hand, when an inflected form is requested, the response is likely to be a frustrating announcement that the word is not in the database; users are required to know the reference form of every word they want to look up. In WordNet, only base forms of words are generally represented. In order to spare users the trouble of affix stripping, and to assist with the creation of programs that use WordNet to automatically process natural language texts, the WordNet software suite includes functions that give WordNet some intelligence about English morphology. At the present time no morphological processes are performed on adverbs.

The WordNet morphological processing functions, Morphy, handle a wide range of morphological transformations. Morphy uses two types of processes to try to convert a word form into a form that is found in the WordNet database. There are lists of inflectional endings, based on syntactic category, that can be detached from individual words in an attempt to find a form of the word that is in WordNet. There are also exception lists for each syntactic category in which a search for an inflected form may be done. Morphy tries to use these two processes in an intelligent manner to translate the word form passed to the form found in WordNet. Morphy first checks for exceptions, then uses the rules of detachment.

The Morphy functions are part of the WordNet library and are used by the retrieval software and various applications. The primary interface function is passed a string (a word form or collocation) and a syntactic category. Since some words, such as *axes* can have more than one base form (*axe* and *axis*), Morphy is set up to work in the following manner. The first time that Morphy is called with a specific string, it returns a base form. For each subsequent lookup of the same string, Morphy returns an alternative base form. Whenever Morphy cannot perform a transformation, NULL is returned.

Exception Lists

There is one exception list for each syntactic category (except adverbs). The exception lists contain the morphological transformations for words that are not regular and therefore cannot be processed in an algorithmic manner. Each line of an exception list contains an inflected form of a word, followed by one or more base forms of the word. The list is kept in alphabetical order and a binary search is used to find words in these lists.

Single Words

In general, single words are relatively easy to process. Morphy first looks for the word form in the exception list. If it is found, then the first base form is returned. Subsequent lookups for the same word form return alternative base forms, if present. A NULL is returned when there are no more base forms of the word.

If the word is not found in the exception list corresponding to the syntactic category, then an algorithmic process that looks for a matching suffix is applied. If a matching suffix is found, a corresponding ending is applied, if necessary, and WordNet is consulted to see if the resulting word is found in WordNet. Refer to Table 4 for a list of suffixes and endings for each syntactic category.

Collocations

As opposed to single words, collocations can be quite challenging to transform into a base form that is present in WordNet. In general, only base forms of words, even those comprising collocations such as *attorney general*, are stored in WordNet. Transforming the collocation *attorneys general* is then simply a matter of finding the base forms of the individual words comprising the collocation. This usually works for nouns, therefore non-conforming nouns, such as *customs duty* are presently entered in the noun exception list (a transformation on each word results in the base form *custom duty*, which is not in WordNet).

Verb collocations that have prepositions, such as *stand in line*, are more difficult. As with single words, the exception list is searched first. If the collocation is not found, special code in Morphy determines whether a verb collocation has a preposition in it. If it does, the following process is applied to try to find the base form. It is assumed that the first word in the collocation is a verb and that the last word is a noun. The algorithm then builds a search string with the base forms of the verb and noun, leaving the

Morphy Suffixes and Endings

Noun		Verb		Adjective	
Sufffix	Ending	Suffix	Ending	Suffix	Ending
s		s		er	
ses	s	ies	y	est	
xes	x	es	e	er	e
zes	z	es		est	e
ches	ch	ed	e		
shes	sh	ed			
		ing	e		
		ing			

Table 4

remainder of the collocation (usually just the preposition, but more words may be involved) in the middle. For example, passed *standing in lines*, the database search would be performed with *stand in line*, which is found in WordNet, and therefore returned from Morphy. If a verb collocation does not contain a preposition, then the base form of each word in the collocation is found and WordNet is searched for the resulting string.

Hyphenation

Hyphenation also presents special difficulties when searching WordNet. It is often a subjective determination whether a word is hyphenated, is closed up, or is a collocation of several words, and which of the various forms are entered into WordNet. When Morphy breaks a string into ‘‘words’’, it looks for both spaces and hyphens as delimiters.

Future Work

Since many noun collocations contains prepositions, such as *line of products*, an algorithm similar to that used for verbs should be written for nouns. In the present scheme, if Morphy is passed *lines of products*, the search string becomes *line of product*, which is not in WordNet. Morphy should also be able to work in both directions — when passed a base form, it should be possible to obtain inflected forms of the word.

Appendix A
Lexicographers' Files

noun.Tops	unique beginners for nouns
noun.act	nouns denoting acts or actions
noun.animal	nouns denoting animals
noun.artifact	nouns denoting man-made objects
noun.attribute	nouns denoting attributes of people and objects
noun.body	nouns denoting body parts
noun.cognition	nouns denoting cognitive processes and contents
noun.communication	nouns denoting communicative processes and contents
noun.event	nouns denoting natural events
noun.feeling	nouns denoting feelings and emotions
noun.food	nouns denoting foods and drinks
noun.group	nouns denoting groupings of people or objects
noun.location	nouns denoting spatial position
noun.motive	nouns denoting goals
noun.object	nouns denoting natural objects (not man-made)
noun.person	nouns denoting people
noun.phenomenon	nouns denoting natural phenomena
noun.plant	nouns denoting plants
noun.possession	nouns denoting possession and transfer of possession
noun.process	nouns denoting natural processes
noun.quantity	nouns denoting quantities and units of measure
noun.relation	nouns denoting relations between people or things or ideas
noun.shape	nouns denoting two and three dimensional shapes
noun.state	nouns denoting stable states of affairs
noun.substance	nouns denoting substances
noun.time	nouns denoting time and temporal relations

verb.body	verbs of grooming, dressing and bodily care
verb.change	verbs of change of size, temperature, intensity, etc.
verb.cognition	verbs of thinking, judging, analyzing, doubting, etc.
verb.communication	verbs of telling, asking, ordering, singing, etc.
verb.competition	verbs of fighting, athletic activities, etc.
verb.consumption	verbs of eating and drinking
verb.contact	verbs of touching, hitting, tying, digging, etc.
verb.creation	verbs of sewing, baking, painting, performing, etc.
verb.emotion	verbs of feeling
verb.motion	verbs of walking, flying, swimming, etc.
verb.perception	verbs of seeing, hearing, feeling, etc.
verb.possession	verbs of buying, selling, owning, and transfer
verb.social	verbs of political and social activities and events
verb.stative	verbs of being, having, spatial relations
verb.weather	verbs of raining, snowing, thawing, thundering, etc.
adj.all	all adjective clusters
adj.pert	relational adjectives (pertainyms)
adv.all	all adverbs

Appendix B
Verb Sentence Frames

1 Something ----s
2 Somebody ----s
3 It is ----ing
4 Something is ----ing PP
5 Something ----s something Adjective/Noun
6 Something ----s Adjective/Noun
7 Somebody ----s Adjective
8 Somebody ----s something
9 Somebody ----s somebody
10 Something ----s somebody
11 Something ----s something
12 Something ----s to somebody
13 Somebody ----s on something
14 Somebody ----s somebody something
15 Somebody ----s something to somebody
16 Somebody ----s something from somebody
17 Somebody ----s somebody with something
18 Somebody ----s somebody of something
19 Somebody ----s something on somebody
20 Somebody ----s somebody PP
21 Somebody ----s something PP
22 Somebody ----s PP
23 Somebody's (body part) ----s
24 Somebody ----s somebody to INFINITIVE
25 Somebody ----s somebody INFINITIVE
26 Somebody ----s that CLAUSE
27 Somebody ----s to somebody
28 Somebody ----s to INFINITIVE
29 Somebody ----s whether INFINITIVE
30 Somebody ----s somebody into V-ing something
31 Somebody ----s something with something
32 Somebody ----s INFINITIVE
33 Somebody ----s VERB-ing
34 It ----s that CLAUSE
35 Something ----s INFINITIVE

Appendix C
Library Functions

Function Name	Input Arguments	Process Performed	Information Returned
findtheinfo findtheinfo_ds	word syntactic category search type sense number	Search of database	Formatted search results in buffer or data structure
traceptrs traceptrs_ds	synset structure pointer type syntactic category depth	Trace pointer hierarchy and place formatted synsets in output buffer or data structure	None or data structure
tracecoords	see traceptrs	Trace coordinate terms	None
traceinherit	see traceptrs	Trace meronyms	None
traceadjant	synset structure	Trace adjective antonyms	None
is_defined	word syntactic category	Find all possible searches for word	Bit mask with one bit set for each possible search type
index_lookup getindex	word syntactic category	Find word in index file	Data structure for index entry
read_synset	syntactic category offset word	Read synset from data file	Data structure for synset
free_syns	synset list pointer	Free synsets from findtheinfo_ds	None
morphstr morphword	word syntactic category	Morphology on collocation or single word	Base form of word or NULL if none
wninit	None	Initialize library functions	None
bin_search	word file descriptor	Binary search algorithm	Line from file containing word or NULL if not found
cntwords	collocation char	Count number of words in string separated by char	Integer number of words in collocation
strtolower	string	Convert string to lower case	Lower case string
strsubst	string ‘from’ char ‘to’ char	Replace all occurrences of ‘from’ with ‘to’ in string	Modified string
getptrtype	pointer name	Convert string to pointer type	Integer pointer type
getpos	string	Convert string to syntactic category	Integer syntactic category
getsstype	string	Convert string to synstet type	Integer synset type

References

Aronoff, M. 1976. *Word Formation in Generative Grammar*. Cambridge, Mass.: MIT Press.

Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Clarendon Press.

Bartning, Inge. 1980. *Syntaxe et sémantique des pseudo-adjectifs français.* Stockholm: Almqvist and Wiksell.

Beckwith, R., Fellbaum, C., Gross, D., and Miller, G. A. (in press). ‘WordNet: A Lexical Database Organized on Psycholinguistic Principles’ in Zernik, U. (ed.). *Using On-line Resources to Build a Lexicon*. Hillsdale, N.J.: Erlbaum.

Berlin, B., Breedlove, D., and Raven, P. H. 1966. ‘Folk Taxonomies and Biological Classification.’ *Science* 154: 273-275.

Berlin, B., Breedlove, D., and Raven, P. H. 1973. ‘General Principles of Classification and Nomenclature in Folk Biology.’ *American Anthropologist* 75: 214-242.

Berlin, B., and Kay, P. 1969. *Basic Color Terms: Their Universality and Evolution*. Berkeley and Los Angeles: University of California Press.

Bever, T. G., and Rosenbaum, P. S. 1970. ‘Some Lexical Structures and Their Empirical Validity’ in Jacobs, R. A., and Rosenbaum, P. S. (eds.). *Readings in English Transformational Grammar*. Waltham, Mass.: Ginn.

Bierwisch, M. 1967. ‘Some Semantic Universals of German Adjectives.’ *Foundations of Language* 3: 1-36.

Bierwisch, M. 1989. ‘The Semantics of Gradation’ in Bierwisch, M., and Lang, E. (eds.). *Dimensional Adjectives: Grammatical Structure and Conceptual Interpretation*. Berlin: Springer-Verlag.

Bolinger, Dwight. 1967. Adjectives in English: attribution and predication. *Lingua*: 18.1-34.

Calzolari, N. 1988. ‘The Dictionary and the Thesaurus Can Be Combined’ in Evens, M. (ed.), *Relational Models of the Lexicon: Representing Knowledge in Semantic Networks*. Cambridge: Cambridge University Press.

Caramazza, A., and Berndt, R. S. 1978. ‘Semantic and Syntactic Processes in Aphasia: A Review of the Literature.’ *Psychological Bulletin* 85: 898-918.

Carnap, R. 1947. *Meaning and Necessity*. Chicago: University of Chicago Press.

Carter, R. 1976. ‘Some Constraints on Possible Words.’ *Semantikos* 1: 27-66.

Chafe, W. 1970. *Meaning and the Structure of Language*. Chicago: University of Chicago Press.

Chaffin, R., Hermann, D. J., and Winston, M. 1988. ‘An Empirical Taxonomy of Part-Whole Relations: Effects of Part-Whole Relation Type on Relation Identification.’ *Language and Cognitive Processes* 3: 17-48.

Charles, W. G., and Miller, G. A. 1989. ‘Contexts of Antonymous Adjectives.’ *Applied Psycholinguistics* 10: 357-375.

Chomsky, N. 1972. *Studies on Semantics in Generative Grammar.* The Hague: Mouton.

Chomsky, N. 1977. ‘Deep Structure, Surface Structure, and Semantic Interpretation’ in Steinberg, D., and Jakobovits, L. (eds). *Semantics: An Interdisciplinary Reader.*

Cambridge: Cambridge University Press.
Cliff, N. 1959. ‘Adverbs as Multipliers.’ *Psychological Review* 66: 27-44.
Collins, A. M., and Quillian, M. R. 1969. ‘Retrieval Time From Semantic Memory.’ *Journal of Verbal Behavior and Verbal Learning* 8: 240-247.
Conrad, C. 1972. ‘Cognitive Economy in Semantic Memory.’ *Journal of Experimental Psychology* 55: 75-84.
Cruse, D. A. 1986. *Lexical Semantics*. New York: Cambridge University Press.
Deese, J. 1964. ‘The Associative Structure of Some English Adjectives.’ *Journal of Verbal Learning and Verbal Behavior* 3: 347-357.
Deese, J. 1965. *The Structure of Associations in Language and Thought*. Baltimore: Johns Hopkins Press.
Dowty, D. 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.
Dunker, K. 1945. ‘On Problem Solving.’ *Psychological Monographs,* 58: Whole No. 270.
Fahlman, S. E. 1979. *NETL: A System for Representing and Using Real-World Knowledge*. Cambridge, Mass.: MIT Press.
the Organization of Verbs in the Mental Lexicon.’ In *Proceedings of the 12th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum.
Fellbaum, C., and Kegl, J. 1988. ‘Taxonomic Hierarchies in the Verb Lexicon.’ Presented at the EURALEX Third International Congress, Budapest, Hungary.
Fellbaum, C., and Kegl, J. 1989. ‘Taxonomic Structures and Cross-Category Linking in the Lexicon’ in de Jong, K., and No, Y. (eds.). *Proceedings of the Sixth Eastern State Conference on Linguistics*. Columbus, Ohio: Ohio State University
Fellbaum, C., and Miller, G. A. (in press). ‘Folk Psychology or Semantic Entailment? A Reply to Rips and Conrad.’ *Psychological Review*.
Fillenbaum, S., and Jones, L. V. 1965. ‘Grammatical Contingencies in Word Association.’ *Journal of Verbal Learning and Verbal Behavior* 4: 248-255.
Fillmore, C. J. 1968. ‘The Case for Case’ in Bach, E., and Harms, R. (eds.). *Universals in Linguistic Theory*. New York: Holt, Reinhart, and Winston.
Fillmore, C. J. 1986. ‘Pragmatically Controlled Zero Anaphora’ in *Proceedings of the Twelfth Annual Meeting of the Berkeley Linguistics Society*.
Fodor, J. 1970. ‘Three Reasons for Not Deriving ‘‘Kill’’ From ‘‘Cause To Die.’’ ’ *Linguistic Inquiry* 1: 429-438.
Francis, W. N., and Kučera, H. 1982. *Frequency Analysis of English Usage: Lexicon and Grammar*. Boston: Houghton Mifflin.
Gardner, H. 1973. ‘The Contribution of Operativity to Naming Capacity in Aphasic Patients.’ *Neuropsychologia* 11: 213-220.
Garrett, M. F. 1982. ‘Production of Speech: Observations from Normal and Pathological Language Use’ in A. Ellis (ed.). *Normality and Pathology in Cognitive Functions*. London: Academic Press.
Gentner, D., and France, I. 1988. ‘The Verb Mutability Effect: Studies of the Combinatorial Semantics of Nouns and Verbs’ in Small, S., Cottrell, G., and Tanenhaus, M. (eds.). *Lexical Ambiguity Resolution*. Los Altos, Calif.: Morgan Kaufmann.

Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

Gleitman, L. 1990. 'The Structural Sources of Verb Meanings.' *Language Acquisition* 1: 1-55.

Gross, D., Fischer, U., and Miller, G. A. 1989. 'The Organization of Adjectival Meanings.' *Journal of Memory and Language* 28: 92-106.

Gross, D., Kegl, J., Gildea, P., and Miller, G. A. 1989. *A Coded Corpus and Commentary on Children's Dictionary Strategies*. Princeton, N.J.: Cognitive Science Laboratory Report 39.

Gruber, J. 1976. *Lexical Structures in Syntax and Semantics*. New York: North-Holland.

Heider, E. R. 1972. 'Universals in Color Naming and Memory.' *Journal of Experimental Psychology* 93: 10-20.

Heider, E. R., and Olivier, D. C. 1972. 'The Structure of Color Space in Naming and Memory for Two Languages.' *Cognitive Psychology* 3: 337-354.

Hoffmann, J., and Ziessler, C. 1983. 'Objektidentifikation in künstlichen Begriffshierarchien.' *Zeitschrift für Psychologie* 191: 135-167.

Horn, L. 1989. 'Morphology, Pragmatics, and the *un-* Verb' in Powers, J., and de Jong, K. (eds.). *Proceedings of the Fifth Eastern State Conference on Linguistics*. Columbus, Ohio: Ohio State University.

Jackendoff, R. S. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, Mass.: MIT Press.

Jackendoff, R. S. 1976. 'Toward an Explanatory Semantic Representation.' *Linguistic Inquiry* 7: 89-150.

Jackendoff, R. S. 1983. *Semantics and Cognition*. Cambridge, Mass.: MIT Press.

Jastrezembski, J. E. 1981. 'Multiple Meanings, Number of Related Meanings, Frequency of Occurrence, and the Lexicon.' *Cognitive Psychology* 13: 278-305.

Johnson-Laird, P. N., and Oatley, K. 1989. 'The Language of Emotion: An Analysis of a Semantic Field.' *Cognition and Emotion* 3: 81-123.

Justeson, John S., and Katz, Slava M. 1991a. Co-occurrences of Antonymous Adjectives and Their Contexts. *Computational Linguistics*: 17.1-19.

Justeson, John S., and Katz, Slava M. 1991b. Redefining Antonymy: The Textual Structure of a Semantic Relation. Proceedings of the Seventh Annual Conference of the UW Centre for the New OED, 138-154. Waterloo, Canada.

Justeson, Johm S., and Katz, Slava, M. 1993. Principled Disambiguation: Discriminating Adjective Senses with Modified Nouns. Ms., IBM T. Watson Research Center, Yorktown Heights, NY.

Katz, J. J. 1964. 'Semantic Theory and the Meaning of ''Good.''' *Journal of Philosophy* 61: 739-766.

Katz, J. J. 1972. *Semantic Theory*. New York: Harper and Row.

Katz, J. J., and Fodor, J. 1963. 'The Structure of a Semantic Theory.' *Language* 39: 170-210.

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

Keil, F. C. 1983. 'On the Emergence of Semantic and Conceptual Distinctions.' *Journal of Experimental Psychology: General* 112: 357-385.

Kempson, R. 1977. *Semantic Theory*. Cambridge: Cambridge University Press.

Lakoff, G. 1970. *Irregularity in Syntax*. New York: Holt, Reinhart, and Winston.

Levi, J. N. 1978. *The Syntax and Semantics of Complex Nominals*. New York: Academic Press.

Levin, B. 1985. 'Introduction' in Levin, B. (ed.). *Lexical Semantics in Review*. Cambridge, Mass.: Center for Cognitive Science, MIT.

Levin, B. 1989. 'Towards a Lexical Organization of English Verbs.' Ms., Evanston: Northwestern University.

Lyons, J. 1977. *Semantics*. 2 vols. New York: Cambridge University Press.

Marchand, H. 1969. *The Categories and Types of Present-Day English Word-Formation*, 2nd ed. Munich: C. H. Beck.

McCawley, J. D. 1968. 'Lexical Insertion in a Transformational Grammar Without Deep Structure' in Darden, B. J., Bailey, C-J. N., and Davison, A. (eds.). *Papers from the fourth regional meeting, Chicago Linguistic Society*. Chicago, Ill.: Department of Linguistics, University of Chicago. Pp. 71-80.

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 Dictionary. Waterloo, Canada: University of Waterloo.

Miller, G. A. 1986. 'Dictionaries in the Mind.' *Language and Cognitive Processes* 1: 171-185.

Miller, G. A. (in press). 'Lexical Echoes of Perceptual Structure' in *The Perception of Structure*, in honor of W. R. Garner. Washington, DC: American Psychological Association.

Miller, G. A., and Charles, W. (in press). 'Contextual Correlates of Semantic Similarity.' *Language and Cognitive Processes*.

Miller, G. A., and Fellbaum, C. (submitted). 'Semantic Networks of English.' *Cognition*.

Miller, G. A., and Gildea, P. M. 1987. 'How Children Learn Words.' *Scientific American* 257: No. 3: 94-99.

Mittwoch, A. 1982. 'On the Difference Between *Eating* and *Eating Something*: Activities vs. Accomplishments.' *Linguistic Inquiry* 13: 113-122.

A New Dictionary Compiled on Historical Principles. 1928. Oxford: Oxford University Press. (OED.)

Murphy, Gregory L., and Andrew, Jane M. (1993). The Conceptual Basis of Antonymy and Synonymy in Adjectives. *Journal of Memory and Language*: 32.301-319.

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

Pinker, S. 1989. *Learnability and Cognition: The Acquisition of Argument Structure*. Cambridge, Mass.: MIT Press.

Quillian, M. R. 1967. 'Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities.' *Behavioral Science* 12: 410-430.

Quillian, M. R. 1968. 'Semantic Memory.' In Minsky, M. (ed.). *Semantic Information Processing*. Cambridge, Mass.: MIT Press.

Rips, L., and Conrad, F. 1989. 'Folk Psychology of Mental Activities.' *Psychological Review* 96: 187-207.

Roget's International Thesaurus, 4th ed. 1977. New York: Harper and Row.

Rosch, E. 1975. 'Cognitive Representations of Semantic Categories.' *Journal of Experimental Psychology*, 104: 192-233.

Rosch, E., Mervis, C. B., Gray, W., Johnson, D., and Boyes-Braem, P. 1976. 'Basic Objects in Natural Categories.' *Cognitive Psychology* 8: 382-439.

Sapir, E. 1944. 'Grading: A Study in Semantics.' *Philosophy of Science* 11: 83-116.

Schachter, P. 1985. 'Parts-of-Speech Systems.' In Shopen, T., (ed.). *Language Typology and Syntactic Description*, vol. 1. Cambridge: Cambridge University Press.

Schank, R. C. 1972. 'Conceptual Dependency: A Theory of Natural Language Understanding.' *Cognitive psychology* 3: 552-631.

Schank, R. C., and Abelson, R. 1977. *Scripts, Goals, Plans, and Understanding*. Hillsdale, NJ: Erlbaum.

Shibatani, M. 1972. 'Three Reasons for Not Deriving ''Kill'' From ''Cause to Die'' in Japanese' in Kimball, J. (ed.). *Syntax and Semantics*, vol.1. New York: Academic Press.

Smith, E. E. 1978. 'Theories of Semantic Memory.' In Estes, W. K. (ed.). *Handbook of Learning and Cognitive Processes*, vol. 5. Hillsdale, NJ: Erlbaum.

The Synonym Finder. 1978. Emmaus, Pa.: Rodale Press.

Talmy, L. 1985. 'Lexicalization Patterns: Semantic Structure in Lexical Forms' in Shopen, T. (ed.). *Language Typology and Syntactic Description*, vol. 3. Cambridge: Cambridge University Press.

Thorndike, E. L., and Lorge, I. 1944. *The Teacher's Word Book of 30,000 Words*. New York: Bureau of Publications, Columbia University.

Touretzky, D. S. 1986. *The Mathematics of Inheritance Systems.* Los Altos, Calif.: Morgan Kaufmann.

Tversky, B., and Hemenway, K. 1984. 'Objects, Parts, and Categories.' *Journal of Experimental Psychology: General* 113: 169-193.

Vizetelly, F. H. 1915. *The Development of the Dictionary of the English Language*. New York: Funk and Wagnalls.

Wierzbicka, A. 1984. 'Apples are Not a ''Kind of Fruit.'' ' *American Ethnologist* 11: 313-328.

Wilkins, A. J. 1971. 'Conjoint Frequency, Category Size, and Categorization Time.' *Journal of Verbal Learning and Verbal Behavior* 10: 382-385.

Winston, M. E., Chaffin, R., and Hermann, D. J. 1987. 'A Taxonomy of Part-Whole Relations.' *Cognitive Science* 11: 417-444.

Zipf, G. K. 1945. 'The Meaning-Frequency Relationship of Words.' *Journal of General Psychology* 33: 251-256.