# Lexical Ambiguity
# and
# The Role of Knowledge Representation in Lexicon Design

**Branimir Boguraev**
Lexical Systems Group
IBM T.J. Watson Research Center
Yorktown Heights, New York 10598

**James Pustejovsky**
Computer Science Department
Brandeis University
Waltham, MA 02254

## Abstract

The traditional framework for ambiguity resolution employs only 'static' knowledge, expressed generally as selectional restrictions or domain specific constraints, and makes no use of any specific knowledge manipulation mechanisms apart from the simple ability to match valences of structurally related words. In contrast, this paper suggests how a theory of lexical semantics making use of a knowledge representation framework offers a richer, more expressive vocabulary for lexical information. In particular, by performing specialized inference over the ways in which aspects of knowledge structures of words in context can be composed, mutually compatible and contextully relevant lexical components of words and phrases are highlighted. In the view presented here, lexical ambiguity resolution is an integral part of the same procedure that creates the semantic interpretation of a sentence itself.

**Keywords:** lexical ambiguity, lexical semantics, compositionality, lexical organization.

## 1 Importing Knowledge Representation into the Lexicon

Our thesis below is that a theory of lexical semantics making use of a knowledge representation (KR) framework offers a richer, more expressive vocabulary for lexical information. Ultimately the goal of this research is to explain the creative use of language and highlight the role of a constantly evolving lexicon, while obviating the current prevalent views of 'static' lexicon design. A side effect of adopting such a theory as the basis of semantic interpretation is that some classically difficult problems in ambiguity — in particular, lexical — are resolved by viewing them from a different perspective. In an implementation such as that proposed here, lexical ambiguity resolution is an integral part of the same procedure that produces the semantic interpretation of a sentence.

There are several methodological motivations for wanting to import tools developed for computational representation and manipulation of knowledge into the study of word meaning, or *lexical semantics*. In particular, we believe that the goals of computational linguistics are the same as those of linguistics: to provide useful, testable and explanatory theories of the nature of language and its relation to human cognition as a whole. It follows that computational linguistics *is* linguistics as it should now be done, and that the computational tools developed and available in the larger context of Artificial Intelligence should not be ignored by linguists. Furthermore, 'shifting' the application area of KR formalisms from their traditional domain (general world knowledge) to a level below words (lexical knowledge), allows us to abstract the notion of lexical meaning away from world knowledge, as well as from other semantic influences (e.g. discourse and pragmatic factors); such a process of abstraction is a crucial prerequisite of any theory of lexical meaning.

Bringing in KR tools provides us with several benefits, which are instrumental to enriching the semantics of lexical expressions. Firstly, it is now possible to systematically incorporate world knowledge into the lexical entry, while still maintaining an awareness of the boundary between lexical and common sense knowledge. Secondly, it is also possible to reason over that knowledge, facilitating the construction of richer semantic interpretations. Finally, having developed a theory incorporating generalizations about the systematic patterning of words, we have a formal language for expressing these generalizations. The interplay of these capabilities results in a generative language for expressing the meanings of words, while providing a different way of capturing multiple word senses through richer composition. Together with a set of principles for *lexical decomposition*, whose central tenet is that semantic expressions for word meaning (in context) are constructed by a fixed number of generative devices (cf. Pustejovsky [1989b]), this language becomes a tool for expressing lexical knowledge, while not presupposing finite enumeration of word senses.

Current dictionaries reflect, through their organization, the traditional view of word senses; in particular, they assume that the space of possible/allowable uses of a word is exhaustively carved out by an enumerable set of senses for that word. Computational lexicons, to date, generally tend to follow this organization. As a result, the natural language interpretation tasks these lexica support acquire (or inherit) similar view to lexical ambiguity, which then necessitates a particular approach to disambiguation. Furthermore, dictionaries and lexicons currently are of a distinctly static nature: the division into separate word senses not only precludes permeability; it also fails to account for the creative use of words in novel contexts. In contrast, rather than taking a 'snapshot' of language at any moment of time and freezing it into lists of word sense specifications, the model of the lexicon proposed here does not preclude extendability: it is open-ended in nature and accounts for the novel, creative, uses of words in a variety of contexts by positing procedures for generating semantic expressions for words on the basis of particular contexts.

In the remainder of this paper we will illustrate a particular theory of lexical semantics, following Pustejovsky [forthcoming] which promotes the notion of a *generative lexicon*. In particular, we briefly discuss certain types of lexical ambiguity, demonstrate how traditional methods of ambiguity resolution fail to scale up for these (and other) cases, and then outline an approach to semantic interpretation embodying richer methods of compositionality.

As we also show below, the lexical model we propose has the effect of greatly reducing the size of the lexicon. Moreover, it bears directly on issues of organization and content of computational lexicons, as the model now embodies strong assumptions about the kinds of lexical aspects of words essential for natural language processing. The generative theory of lexical semantics, then, imposes a strong focus on current efforts to extract lexical data from large online text resources (dictionaries and corpora): it not only offers a uniform representational framework for expressing the data extracted by the tools and methods of computational lexicography (cf. Boguraev and Briscoe, [1989]), but also offers guidance on the kinds of lexical data — or distinctions in the lexical behavior of words — which should be sought in such resources (cf. Boguraev *et al.*, [1990]).

## 2 The Nature of Lexical Ambiguity

One of the most pervasive phenomena in natural language, and one which every realistic language processing application faces, is that of ambiguity. Consequently, resolution of lexical ambiguity becomes an essential task, without which deeper (or perhaps *any*) language understanding and interpretation is impossible.

### 2.1 Inadequacy of Resolving Ambiguity by Enumeration

As we pointed out earlier, all current computational lexicons fit into a particular processing framework for dealing with this problem. Assuming a partitioning of the space of possible uses of a word into word senses — as postulated and defined by the entry for that word — the problem becomes that of selecting, on the basis of various contextual factors (typically subsumed by, but not necessarily limited to, the notion of selectional restrictions), the word sense closest to the use of the word in the given text. Computationally, this reduces to a search within a finite space of possibilities.

Realistically, however, this approach fails on several accounts — both in terms of what information is made available in a lexicon to drive the disambiguation process and how an automated sense selection procedure might make use of this information. Typically, external contextual factors alone are not sufficient for precise selection of a word sense; additionally, often the lexical entry does not provide enough reliable pointers to critically discriminate between word senses. Secondly, the search process becomes computationally expensive, if not in effect intractable, when it has to account for longer phrases made up of individually ambiguous words. Finally, the assumption that an exhaustive listing can be assigned to the different uses of a word lacks the explanatory power necessary for making generalizations and/or predictions about how words used in a novel way can be reconciled with their currently existing lexical definitions.

To illustrate this last point, below we present some examples of problematic nature for current ambiguity resolution frameworks.

**Creative Use of Words** Consider the ambiguity and context-dependence of adjectives such as *fast* and *slow*, where the meaning of the predicate varies depending on the head being modified. Typically, a lexicon requires an enumeration of different senses for such words, in order to account for the ambiguity illustrated below:

1: *a fast car*: Ambiguous: a car driven quickly / one that is inherently fast.
2: *a fast typist*: the person performs the act of typing quickly.
3: *a fast waltz*: the motion of the dance is quick.
4: *a fast book*: one that can be read in a short time.
5: *a fast reader*: one who reads quickly.

These examples involve at least three distinct word senses for the word *fast*:

`fast(1)`: to move quickly;
`fast(2)`: to perform some act quickly;
`fast(3)`: to do something that takes little time.

(Note that in a real lexicon, word senses would be further annotated with selectional restrictions; these are omitted here for brevity.) Upon closer analysis, each occurrence of *fast* above predicates in a slightly different way. In fact, any finite enumeration of word senses will not account for creative applications of this adjective in the language. For example, *fast* in the phrase *a fast motorway* refers to the ability of vehicles on the motorway to sustain high speed. As a novel use of *fast*, we are clearly looking at a new sense that is not covered by the enumeration above.

**Permeability of Word Senses** Part of our argument for a different organization of the lexicon is based on a claim that the boundaries between the word senses in the analysis of *fast* above are too rigid. Still, even if we assume that enumeration is adequate as a descriptive mechanism, it is not always obvious how to select the correct word sense in any given context: consider the systematic ambiguity of verbs like *bake* (discussed by Atkins *et al.*, [1988]), which require discrimination with respect to *change-of-state* versus *create* readings:

6: *John baked the potato.*
7: *John baked the cake.*

The problem here is that there is too much overlap in the 'core' semantic components of the different readings; hence, it is not possible to guarantee correct word sense selection on the basis of selectional restrictions alone. Furthermore, as language evolves, partial overlaps of core and peripheral components of different word meanings make the traditional notion of word sense, as implemented in current dictionaries, inadequate (see Atkins [1990] for a critique of the flat, linear enumeration-based organization of dictionary entries). The only feasible approach would be to employ considerably more refined distinctions in the semantic content of the complement than is conventionally provided by e.g. the mechanism of selectional restrictions.

**Difference in Syntactic Forms** It is equally arbitrary to create separate word senses for a lexical item just because it can participate in distinct syntactic realizations — and yet this has been the only approach open to computational lexicons which assume the ambiguity resolution framework outlined above. A striking example of this is provided by verbs such as *believe* and *forget*. Observe in (8–11) below that the syntactic realization of the complement determines both the factivity of the proposition in the complement or whether an NP is interpreted as a concealed question (see, for example, Grimshaw [1979]).

8: *Mary forgot that she locked the door.*
9: *Mary forgot to lock the door.*
10: *Mary forgot the answer.*
11: *Mary forgot her wallet.*

Sensitivity to factivity would affect, for instance, the interpretation by a question-answering system, when asked *Did Mary lock the door?* Since sentence (8) is factive and (9) is nonfactive, the answers should be *Yes* and *No* respectively. Such a distinction could be easily accounted for by simply positing separate word senses for each syntactic type, but this misses the obvious relatedness between the two instances of *forget*. If it were possible to make the use of *forget* in (8) and (9) sensitive to the syntactic type of its complement, then we could also explain the parallel cases in (10) and (11). This would allow us to have essentially one definition for *forget* which could, by suitable composition with the different complement types, generate all the allowable readings (cf. Pustejovsky [1989a]).[1]

## 2.2 Towards a Dynamic Model of the Lexicon

The major thrust of our analysis has attempted to show that the ambiguities shown above cannot be adequately handled by exhaustive enumeration of what are regarded as different word senses. It follows that the conventional computational framework for lexical ambiguity resolution, and in particular, the format for lexical entries in current computational lexicons, fails in at least two respects. It is unable to describe all the 'senses' of a word through finite enumeration; and it is also unable to capture interesting generalizations between 'senses' of the same word.

Such failures are partially due to limited (lexical) knowledge made available to natural language processing systems, as well as to an impoverished notion of (lexical) inference. Thus, the traditional framework for ambiguity resolution only employs 'static' knowledge, expressed as e.g. selectional restrictions, and no specific knowledge manipulation mechanisms apart from the simple ability to match valences of connected words. In contrast, we show below how a lexical entry can be assigned a richer knowledge structure and how, by performing specialized inference over the ways in which aspects of knowledge structures of words in context can be composed, mutually compatible and relevant lexical components of words and phrases are highlighted. This process, licensed by constraints operating through the inference mechanisms, in fact, results in generating a semantic interpretation of a phrase, resolving en route the ambiguity of lexical items at their source.

# 3 Ambiguity and Compositionality

The richer structure for the lexical entry proposed here takes to an extreme the established notions of *predicate-argument structure, primitive decomposition* and *conceptual organization*; these are then viewed as defining a space of possible contexts in which a word can be used. Rather than committing to an enumeration of a pre-determined number of different word senses, a lexical entry for a word now encodes a range of deeper aspects of lexical meaning. Looking at a word in isolation, these meaning components simply denote the semantic boundaries appropriate to its use. Viewing a word in the context of other words, mutually compatible aspects in the respective lexical decompositions become more prominent, thus forcing a specific interpretation of each individual word. It is important to realize that this is a generative process, which goes well beyond the simple matching of features. On the contrary, such a framework requires, in addition to a flexible notation for expressing semantic generalizations at the lexical level, a mechanism for composing these individual entries on the phrasal level.

To get a better understanding of how the distinctions in lexical meaning manifest themselves, it is im-

[1] Note that such a treatment is different, in effect, from proposals to conflate more than one syntactic realization of the same word sense through the mechanism of *type coercion* — see Briscoe *et al.* [1990] for a computational approach, based on a suitably enriched lexical representation and utilizing the notions of type coercion (Pustejovsky, [1989a]) and qualia structure (see below).

portant to study and define the role that *all* lexical types play in contributing to the overall meaning of a phrase. This is not just a methodological point: crucial to the processes of semantic interpretation which the lexicon is targeted for is the notion of *compositionality*, necessarily different from the more conventional pairing of e.g. verbs (as functions) and nouns (as arguments). As we indicated earlier, if the semantic load in the lexicon is entirely spread among the verb entries — as many existing computational systems assume — differences like those exemplified in (6–7) and (8–11) can only be accounted for by treating *bake*, *forget*, and so forth as polysemous verbs. If, on the other hand, elaborate lexical meanings of e.g. verbs and adjectives could be made sensitive to components of equally elaborate decompositions of e.g. nouns, the notion of spreading the semantic load evenly across the lexicon becomes the key organizing principle in expressing the knowledge necessary for disambiguation.

In order to be able to express the distinctions, at lexical level, required for analyzing the examples in the last section, it is necessary to go beyond viewing lexical decomposition as based only on a predetermined set of primitives; rather, what is needed is the conjunction of being able to specify (e.g. by means of sets of predicates) different levels, or perspectives, of lexical representation and being able to compose (via a fixed number of generative devices) these predicates. A 'static' definition of a word now only provides its literal meaning; suitable compositions of appropriately 'highlighted' projections of (syntactically) related words, generate meanings in context.

In such a way, many of the shortcomings — in particular those from the perspective of automatic language processing — of the more descriptive approach inherent in exhaustive enumeration of word senses can be overcome. What makes this possible is the combination of two notions, both of them following from general principles of KR theory. First, by positing a language for describing different levels of word meanings, we are no longer confined to the constraints following from operating with a fixed inventory of primitives; moreover, we now also have a way of incorporating in this language the set of rules governing the generative processes. Secondly, through the very nature of these rules, we are assured that the semantic representations ultimately associated with text (fragments) are going to be well-formed.

### 3.1 Levels of Lexical Meaning

Pustejovsky [1989b] proposes several levels of lexical representation. Following an analysis of a broad range of (traditionally ambiguous) constructions, and in particular the aspects of word meanings which account for the ambiguities, he argues for four structures that a theory of computational lexical semantics needs to capture.

**Argument Structure** This defines the conventional mapping from a word to a function, and relates the syntactic realization of a word to the number and type of arguments that are identified at the level of syntax and made use of at the level of semantics (cf. Grimshaw [1990]).

**Event Structure** This identifies the particular event type for a verb or a phrase. There are essentially three components to this structure: the primitive event type — state (S), process (P) or transition (T); the focus of the event; and the rules for event composition.

**Qualia Structure** This defines the essential attributes of an object associated with a lexical item. By positing separate components (see below) in what is, in essence, argument structure for nominals, nouns are elevated from the status of being passive arguments to active functions.

**Lexical Inheritance Structure** This determines the way(s) in which a word is related to other concepts in the lexicon. In addition to answering questions concerning the organization of a (lexical) knowledge base, this level of word meaning makes it possible to link lexical knowledge with general world (common sense) knowledge.

Since the only level of lexical representation not extensively discussed in the literature is that of qualia structure, we briefly outline its components below.

### 3.2 Qualia Structure

The essence of the proposal is that there is a system of relations that characterizes the semantics of nominals, very much like the argument structure of a verb. Pustejovsky [1989b] calls this the *Qualia Structure*, inspired by Aristotle's theory of explanation and ideas from Moravcsik [1975]. In effect, the qualia structure of a noun determines its meaning as much as the list of arguments determines a verb's meaning. The elements that make up a qualia structure include notions such as container, space, surface, figure, artifact, and so on.[2]

Briefly, the Qualia Structure of a word specifies four aspects of its meaning:

- the relation between an object and its constituent parts;
- that which distinguishes it within a larger domain;
- its purpose and function;
- factors involved in its origin or "bringing it about".

These aspects of a word's meaning are called its *Constitutive Role*, *Formal Role*, *Telic Role*, and *Agentive Role*, respectively.[3] The motivation for

[2]These components of an object's denotation have long been considered crucial for our commonsense understanding of how things interact in the world. Cf. Hayes [1979], Hobbs *et al.* [1987a], and Croft [1986] for discussion of these qualitative aspects of meaning.

[3]Some of these roles are reminiscent of descriptors used by various computational researchers, such as Wilks [1975], Hayes [1979], and Hobbs *et al.* [1987a]. Within the theory outlined here, these roles determine a minimal semantic description of a word which has both semantic and grammatical consequences.

positing such characterizations of word meaning is that by enriching the semantic descriptions of nominal types, we will be able to "spread the semantic load" more evenly through the lexicon, while accounting for novel word senses arising in syntactic composition.

### 3.3 Lexical Ambiguity Resolution

Let us examine how this view is able to account for the ambiguities discussed in the previous section. Consider first the example with *fast*. We can capture the general behavior of how such adjectives predicate by making reference to the richer internal structure for nominals suggested above. That is, we can view *fast* as always predicating of the *Telic* role of a nominal. To illustrate this, consider the qualia structure for a noun such as *car*:

```
car(*x*)
    [Const:  {body, engine,...}]
    [Form:   car-shape(*x*)]
    [Telic:  move(P,*x*), drive(P,y,*x*)]
    [Agent:  artifact(*x*)]
```

Notice that the Telic role specifies the purpose and function of the noun. In the phrase, *a fast car*, it is the relation specified there (seen as an event — namely, a process, P) which is modified by the adjective as being fast. Similarly, for the nouns *typist*, *waltz*, *book*, and *reader*, it is their Telic role that is interpreted as being fast (without going into details, we note here that the Telic role of *typist* determines the activity being performed, namely typing; similarly for *waltz*, its Telic role refers to dancing). Hence, the interpretations of *fast* in the examples (1–5) above can all be derived from a single word sense, and there is no need for enumerating the different senses (cf. Pustejovsky [forthcoming]). The lexical semantics for this adjective will indicate that it acts as an event predicate, modifying the Telic role of the noun, as illustrated in the minimal lexical semantic structure for *fast* below:

$$fast(*x*) \Rightarrow (Telic : \lambda P \exists E[fast(E) \wedge P(E, *x*)])$$

Notice that, in addition to obviating the need for separate senses, we can generate the novel use of *fast* mentioned above in the phrase *a fast motorway*, since the Telic role of *motorway* specifies its purpose, and it is this activity which is interpreted as fast:

$$[Telic : travel(P, cars) \wedge on(P, *x*)].$$

The composition of the expression defining *fast* with the lexical aspect it specifies as its 'target' — the Telic role of its argument (*motorway*) — results in an interpretation corresponding to a use of the word when referring to a road: one that allows for fast travel by cars.

### 3.4 Implications for Natural Language Processing

The framework proposed above is very attractive for NLP, for at least two reasons. Firstly, it can be formalized, and thus make the basis for a computational procedure for word interpretation in context. Secondly, it does not require the notion of exhaustive enumeration of all the different ways in which a word can behave, in particular in collocations with other words. Consequently, the framework can naturally cope with the 'creative' use of language; that is, the open-ended nature of word combinations and their associated meanings.

The method of fine-grained characterization of lexical entries, as proposed here, effectively allows us to conflate different word senses (in the traditional meaning of this term) into a single *meta-entry*, thereby offering great potential not only for systematically encoding regularities of word behavior dependent on context, but also for greatly reducing the size of the lexicon. The theoretical claim here is that such a characterization constrains what a possible word meaning can be, through the mechanism of logically well-formed semantic expressions. The expressive power of a KR formalism can then be viewed as simply a tool which gives substance to this claim.

## 4 Knowledge Representation and Global Organization of Lexicon

So far we have looked at the "classical" problem of ambiguity of words, manifested in the problem of how to select suitable word senses for a word in running text, according to some notion of context. As we pointed out just now, the solution outlined in the previous section, in addition to offering an alternative way of approachng the problem, has the important 'side effect' on the size of the lexicon.

In this section we address, at more depth, the question of how the techniques and methods of KR relate directly to the *problem* of lexical ambiguity resolution, the *information* to bear on it, and the *methods* for solving it. The discussion is carried out in the context of an alternative manifestation of the same problem, which we refer to as "hidden" lexical ambiguity. We also use the same context for presenting, informally, some aspects of the lexical inheritance structure as another level of lexical meaning.[4]

One of the implications of positing qualia structures is the necessity to have, superimposed on the lexicon, a realization of more than one lattice structure. Earlier attempts at conceptual hierarchies faced this problem all the time: conceptual models typically make heavy use of multiple inheritance: "book" is_a "literature", "book" is_a "object", "dictionary" is_a "object", "dictionary" is_a "reference", "car" is_a "vehicle" "car" is_a "artifact", and so

[4] Introducing inheritance into the lexicon is not an entirely new idea. For example, Flickinger *et al.* [1985] discuss the value of inheritance as a representational device for capturing generalizations across classes of lexical entries. A further argument for the usefulness of inheritance mechanisms is provided by Briscoe *et al.* [1990], who show how a mechanism of lexical inference can augment a text analysis system which performs syntactic analysis and semantic interpretation by making reference to detailed lexical decomposition of entries in the style of Pustejovsky [1989a].

forth. Still, as descriptive as such relations may appear, models like these suffer from a very limited notion of lexical structure. Thus, even though elaborate mechanisms have been proposed to control and limit the flow of information along the e.g. generalization/specialization links, there has been no theory to either (a) explain how to assign structure to lexical items, or (b) specify lexical relations between lexical items in terms of links between only certain aspects of their respective lexical structures. Pustejovsky's theory of lexical semantics [1989b], with its several distinct levels of semantic description, and in particular the qualia structure, are relevant to just this issue.

On this view, a lexical item inherits information according the qualia structure it carries. In this way, the different senses for words can be rooted into suitable lattices. To illustrate this point, consider the two `is_a` relations below, and the differences in what relations the objects enter into.

| play | (is_a book) | (dictionary is_a book) |
|---|---|---|
| read | ok | no |
| buy | ok | ok |
| consult | no | ok |
| begin | ok(?) | no |

Lexical inheritance theory, on the other hand, posits a separate lattice per role in the qualia structure. Briefly, inheritance through qualia amounts to the following relations for this example: book **is_form** phys-object, book **is_telic** literature, dictionary **is_form** book, dictionary **is_telic** reference, book **is_agent** literature, dictionary **is_agent** compiled-material, play **is_agent** literature, play **is_telic** performance. With the qualia roles differentiating the lattice structures, we can exclude the unwanted inferences listed above.

## 5 Conclusion

We have outlined a framework for lexical semantic research that we believe can be useful for both computational linguists and theoretical linguists alike. We argued against the view that word meanings are fixed and inflexible, where lexical ambiguity must be treated by multiple word entries in the lexicon. Rather, the lexicon can be seen as a generative system, where word senses are related by logical operations defined by the well-formedness rules of the semantics. In this view, much of the lexical ambiguity of highly ambiguous lexical items is eliminated because the semantic load is spread more evenly throughout the lexicon to the other lexical categories; furthermore, the lexical knowledge we propose as necessary for ambiguity resolution is seen as factored out at different levels of lexical representation. We looked at two of these levels, qualia structure and lexical inheritance, as they turn out to be of particular relevance to the structuring of the semantic information carried by e.g. nouns and adjectives, and applying it, via composition, to the construction of semantic interpretation of complex expressions. The methods underlying the analysis of ambiguous phrases and the construction of corresponding semantic expressions make extended use of KR devices and techniques.

## References


[Atkins, 1990] B. Atkins. "Building a Lexicon: Reconciling Anisomorphic Sense Differentiations in Machine-Readable Dictionaries", presented at BBN Symposium: *Natural Language in the 90's—Language and Action in the World.*

[Atkins *et al.*, 1988] B. Atkins, J. Kegl, and B. Levin. "Anatomy of a Verb Entry", *International Journal of Lexicography,* 1(1).

[Boguraev and Briscoe, 1989] B. Boguraev and E. Briscoe (Eds). *Computational Lexicography for Natural Language Processing,* Harlow, Essex, Longman.

[Boguraev *et al.*, 1990] B. Boguraev, R. Byrd, J. Klavans, and M. Neff. "From Structural Analysis of Lexical Resources to Semantics in a Lexical Knowledge Base", RC 15427, IBM T.J.Watson Research Center.

[Briscoe *et al.*, 1990] E. Briscoe, A. Copestake, and B. Boguraev. "Lexical Semantics via Lexicology", ms. University of Cambridge Computer Laboratory and IBM T.J.Watson Research Center.

[Croft, 1986] W. Croft. *Categories and Relations in Syntax: The Clause-Level Organization of Information,* Ph.D. Diss., Stanford University, 1986.

[Flickinger *et al.*, 1985] D. Flickinger, C. Pollard, and T. Wasow. "Structure-Sharing in Lexical Representation", *Proceedings of 23rd Annual Meeting of the ACL,* Chicago, IL.

[Grimshaw, 1979] J. Grimshaw. "Complement Selection and the Lexicon", *Linguistic Inquiry.*

[Hayes, 1979] P.Hayes. "Naive Physics Manifesto", in *Expert Systems in the Micro-Electronic Age,* Donald Mitchie (Ed), Edinburgh University Press, Scotland.

[Hobbs *et al.*, 1987a] J. Hobbs, W. Croft, T. Davies, D. Edwards, and K. Laws. "Commonsense Metaphysics and Lexical Semantics", *Computational Linguistics,* 13(3-4).

[Levin and Rapoport, 1988] B. Levin and T. Rapoport. "Lexical Subordination", *Proceedings of CLS 24,* 275-289.

[Moravcsik, 1975] J.M. Moravcsik. "Aita as Generative Factor in Aristotle's Philosophy", *Dialogue.*

[Pustejovsky, 1989a] J. Pustejovsky. "Type Coercion and Selection", paper at *West Coast Conference on Formal Linguistics,* Vancouver, 1989.

[Pustejovsky, 1989b] J. Pustejovsky. "The Generative Lexicon", ms. Brandeis University.

[Pustejovsky, forthcoming] J. Pustejovsky. *The Generative Lexicon: A Theory of Computational Lexical Semantics,* MIT Press, Cambridge, MA.

[Wilks, 1975] Y. Wilks. "An Intelligent Analyzer and Understander for English", *CACM,* 18, 264-274.