

THE GENERATIVE LEXICON

*James Pustejovsky*Source: *Computational Linguistics* 17(4) (1991): 409–441.

In this paper, I will discuss four major topics relating to current research in lexical semantics: methodology, descriptive coverage, adequacy of the representation, and the computational usefulness of representations. In addressing these issues, I will discuss what I think are some of the central problems facing the lexical semantics community, and suggest ways of best approaching these issues. Then, I will provide a method for the decomposition of lexical categories and outline a theory of lexical semantics embodying a notion of *cocompositionality* and *type coercion*, as well as several levels of semantic description, where the semantic load is spread more evenly throughout the lexicon. I argue that lexical decomposition is possible if it is performed *generatively*. Rather than assuming a fixed set of primitives, I will assume a fixed number of generative devices that can be seen as constructing semantic expressions. I develop a theory of *Qualia Structure*, a representation language for lexical items, which renders much lexical ambiguity in the lexicon unnecessary, while still explaining the systematic polysemy that words carry. Finally, I discuss how individual lexical structures can be integrated into the larger lexical knowledge base through a theory of lexical *inheritance*. This provides us with the necessary principles of global organization for the lexicon, enabling us to fully integrate our natural language lexicon into a conceptual whole.

1. Introduction

I believe we have reached an interesting turning point in research, where linguistic studies can be informed by computational tools for lexicology as well as an appreciation of the computational complexity of large lexical databases. Likewise, computational research can profit from an awareness of the grammatical and syntactic distinctions of lexical items; natural language processing systems must account for these differences in their lexicons and grammars. The wedding of these disciplines is so important, in fact,

that I believe it will soon be difficult to carry out serious computational research in the fields of linguistics and NLP without the help of electronic dictionaries and computational lexicographic resources (cf. Walker *et al.* [forthcoming] and Boguraev and Briscoe [1988]). Positioned at the center of this synthesis is the study of word meaning, lexical semantics, which is currently witnessing a revival.

In order to achieve a synthesis of lexical semantics and NLP, I believe that the lexical semantics community should address the following questions:

1. Has recent work in lexical semantics been methodologically sounder than the previous work in the field?
2. Do theories being developed today have broader coverage than the earlier descriptive work?
3. Do current theories provide any new insights into the representation of knowledge for the global structure of the lexicon?
4. Finally, has recent work provided the computational community with useful resources for parsing, generation, and translation research?

Before addressing these questions, I would like to establish two basic assumptions that will figure prominently in my suggestions for a lexical semantics framework. The first is that, without an appreciation of the syntactic structure of a language, the study of lexical semantics is bound to fail. There is no way in which meaning can be completely divorced from the structure that carries it. This is an important methodological point, since grammatical distinctions are a useful metric in evaluating competing semantic theories.

The second point is that the meanings of words should somehow reflect the deeper, conceptual structures in the system and the domain it operates in. This is tantamount to stating that the semantics of natural language should be the image of nonlinguistic conceptual organizing principles (whatever their structure).

Computational lexical semantics should be guided by the following principles. First, a clear notion of semantic well-formedness will be necessary to characterize a theory of possible word meaning. This may entail abstracting the notion of lexical meaning away from other semantic influences. For instance, this might suggest that discourse and pragmatic factors should be handled differently or separately from the semantic contributions of lexical items in composition.¹ Although this is not a necessary assumption and may in fact be wrong, it may help narrow our focus on what is important for lexical semantic descriptions.

Secondly, lexical semantics must look for representations that are richer than thematic role descriptions (Gruber 1965; Fillmore 1968). As argued in Levin and Rappaport (1986), named roles are useful at best for establishing

fairly general mapping strategies to the syntactic structures in language. The distinctions possible with *theta*-roles are much too coarse-grained to provide a useful semantic interpretation of a sentence. What is needed, I will argue, is a principled method of lexical decomposition. This presupposes, if it is to work at all, (1) a rich, recursive theory of semantic composition, (2) the notion of semantic well-formedness mentioned above, and (3) an appeal to several levels of interpretation in the semantics (Scha 1983).

Thirdly, and related to the point above, **the lexicon is not just verbs**. Recent work has done much to clarify the nature of verb classes and the syntactic constructions that each allows (Levin 1985, 1989). Yet it is not clear whether we are any closer to understanding the underlying nature of verb meaning, why the classes develop as they do, and what consequences these distinctions have for the rest of the lexicon and grammar. The curious thing is that there has been little attention paid to the other lexical categories (but see Miller and Johnson-Laird [1976], Miller and Fellbaum [1991], and Fass [1988]). That is, we have little insight into the semantic nature of adjectival predication, and even less into the semantics of nominals. Not until all major categories have been studied can we hope to arrive at a balanced understanding of the lexicon and the methods of composition.

Stepping back from the lexicon for a moment, let me say briefly what I think the position of lexical research should be within the larger semantic picture. Ever since the earliest attempts at real text understanding, a major problem has been that of controlling the inferences associated with the interpretation process. In other words, how deep or shallow is the understanding of a text? What is the unit of well-formedness when doing natural language understanding: the sentence, utterance, paragraph, or discourse? There is no easy answer to this question because, except for the sentence, these terms are not even formalizable in a way that most researchers would agree on.

It is my opinion that the representation of the context of an utterance should be viewed as involving many different *generative factors* that account for the way that language users create and manipulate the context under constraints, in order to be understood. Within such a theory, where many separate semantic levels (e.g. lexical semantics, compositional semantics, discourse structure, temporal structure) have independent interpretations, the global interpretation of a "discourse" is a highly flexible and malleable structure that has no single interpretation. The individual sources of semantic knowledge compute local inferences with a high degree of certainty (cf. Hobbs *et al.* 1988; Charniak and Goldman 1988). When integrated together, these inferences must be globally coherent, a state that is accomplished by processes of cooperation among separate semantic modules. The basic result of such a view is that semantic interpretation proceeds in a principled fashion, always aware of what the source of a particular inference is, and what the certainty of its value is. Such an approach allows the reasoning

process to be both tractable and computationally efficient. The representation of lexical semantics, therefore, should be seen as just one of many levels in a richer characterization of contextual structure.

2. Methods in lexical semantics

Given what I have said, let us examine the questions presented above in more detail. First, let us turn to the issue of methodology. How can we determine the soundness of our method? Are new techniques available now that have not been adequately explored? Very briefly, one can summarize the most essential techniques assumed by the field, in some way, as follows (see, for example Cruse [1986]):

- On the basis of categorial distinctions, establish the fundamental differences between the grammatical classes; the typical semantic behavior of a word of category *X*. For example, verbs typically behave as predicators, nouns as arguments.
- Find distinctions between elements of the same word class on the basis of collocation and cooccurrence tests. For example, the nouns *dog* and *book* partition into different selectional classes because of contexts involving *animacy*, while the nouns *book* and *literature* partition into different selectional classes because of a mass/count distinction.
- Test for distinctions of a grammatical nature on the basis of diathesis; i.e. alternations that are realized in the syntax. For example, *break* vs. *cut* in (1) and (2) below (Fillmore 1968; Lakoff 1970; Hale and Keyser 1986):

Example 1

- The glass **broke**.
- John **broke** the glass.

Example 2

- *The bread **cut**.
- John **cut** the bread.

Such alternations reveal subtle distinctions in the semantic and syntactic behavior of such verbs. The lexical semantic representations of these verbs are distinguishable on the basis of such tests.

- Test for entailments in the word senses of a lexical item, in different grammatical contexts. One can distinguish, for example, between *context-free* and *context-sensitive* entailments. When the use of a word always entails a certain proposition, we say that the resulting entailment is not dependent on the syntactic context (cf. Katz and Fodor 1963; Karttunen 1971, 1974; Seuren 1985). This is illustrated in Example 3, where a *killing* always entails a *dying*.

Example 3

- John **killed** Bill.
- Bill **died**.

When the same lexical item may carry different entailments in different contexts, we say that the entailments are sensitive to the syntactic contexts; for example, *forget* in Example 4,

Example 4

- John **forgot** that he locked the door.
- John **forgot** to lock the door.

Example 4a has a factive interpretation of *forget* that 4b does not carry; in fact, 4b is counterfactive. Other cases of contextual specification involve aspectual verbs such as *begin* and *finish* as shown in Example 5.

Example 5

- Mary **finished** the cigarette.
- Mary **finished** her beer.

The exact meaning of the verb *finish* varies depending on the object it selects, assuming for these examples the meanings *finish smoking* or *finish drinking*.

- Test for the ambiguity of a word. Distinguish between homonymy and polysemy, (cf. Hirst 1987; Wilks 1975b); that is, from the accidental and logical aspects of ambiguity. For example, the homonymy between the two senses of *bank* in Example 6 is accidental.²

Example 6

- the **bank** of the river
- the richest **bank** in the city

In contrast, the senses in Example 7 exhibit a polysemy (cf. Weinreich 1972; Lakoff 1987).

Example 7

- The **bank** raised its interest rates yesterday (i.e. the *institution*).
- The store is next to the new **bank** (i.e. the *building*).

- Establish what the compositional nature of a lexical item is when applied to other words. For example, *alleged* vs. *female* in Example 8.

Example 8

- the **alleged** suspect
- the **female** suspect

While *female* behaves as a simple intersective modifier in 8b, certain modifiers such as *alleged* in 8a cannot be treated as simple attributes; rather, they create an intensional context for the head they modify. An



even more difficult problem for compositionality arises from phrases containing frequency adjectives (cf. Stump 1981), as shown in 8c and 8d.

Example 8

- c. An **occasional** sailor walks by on the weekend.
- d. Caution: may contain an **occasional** pit (notice on a box of prunes).

The challenge here is that the adjective doesn't modify the nominal head, but the entire proposition containing it (cf. Partee [1985] for discussion). A similar difficulty arises with the interpretation of scalar predicates such as *fast* in Example 9. Both the scale and the relative interpretation being selected for depends on the noun that the predicate is modifying.

Example 9

- a. a **fast** typist: one who types quickly
- b. a **fast** car: one which can move quickly
- c. a **fast** waltz: one with a fast tempo

Such data raise serious questions about the principles of compositionality and how ambiguity should be accounted for by a theory of semantics.

This just briefly characterizes some of the techniques that have been useful for arriving at pre-theoretic notions of word meaning. What has changed over the years are not so much the methods themselves as the descriptive details provided by each test. One thing that has changed, however — and this is significant — is the way computational lexicography has provided stronger techniques and even new tools for lexical semantics research: see Atkins (1987) for sense discrimination tasks; Amsler (1985), Atkins *et al.* (forthcoming) for constructing concept taxonomies; Wilks *et al.* (1988) for establishing semantic relatedness among word senses; and Boguraev and Pustejovsky (forthcoming) for testing new ideas about semantic representations.

3. Descriptive adequacy of existing representations

Turning now to the question of how current theories compare with the coverage of lexical semantic data, there are two generalizations that should be made. First, the taxonomic descriptions that have recently been made of verb classes are far superior to the classifications available twenty years ago (see Levin [1985] for review). Using mainly the descriptive vocabulary of Talmy (1975, 1985) and Jackendoff (1983), fine and subtle distinctions are drawn that were not captured in the earlier, primitives-based approach of Schank (1972, 1975) or the frame semantics of Fillmore (1968).

As an example of the verb classifications developed by various researchers (and compiled by the MIT Lexicon Project; see Levin [1985, 1989]), consider the grammatical alternations in the example sentences below (cf. Dowty 1991).

Example 10

- a. John **met** Mary.
- b. John and Mary **met**.

Example 11

- a. A car **ran** into a truck.
- b. A car and a truck **ran** into each other.

Example 12

- a. A car **ran** into a tree.
- b. *A car and a tree **ran** into each other.

These three pairs show how the semantics of transitive motion verbs (e.g. *run into*) is similar in some respects to reciprocal verbs such as *meet*. The important difference, however, is that the reciprocal interpretation requires that both subject and object be animate or moving; hence 12b is ill-formed. (cf. Levin 1989; Dowty 1991).

Another example of how diathesis reveals the underlying semantic differences between verbs is illustrated in Examples 13 and 14 below. A construction called *the conative* (see Hale and Keyser [1986] and Levin [1985]) involves adding the preposition *at* to the verb, changing the verb meaning to *an action directed toward an object*.

Example 13

- a. Mary **cut** the bread.
- b. Mary **cut** at the bread.

Example 14

- a. Mary **broke** the bread.
- b. *Mary **broke** at the bread.

What these data indicate is that the conative is possible only with verbs of a particular semantic class; namely, verbs that *specify the manner of an action that results in a change of state of an object*.

As useful and informative as the research on verb classification is, there is a major shortcoming with this approach. Unlike the theories of Katz and Fodor (1963), Wilks (1975a), and Quillian (1968), there is no general coherent view on what the entire lexicon will look like when semantic structures for other major categories are studied. This can be essential for establishing

a globally coherent theory of semantic representation. On the other hand, the semantic distinctions captured by these older theories were often too coarse-grained. It is clear, therefore, that the classifications made by Levin and her colleagues are an important starting point for a serious theory of knowledge representation. I claim that lexical semantics must build upon this research toward constructing a theory of word meaning that is integrated into a linguistic theory, as well as interpreted in a real knowledge representation system.

4. Explanatory adequacy of existing representations

In this section I turn to the question of whether current theories have changed the way we look at representation and lexicon design. The question here is whether the representations assumed by current theories are adequate to account for the richness of natural language semantics. It should be pointed out here that a theory of lexical meaning will affect the general design of our semantic theory in several ways. If we view the goal of a semantic theory as being able to recursively assign meanings to expressions, accounting for phenomena such as synonymy, antonymy, polysemy, metonymy, etc., then our view of compositionality depends ultimately on what the basic lexical categories of the language denote. Conventional wisdom on this point paints a picture of words behaving as either active functors or passive arguments (Montague 1974). But we will see that if we change the way in which categories can denote, then the form of compositionality itself changes. Therefore, if done correctly, lexical semantics can be a means to reevaluate the very nature of semantic composition in language.

In what ways could lexical semantics affect the larger methods of composition in semantics? I mentioned above that most of the careful representation work has been done on verb classes. In fact, the semantic weight in both lexical and compositional terms usually falls on the verb. This has obvious consequences for how to treat lexical ambiguity. For example, consider the verb *bake* in the two sentences below.

Example 15

- a. John **baked** the potato.
- b. John **baked** the cake.

Atkins, Kegl, and Levin (1988) demonstrate that verbs such as *bake* are systematically ambiguous, with both a *change-of-state* sense (15a) and a *create* sense (15b).

A similar ambiguity exists with verbs that allow the resultative construction, shown in Examples 16 and 17, and discussed in Dowty (1979), Jackendoff (1983), and Levin and Rapoport (1988).

Example 16

- a. Mary **hammered** the metal.
- b. Mary **hammered** the metal flat.

Example 17

- a. John **wiped** the table.
- b. John **wiped** the table clean.

On many views, the verbs in Examples 16 and 17 are ambiguous, related by either a lexical transformation (Levin and Rapoport 1988), or a meaning postulate (Dowty 1979). In fact, given strict requirements on the way that a verb can project its lexical information, the verb *run* in Example 18 will also have two lexical entries, depending on the syntactic environment it selects (Talmy 1985; Levin and Rappaport 1988).

Example 18

- a. Mary **ran** to the store yesterday.
- b. Mary **ran** yesterday.

These two verbs differ in their semantic representations, where *run* in 18a means *go-to-by-means-of-running*, while in 18b it means simply *move-by-running* (cf. Jackendoff 1983).

The methodology described above for distinguishing word senses is also assumed by those working in more formal frameworks. For example, Dowty (1985) proposes multiple entries for control and raising verbs, and establishes their semantic equivalence with the use of meaning postulates. That is, the verbs in Examples 19 and 20 are lexically distinct but semantically related by rules.³

Example 19

- a. It **seems** that John likes Mary.
- b. John **seems** to like Mary.

Example 20

- a. Mary **prefers** that she come.
- b. Mary **prefers** to come.

Given the conventional notions of function application and composition, there is little choice but to treat all of the above cases as polysemous verbs. Yet, something about the systematicity of such ambiguity suggests that a more general and simpler explanation should be possible. By relaxing the conditions on how the meaning of a complex expression is derived from its parts, I will, in fact, propose a very straight-forward explanation for these cases of *logical polysemy*.

5. A framework for computational semantics

In this section, I will outline what I think are the basic requirements for a theory of computational semantics. I will present a conservative approach to decomposition, where lexical items are minimally decomposed into structured forms (or templates) rather than sets of features. This will provide us with a generative framework for the *composition* of lexical meanings, thereby defining the well-formedness conditions for semantic expressions in a language.

We can distinguish between two distinct approaches to the study of word meaning: *primitive-based* theories and *relation-based* theories. Those advocating primitives assume that word meaning can be exhaustively defined in terms of a fixed set of primitive elements (e.g. Wilks 1975a; Katz 1972; Lakoff 1971; Schank 1975). Inferences are made through the primitives into which a word is decomposed. In contrast to this view, a relation-based theory of word meaning claims that there is no need for decomposition into primitives if words (and their concepts) are associated through a network of explicitly defined links (e.g. Quillian 1968; Collins and Quillian 1969; Fodor 1975; Carnap 1956; Brachman 1979). Sometimes referred to as *meaning postulates*, these links establish any inference between words as an explicit part of a network of word concepts.⁴ What I would like to do is to propose a new way of viewing primitives, looking more at the generative or *compositional* aspects of lexical semantics, rather than the decomposition into a specified *number* of primitives.

Most approaches to lexical semantics making use of primitives can be characterized as using some form of *feature-based* semantics, since the meaning of a word is essentially decomposable into a set of features (e.g. Katz and Fodor 1963; Katz 1972; Wilks 1975; Schank 1975). Even those theories that rely on some internal structure for word meaning (e.g. Dowty 1979; Fillmore 1985) do not provide a complete characterization for all of the well-formed expressions in the language. Jackendoff (1983) comes closest, but falls short of a comprehensive semantics for *all* categories in language. No existing framework, in my view, provides a *method for the decomposition of lexical categories*.

What exactly would a method for lexical decomposition give us? Instead of a taxonomy of the concepts in a language, categorized by sets of features, such a method would tell us the minimal semantic configuration of a lexical item. Furthermore, it should tell us the compositional properties of a word, just as a grammar informs us of the specific syntactic behavior of a certain category. What we are led to, therefore, is a *generative* theory of word meaning, but one very different from the generative semantics of the 1970s.

To explain why I am suggesting that lexical decomposition proceed in a *generative* fashion rather than the traditional *exhaustive* approach, let me take as a classic example, the word *closed* as used in Example 21 (see Lakoff 1970).

Example 21

- a. The door is **closed**.
- b. The door **closed**.
- c. John **closed** the door.

Lakoff (1970), Jackendoff (1972), and others have suggested that the sense in 21c must incorporate something like *cause-to-become-not-open* for its meaning. Similarly, a verb such as *give* specifies a transfer from one person to another, e.g., *cause-to-have*. Most decomposition theories assume a set of primitives and then operate within this set to capture the meanings of all the words in the language. These approaches can be called *exhaustive* since they assume that with a fixed number of primitives, complete definitions of lexical meaning can be given. In the sentences in 21, for example, *close* is defined in terms of the negation of a primitive, *open*. Any method assuming a fixed number of primitives, however, runs into some well-known problems with being able to capture the full expressiveness of natural language.

These problems are not, however, endemic to all decomposition approaches. I would like to suggest that lexical (and conceptual) decomposition is possible if it is performed *generatively*. Rather than assuming a fixed set of *primitives*, let us assume a fixed number of *generative devices* that can be seen as constructing semantic expressions.⁵ Just as a formal language is described more in terms of the productions in the grammar than its accompanying vocabulary, a semantic language is definable by the rules generating the structures for expressions rather than the vocabulary of primitives itself.⁶

How might this be done? Consider the sentences in Example 21 again. A minimal decomposition on the word *closed* is that it introduces an *opposition* of terms: *closed* and *not-closed*. For the verbal forms in 21b and 21c, both terms in this opposition are predicated of different subevents denoted by the sentences. In 21a, this opposition is left implicit, since the sentence refers to a single state. Any minimal analysis of the semantics of a lexical item can be termed a *generative* operation, since it operates on the predicate(s) already literally provided by the word. This type of analysis is essentially Aristotle's *principle of opposition* (cf. Lloyd 1968), and it will form the basis of one level of representation for a lexical item. The essential opposition denoted by a predicate forms part of what I will call the *qualia structure* of that lexical item. Briefly, the qualia structure of a word specifies four *aspects* of its meaning:

- the relation between it and its constituent parts;
- that which distinguishes it within a larger domain (its physical characteristics);
- its purpose and function;
- whatever brings it about.

I will call these aspects of a word's meaning its *Constitutive Role*, *Formal Role*, *Telic Role*, and its *Agentive Role*, respectively.⁷

This minimal semantic distinction is given expressive force when combined with a theory of event types. For example, the predicate in 21a denotes the *state* of the door being closed. No opposition is expressed by this predicate. In 21b and 21c, however, the opposition is explicitly part of the meaning of the predicate. Both these predicates denote what I will call *transitions*. The intransitive use of *close* in 21b makes no mention of the causer, yet the transition from *not-closed* to *closed* is still entailed. In 21c, the event that brings about the *closed* state of the door is made more explicit by specifying the actor involved. These differences constitute what I call the *event structure* of a lexical item. Both the opposition of predicates and the specification of causation are part of a verb's semantics, and are structurally associated with slots in the event template for the word. As we will see in the next section, there are different inferences associated with each event type, as well as different syntactic behaviors (cf. Grimshaw 1990 and Pustejovsky 1991).

Because the lexical semantic representation of a word is not an isolated expression, but is in fact linked to the rest of the lexicon, in Section 7, I suggest how the global integration of the semantics for a lexical item is achieved by structured inheritance through the different *qualia* associated with a word. I call this the *lexical inheritance structure* for the word.

Finally, we must realize that part of the meaning of a word is how it translates the underlying semantic representations into expressions that are utilized by the syntax. This is what many have called the *argument structure* for a lexical item. I will build on Grimshaw's recent proposals (Grimshaw 1990) for how to define the mapping from the lexicon to syntax.

This provides us with an answer to the question of what levels of semantic representation are necessary for a computational lexical semantics. In sum, I will argue that lexical meaning can best be captured by assuming the following levels of representation.

1. **Argument Structure:** The behavior of a word as a function, with its arity specified. This is the predicate argument structure for a word, which indicates how it maps to syntactic expressions.
2. **Event Structure:** Identification of the particular event type (in the sense of Vendler [1967]) for a word or phrase: e.g. as state, process, or transition.
3. **Qualia Structure:** The essential attributes of an object as defined by the lexical item.
4. **Inheritance Structure:** How the word is globally related to other concepts in the lexicon.

These four structures essentially constitute the different levels of semantic expressiveness and representation that are needed for a computational theory of lexical semantics. Each level contributes a different kind of information to the meaning of a word. The important difference between this highly configurational approach to lexical semantics and feature-based approaches is that the recursive calculus defined for word meaning here also provides the foundation for a fully compositional semantics for natural language and its interpretation into a knowledge representation model.

5.1 Argument structure

A logical starting point for our investigations into the meaning of words is what has been called the functional structure or argument structure associated with verbs. What originally began as the simple listing of the parameters or arguments associated with a predicate has developed into a sophisticated view of the way arguments are mapped onto syntactic expressions (for example, the *f-structure* in Lexical Functional Grammar [Bresnan 1982] and the *Projection Principle* in Government-Binding Theory [Chomsky 1981]).

One of the most important contributions has been the view that argument structure is highly structured independent of the syntax. Williams's (1981) distinction between external and internal arguments and Grimshaw's proposal for a hierarchically structured representation (Grimshaw 1990) provide us with the basic syntax for one aspect of a word's meaning.

The argument structure for a word can be seen as a minimal specification of its lexical semantics. By itself, it is certainly inadequate for capturing the semantic characterization of a lexical item, but it is a necessary component.

5.2 Event structure

As mentioned above, the theory of decomposition being outlined here is based on the central idea that word meaning is highly structured, and not simply a set of semantic features. Let us assume this is the case. Then the lexical items in a language will essentially be generated by the recursive principles of our semantic theory. One level of semantic description involves an event-based interpretation of a word or phrase. I will call this level the *event structure* of a word (cf. Pustejovsky 1991; Moens and Steedman 1988). The event structure of a word is one level of the semantic specification for a lexical item, along with its argument structure, qualia structure, and inheritance structure. Because it is recursively defined on the syntax, it is also a property of phrases and sentences.⁸

I will assume a sortal distinction between three classes of events: states (e^S), processes (e^P), and transitions (e^T). Unlike most previous sortal classifications for events, I will adopt a *subeventual* analysis of predicates, as

argued in Pustejovsky (1991) and independently proposed in Croft (1991). In this view, an event sort such as e^7 may be decomposed into two sequentially structured subevents, (e^p, s^s) . Aspects of the proposal will be introduced as needed in the following discussion.

6. A theory of qualia

In Section 5, I demonstrated how most of the lexical semantics research has concentrated on verbal semantics. This bias influences our analyses of how to handle ambiguity and certain noncompositional structures. Therefore, the only way to relate the different senses for the verbs in the examples below was to posit separate entries.

Example 22

- a. John **baked** the potato.
($bake_1 = change(x, State(y))$)
- b. John **baked** the cake.
($bake_2 = create(x, y)$)

Example 23

- a. Mary **hammered** the metal.
($hammer_1 = change(x, State(y))$)
- b. Mary **hammered** the metal flat.
($hammer_2 = cause(x, Become(flat(y)))$)

Example 24

- a. John **wiped** the table.
($wipe_1 = change(x, State(y))$)
- b. John **wiped** the table clean.
($wipe_2 = cause(x, Become(clean(y)))$)

Example 25

- a. Mary **ran** yesterday.
($run_1 = move(x)$)
- b. Mary **ran** to the store yesterday.
($run_2 = go-to(x, y)$)

Although the complement types selected by *bake* in 22, for example, are semantically related, the two word senses are clearly distinct and therefore must be lexically distinguished. According to the sense enumeration view, the same argument holds for the verbs in 23–25 as well.

A similar philosophy has lead linguists to multiply word senses in constructions involving Control and Equi-verbs, where different syntactic contexts necessitate different semantic types.⁹

Example 26

- a. It **seems** that John likes Mary.
- b. John **seems** to like Mary.

Example 27

- a. Mary **prefers** that she come.
- b. Mary **prefers** to come.

Normally, compositionality in such structures simply refers to the application of the functional element, the verb, to its arguments. Yet, such examples indicate that in order to capture the systematicity of such ambiguity, something else is at play, where a richer notion of composition is operative. What then accounts for the polysemy of the verbs in the examples above?

The basic idea I will pursue is the following. Rather than treating the expressions that behave as arguments to a function as simple, passive objects, imagine that they are as active in the semantics as the verb itself. The product of function application would be sensitive to both the function and its active argument. Something like this is suggested in Keenan and Faltz (1985), as the *Meaning-Form Correlation Principle*. I will refer to such behavior as *cocompositionality* (see below). What I have in mind can best be illustrated by returning to the examples in 28.

Example 28

- a. John **baked** the potato.
- b. John **baked** the cake.

Rather than having two separate word senses for a verb such as *bake*, suppose there is simply one, a *change-of-state* reading. Without going into the details of the analysis, let us assume that *bake* can be lexically specified as denoting a *process* verb, and is minimally represented as Example 29.¹⁰

Example 29

Lexical Semantics for *bake*:¹¹
 $\lambda y \lambda x \lambda e^p [bake(e^p \wedge agent(e^p, x) \wedge object(e^p, y))]$

In order to explain the shift in meaning of the verb, we need to specify more clearly what the lexical semantics of a noun is. I have argued above that lexical semantic theory must make a logical distinction between the following qualia roles: the *constitutive*, *formal*, *telic*, and *agentive* roles. Now let us examine these roles in more detail. One can distinguish between *potato* and *cake* in terms of how they come about; the former is a natural kind, while the latter is an artifact. Knowledge of an object includes not just being able to identify or refer, but more specifically, being able to explain how an artifact comes into being, as well as what it is used for; the denotation of an object must identify these roles. Thus, any artifact can be identified with the state of being that object, relative to certain predicates.


As is well known from work on event semantics and *Aktionsarten*, it is a general property of *processes* that they can shift their event type to become a *transition* event (cf. Hinrichs 1985; Moens and Steedman 1987; and Krifka 1987). This particular fact about event structures, together with the semantic distinction made above between the two object types, provides us with an explanation for what I will refer to as the *logical polysemy* of verbs such as *bake*.

As illustrated in Example 30a, when the verb takes as its complement a natural kind such as *potato*, the resulting semantic interpretation is unchanged; i.e., a process reading of a state-change. This is because the noun does not “project” an event structure of its own. That is, relative to the process of *baking*, *potato* does not denote an event-type.¹²

Example 30

a. **bake** as *Process*:

$\exists e^p [bake(e^p) \wedge agent(e^p, j) \wedge object(e^p, a-potato)]$



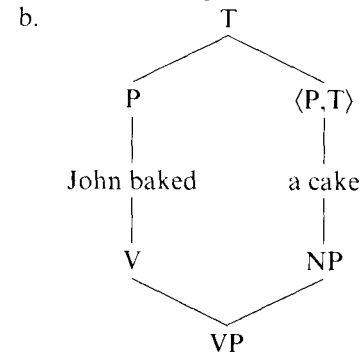
What is it, then, about the semantics of *cake* that shifts this core meaning of *bake* from a state-change predicate to its creation sense? As just suggested, this additional meaning is contributed by specific lexical knowledge we have about artifacts, and *cake* in particular: namely, there is an event associated with that object’s “coming into being,” in this case the process of baking. Thus, just as a verb can select for an argument-type, we can imagine that an argument is itself able to select the predicates that govern it. I will refer to such constructions as *cospecifications*. Informally, relative to the process *bake*, the noun *cake* carries the selectional information that it is a process of “baking” that brings it about.¹³

We can illustrate this schematically in Example 31, where the complement effectively acts like a “stage-level” event predicate (cf. Carlson 1977) relative to the process event-type of the verb (i.e. a function from processes to transitions, $\langle P, T \rangle$).¹⁴ The change in meaning in 31 comes not from the semantics of *bake*, but rather in composition with the complement of the verb, at the level of the entire verb phrase. The “creation” sense arises from the semantic role of *cake* that specifies it is an artifact (see below for discussion).

Example 31

a. **bake** as a derived *Transition*:¹⁵


$\exists e^p, e^s [create(e^p, e^s) \wedge bake(e^p) \wedge agent(e^p, j) \wedge object(e^p, x) \wedge cake(e^s \wedge object(e^s, x))]$



Thus, we can derive both word senses of verbs like *bake* by putting some of the semantic weight on the NP. This view suggests that, in such cases, the verb itself is not polysemous. Rather, the sense of “create” is part of the meaning of *cake* by virtue of it being an artifact. The verb appears polysemous because certain complements add to the basic meaning by virtue of what they denote. We return to this topic below, and provide a formal treatment for how the nominal semantics is expressed in these examples.

Similar principles seem to be operating in the resultative constructions in Examples 23 and 24: namely, a systematic ambiguity is the result of principles of semantic composition rather than lexical ambiguity of the verbs. For example, the resultative interpretations for the verbs *hammer* in 23(b) and *wipe* in 24(b) arise from a similar operation, where both verbs are underlyingly specified with an event type of *process*. The adjectival phrases *flat* and *clean*, although clearly *stative* in nature, can also be interpreted as stage-level event predicates (cf. Dowty 1979). Notice, then, how the resultative construction requires no additional word sense for the verb, nor any special semantic machinery for the resultative interpretation to be available. Schematically, this is shown in Example 32.

Example 32



In fact, this analysis explains why it is that only process verbs participate in the resultative construction, and why the resultant phrase (the adjectival phrase) must be a subset of the states, namely, stage-level event predicates. Because the meaning of the sentence in 32 is determined by both function application of *hammer* to its arguments and function application of *flat* to the event-type of the verb, this is an example of *cocompositionality* (cf. Pustejovsky [forthcoming] for discussion).

Having discussed some of the behavior of logical polysemy in verbs; let us continue our discussion of lexical ambiguity with the issue of *metonymy*. Metonymy, where a subpart or related part of an object “stands for” the object itself, also poses a problem for standard denotational theories of semantics. To see why, imagine how our semantics could account for the “reference shifts” of the complements shown in Example 33.¹⁶

Example 33

- Mary enjoyed the book.
- Thatcher vetoed the channel tunnel. (Cf. Hobbs 1987)
- John began a novel.

The complements of *enjoy* in 33(a) and *begin* in 33(c) are not what these verbs normally select for semantically, namely a property or action. Similarly, the verb *veto* normally selects for an object that is a legislative bill or a suggestion. Syntactically, these may simply be additional subcategorizations, but how are these examples related semantically to the normal interpretations?

I suggest that these are cases of semantic *type coercion* (cf. Pustejovsky 1989a), where the verb has coerced the meaning of a term phrase into a different semantic type. Briefly, type coercion can be defined as follows:¹⁷

Definition

Type Coercion: A semantic operation that converts an argument to the type that is expected by a function, where it would otherwise result in a type error.

In the case of 33(b), it is obvious that what is vetoed is some proposal relating to the object denoted by *the tunnel*. In 33(a), the book is enjoyed only by virtue of some event or process that involves the book, performed by Mary. It might furthermore be reasonable to assume that the semantic structure of *book* specifies what the artifact is used for; i.e. reading. Such a coercion results in a word sense for the NP that I will call *logical metonymy*. Roughly, logical metonymy occurs when a logical argument (i.e. subpart) of a semantic type that is selected by some function denotes the semantic type itself.

Another interesting set of examples involves the possible subjects of causative verbs.¹⁸ Consider the sentences in Examples 34 and 35.

Example 34

- Driving a car in Boston **frightens** me.
- To drive a car in Boston **frightens** me.
- Driving **frightens** me.
- John's driving **frightens** me.
- Cars **frighten** me.
- Listening to this music **upsets** me.
- This music **upsets** me.
- To listen to this music would **upset** me.

Example 35

- John **killed** Mary.
- The gun **killed** Mary.
- John's stupidity **killed** Mary.
- The war **killed** Mary.
- John's pulling the trigger **killed** Mary.

As these examples illustrate, the syntactic argument to a verb is not always the same logical argument in the semantic relation. Although superficially similar to cases of general metonymy (cf. Lakoff and Johnson 1982; Nunberg 1978), there is an interesting systematicity to such shifts in meaning that we will try to characterize below as logical metonymy.

The sentences in 34 illustrate the various syntactic consequences of metonymy and coercion involving experiencer verbs, while those in 35 show the different *metonymic extensions* possible from the causing event in a killing. The generalization here is that when a verb selects an *event* as one of its arguments, type coercion to an event will permit a limited range of logical metonymies. For example, in sentences 34(a,b,c,d,f,h), the entire

event is directly referred to, while in 34(c,g) only a participant from the coerced event reading is directly expressed. Other examples of coercion include “concealed questions” 36 and “concealed exclamations” 37 (cf. Grimshaw 1979; Elliott 1974).

Example 36

- a. John knows *the plane's arrival time*.
(= what time the plane will arrive)
- b. Bill figured out *the answer*.
(= what the answer is)

Example 37

- a. John shocked me with *his bad behavior*.
(= how bad his behavior is)
- b. You'd be surprised at *the big cars he buys*.
(= how big the cars he buys are)

That is, although the italicized phrases syntactically appear as NPs, their semantics is the same as if the verbs had selected an overt question or exclamation.

In explaining the behavior of the systematic ambiguity above, I made reference to properties of the noun phrase that are not typical semantic properties for nouns in linguistics; e.g., artifact, natural kind. In Pustejovsky (1989b) and Pustejovsky and Anick (1988), I suggest that there is a system of relations that characterizes the semantics of nominals, very much like the argument structure of a verb. I called this the *Qualia Structure*, inspired by Aristotle's theory of explanation and ideas from Moravcsik (1975). Essentially, 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.¹⁹

As stated earlier, there are four basic roles that constitute the qualia structure for a lexical item. Here I will elaborate on what these roles are and why they are useful. They are given in Example 38, where each role is defined, along with the possible values that these roles may assume.

Example 38

The Structure of Qualia:

1. **Constitutive Role:** the relation between an object and its constituents, or proper parts.
 - Material
 - Weight
 - Parts and component elements

2. **Formal Role:** that which distinguishes the object within a larger domain.
 - Orientation
 - Magnitude
 - Shape
 - Dimensionality
 - Color
 - Position
3. **Telic Role:** purpose and function of the object.
 - Purpose that an agent has in performing an act
 - Built-in function or aim that specifies certain activities
4. **Agentive Role:** factors involved in the origin or “bringing about” of an object.
 - Creator
 - Artifact
 - Natural Kind
 - Causal Chain

When we combine the qualia structure of a NP with the argument structure of a verb, we begin to see a richer notion of compositionality emerging, one that looks very much like object-oriented approaches to programming (cf. Ingria and Pustejovsky 1990).

To illustrate these structures at play, let us consider a few examples. Assume that the decompositional semantics of a nominal includes a specification of its qualia structure:

Example 39

Object(Const,Form,Telic,Agent)

For example, a minimal semantic description for the noun *novel* will include values for each of these roles, as shown in Example 40, where *x* can be seen as a distinguished variable, representing the object itself.

Example 40

novel(*x*)
 Const: narrative(*x*)
 Form: book(*x*), disk(*x*)
 Telic: read(T,y,*x*)
 Agentive: artifact(*x*), write(T,z,*x*)

This structures our basic knowledge about the object: it is a narrative; typically in the form of a book; for the purpose of reading (whose event type is a *transition*); and is an artifact created by a *transition* event of writing. Observe how this structure differs minimally, but significantly, from the qualia structure for the noun *dictionary* in Example 41.

Example 41

dictionary(*x*)

Const: alphabetized-listing(*x*)

Form: book(*x*), disk(*x*)

Telic: reference(P,y,*x*)

Agentive: artifact(*x*), compile(T,z,*x*)

Notice the differences in the values for the *constitutive* and *telic* roles. The purpose of a dictionary is an activity of referencing, which has an event structure of a *process*.

I will now demonstrate that such structured information is not only useful for nouns, but necessary to account for their semantic behavior. I suggested earlier, that for cases such as 33, repeated below, there was no need to posit a separate lexical entry for each verb, where the syntactic and semantic types had to be represented explicitly.

Example 42

- a. Mary enjoyed the book.
- b. Thatcher vetoed the channel tunnel.
- c. John began a novel.

Rather, the verb was analyzed as *coercing* its complement to the semantic type it expected. To illustrate this, consider 42(c). The type for *begin* within a standard typed intensional logic is $\langle VP, \langle NP, S \rangle \rangle$, and its lexical semantics is similar to that of other subject control verbs (cf. Klein and Sag [1985] for discussion).

Example 43 $\lambda P \lambda \mathcal{P} \mathcal{P} \lambda x [begin'(P(x^*))(x^*)]$

Assuming an event structure such as that of Krifka (1987) or Pustejovsky (1991), we can convert this lexical entry into a representation consistent with a logic making use of event-types (or sorts) by means of the following meaning postulat.²⁰

Example 44 $\forall P \forall x_1 \dots x_n \square [P_\sigma(x_1) \dots (x_n) \leftrightarrow \exists e^\sigma [P(x_1) \dots (x_n)(e^\sigma)]]$

This allows us to type the verb *begin* as taking a transition event as its first argument, represented in Example 45.

Example 45 $\lambda P_T \lambda \mathcal{P} \mathcal{P} \lambda x [begin'(P_T(x^*))(x^*)]$

Because the verb requires that its first argument be of type *transition* the complement in 33(c) will not match without some sort of shift. It is just this kind of context where the complement (in this case *a novel*) is *coerced* to another type. The coercion dictates to the complement that it must conform to its type specification and the qualia roles may in fact have values matching the correct type. For purposes of illustration, the qualia structure for *novel* from 41 can be represented as the logical expression in Example 46.

Example 46**novel** translates into: $\lambda x [novel(x) \wedge Const(x) = narrative'(x) \wedge$ $Form(x) = book'(x) \wedge$ $Telic(x) = \lambda y. e^T [read'(x)(y)(e^T)] \wedge$ $Agent(x) = \lambda y. e^T [write'(x)(y)(e^T)]]$

The coercion operation on the complement in the above examples can be seen as a request to find any transition event associated with the noun. As we saw above, the qualia structure contains just this kind of information.

We can imagine the qualia roles as partial functions from a noun denotation into its subconstituent denotations. For our present purposes, we abbreviate these functions as Q_F, Q_C, Q_T, Q_A . When applied, they return the value of a particular qualia role. For example, the purpose of a novel is for reading it, shown in 47(a), while the mode of creating a novel is by writing it, represented in 47(b).

Example 47a. $Q_T(\mathbf{novel}) = \lambda y. e^T [read(x)(y)(e^T)]$ b. $Q_A(\mathbf{novel}) = \lambda y. e^T [write(x)(y)(e^T)]$


As the expressions in 47 suggest, there are, in fact, two obvious interpretations for this sentence in 42(c).

Example 48a. John began to **read** a novel.b. John began to **write** a novel.

One of these is selected by the coercing verb, resulting in a complement that has a event-predicate interpretation, without any syntactic transformations (cf. Pustejovsky [1989a] for details).²¹ The derivation in 49(a) and the structure in 49(b) show the effects of this coercion on the verb's complement, using the *telic* value of *novel*.²²

Example 49

- a. John began a novel.
- b. $\text{begin}'(Q_T(\text{a novel}))(\text{John}) \Rightarrow$
- c. $\text{begin}'(\lambda x, e^T[\text{read}(\text{a novel})(x)(e^T)])(\text{John}) \Rightarrow$
- d. $\text{John}\{\lambda x[\text{begin}'(\lambda x, e^T[\text{read}(\text{a novel})(x)(e^T)](x^*))(x^*)]\} \Rightarrow$
- e. $\text{John}\{\lambda x[\text{begin}'(\lambda e^T[\text{read}(\text{a novel})(x^*)(e^T)])(x^*)]\} \Rightarrow$
- f. $\text{begin}'(\lambda e^T[\text{read}(\text{a novel})(\text{John})(e^T)])(\text{John})$



The fact that this is not a unique interpretation of the elliptical event predicate is in some ways irrelevant to the notion of type coercion. That there is *some* event involving the complement is required by the lexical semantics of the governing verb and the rules of type well-formedness, and although there are many ways to act on a novel, I argue that certain relations are “privileged” in the lexical semantics of the noun. It is not the role of a lexical semantic theory to say what readings are preferred, but rather which are available.²³

Assuming the semantic selection given above for *begin* is correct, we would predict that, because of the process event-type associated with the telic role for *dictionary*, there is only one default interpretation for the sentence in 50; namely, the *agentive* event of “compiling.”

Example 50

- a. Mary began a dictionary. (Agentive)
- b. ?? Mary began a dictionary. (Telic)

Not surprisingly, when the noun in complement position has no default interpretation within an event predicate — as given by its qualia structure — the resulting sentence is extremely odd.

Example 51

- a. *Mary began a rock.
- b. ??John finished the flower.

The semantic distinctions that are possible once we give semantic weight to lexical items other than verbs are quite wide-ranging. The next example I will consider concerns scalar modifiers, such as *fast*, that modify different predicates depending on the head they modify. If we think of certain modifiers as modifying only a subset of the qualia for a noun, then we can view *fast* as modifying only the telic role of an object. This allows us to go beyond treating adjectives such as *fast* as intersective modifiers — for example, as $\lambda x[\text{car}'(x) \wedge \text{fast}'(x)]$. Let us assume that an adjective such as *fast* is a member of the general type $\langle N, N \rangle$, but can be subtyped as applying to the Telic role of the noun being modified. That is, it has as its type, $\langle [N \text{ Telic}], N \rangle$. This gives rise directly to the different interpretations in Example 52.

Example 52

- a. **a fast car:** driving
 $Q_T(\text{car}) = \lambda x \lambda y \lambda e^P [\text{drive}(x)(y)(e^P)]$
- b. **a fast typist:** typing
 $Q_T(\text{typist}) = \lambda x \lambda e^P [\text{type}(x)(e^P)]$
- c. **a fast motorway:** traveling
 $Q_T(\text{motorway}) = \lambda x \lambda e^P [\text{travel}(\text{cars})(e^P) \wedge \text{on}(x)(\text{cars})(e^P)]$

These interpretations are all derived from a single word sense for *fast*. Because the lexical semantics for this adjective indicates that it modifies the telic role of the noun, it effectively acts as an event predicate rather than an attribute over the entire noun denotation, as illustrated in Example 53 for *fast motorway* (cf. Pustejovsky and Boguraev [1991] for discussion).

Example 53

$$\lambda x[\text{motorway}(x) \dots [\text{Telic}(x) = \lambda e^P[\text{travel}(\text{cars})(e^P) \wedge \text{on}(x)(\text{cars})(e^P) \wedge \text{fast}(e^P)]]]$$

As our final example of how the qualia structure contributes to the semantic interpretation of a sentence, observe how the nominals *window* and *door* in Examples 54 and 55 carry two interpretations (cf. Lakoff [1987] and Pustejovsky and Anick [1988]):

Example 54

- a. John crawled through *the window*.
- b. *The window* is closed.

Example 55

- a. Mary painted *the door*.
- b. Mary walked through *the door*.

Each noun appears to have two word senses: a physical object denotation and an aperture denotation. Pustejovsky and Anick (1988) characterize the meaning of such “Double Figure-Ground” nominals as inherently relational, where both parameters are logically part of the meaning of the noun. In terms of the qualia structure for this class of nouns, the formal role takes as its value the *Figure* of a physical object, while the constitutive role assumes the *Invert-Figure* value of an aperture.²⁴

Example 56

Lexical Semantics for **door**:

door(*x*,*y*)
 Const: aperture(*y*)
 Form: phys-obj(*x*)
 Telic: pass-through(T,z,*y*)
 Agentive: artifact(*x*)

The foregrounding or backgrounding of a nominal’s qualia is very similar to argument structure-changing operations for verbs. That is, in 55(a), *paint* applies to the formal role of *the door*, while in 55(b), *through* will apply to the constitutive interpretation of the same NP. The ambiguity with such nouns is a logical one, one that is intimately linked to the semantic representation of the object itself. The qualia structure, then, is a way of capturing this logical polysemy.

In conclusion, it should be pointed out that the entire lexicon is organized around such logical ambiguities, which Pustejovsky and Anick (1988) call *Lexical Conceptual Paradigms*. Pustejovsky (forthcoming) distinguishes the following systems and the paradigms that lexical items fall into:

Example 57

- a. Count/Mass Alternations
- b. Container/Containee Alternations
- c. Figure/Ground Reversals
- d. Product/Producer Diathesis
- e. Plant/Fruit Alternations
- f. Process/Result Diathesis
- g. Object/Place Reversals
- h. State/Thing Alternations
- i. Place/People

Such paradigms provide a means for accounting for the systematic ambiguity that may exist for a lexical item. For example, a noun behaving according to paradigm 57(a) exhibits a logical polysemy involving packaging or grinding operators; e.g., *haddock* or *lamb* (cf. Copestake and Briscoe [1991] for details).

7. Lexical inheritance theory

In previous sections, I discussed lexical ambiguity and showed how a richer view of lexical semantics allows us to view a word’s meaning as being flexible, where word senses could arise generatively by composition with other words. The final aspect of this flexibility deals with the logical associations a word has in a given context: that is, how this semantic information is organized as a global knowledge base. This involves capturing both the inheritance relations between concepts and, just as importantly, how the concepts are integrated into a coherent expression in a given sentence.

I will assume that there are two inheritance mechanisms at work for representing the conceptual relations in the lexicon: *fixed* inheritance and *projective* inheritance. The first includes the methods of inheritance traditionally assumed in AI and lexical research (e.g. Roberts and Goldstein 1977; Brachman and Schmolze 1985; Bobrow and Winograd 1977); that is, a *fixed* network of relations, which is traversed to discover existing related and associated concepts (e.g. hyponyms and hypernyms). In order to arrive at a comprehensive theory of the lexicon, we need to address the issue of global organization, and this involves looking at the various modes of inheritance that exist in language and conceptualization. Some of the best work addressing the issue of how the lexical semantics of a word ties into its deeper conceptual structure includes that of Hobbs *et al.* (1987) and Wilks (1975), while interesting work on shared information structures in NLP domains is that of Flickinger *et al.* (1985) and Evans and Gazdar (1989, 1990).

In addition to this static representation, I will introduce another mechanism for structuring lexical knowledge, the *projective* inheritance, which operates *generatively* from the qualia structure of a lexical item to create a relational structure for ad hoc categories. Both are necessary for projecting the semantic representations of individual lexical items onto a sentence level interpretation. The discussion here, however, will be limited to a description of projective inheritance and the notion of “degrees of prototypicality” of predication. I will argue that such degrees of salience or coherence relations can be explained in structural terms by examining a network of related lexical items.²⁵

I will illustrate the distinction between these mechanisms by considering the two sentences in Example 58, and their relative prototypicality.

Example 58

- a. The prisoner *escaped* last night.
- b. The prisoner *ate* dinner last night.

Both of these sentences are obviously well-formed syntactically, but there is a definite sense that the predication in 58(a) is “tighter” or more prototypical

than that in 58(b). What would account for such a difference? Intuitively, we associate prisoner with an *escaping* event more strongly than an *eating* event. Yet this is not information that comes from a fixed inheritance structure, but is rather usually assumed to be commonsense knowledge. In what follows, however, I will show that such distinctions can be captured within a theory of lexical semantics by means of generating ad hoc categories.

First, we give a definition for the *fixed inheritance* structure of a lexical item (cf. Touretzky 1986). Let Q and P be concepts in our model of lexical organization. Then:

Definition

A sequence $\langle Q_1, P_1, \dots, P_n \rangle$ is an *inheritance path*, which can be read as the conjunction of ordered pairs $\{\langle x_i, y_i \rangle \mid 1 \leq i \leq n\}$.

Furthermore, following Touretzky, from this we can define the set of concepts that lie on an inheritance path, the *conclusion space*.

Definition

The *conclusion space* of a set of sequences Φ is the set of all pairs $\langle Q, P \rangle$ such that a sequence $\langle Q, \dots, P \rangle$ appears in Φ .

From these two definitions we can define the traditional is-a relation, relating the above pairs by a *generalization* operator, \leq_G ,²⁶ as well as other relations that I will not discuss.²⁷

Let us suppose that, in addition to these fixed relational structures, our semantics allows us to dynamically create arbitrary concepts through the application of certain transformations to lexical meanings. For example, for any predicate, Q — e.g. the value of a qualia role — we can generate its opposition, $\neg Q$ (cf. Pustejovsky 1991). By relating these two predicates temporally we can generate the arbitrary transition events for this opposition (cf. Wright 1963):

Example 59

- a. $\neg Q(x) \leq Q(x)$
- b. $Q(x) \leq \neg Q(x)$
- c. $Q(x) \leq Q(x)$
- d. $\neg Q(x) \leq \neg Q(x)$

Similarly, by operating over other qualia role values we can generate semantically related concepts. I will call any operator that performs such an operation a *projective transformation*, and define them below:

Definition

A *projective transformation*, π , on a predicate Q_1 generates a predicate, Q_2 , such that $\pi(Q_1) = Q_2$, where $Q_2 \notin \Phi$. The set of transformations includes: \neg , negation, \leq , temporal precedence, \geq , temporal succession, $=$, temporal equivalence, and *act*, an operator adding agency to an argument.

Intuitively, the space of concepts traversed by the application of such operators will be related expressions in the neighborhood of the original lexical item. This space can be characterized by the following two definitions:

Definition

A series of applications of transformations, π_1, \dots, π_n , generates a sequence of predicates, $\langle Q_1, \dots, Q_n \rangle$, called the *projective expansion* of Q_1 , $P(Q_1)$.

Definition

The *projective conclusion space*, $P(\Phi_R)$, is the set of projective expansions generated from all elements of the conclusion space, Φ , on role R of predicate Q : as: $P(\Phi_R) = \{\langle P(Q_1), P(Q_n) \rangle \mid \langle Q_1, \dots, Q_n \rangle \in \Phi_R\}$.

From this resulting representation, we can generate a relational structure that can be considered the set of ad hoc categories and relations associated with a lexical item (cf. Barselou 1983).

Using these definitions, let us return to the sentences in Example 58. I will assume that the noun *prisoner* has a qualia structure such as that shown in 60.

Example 60

Qualia Structure of **prisoner(x)**:

prisoner(*x*)
 Form: human(*x*)
 Telic: [confine(y,*x*) & location(*x*,prison)]

Furthermore, I assume the following lexical structure for *escape*.

Example 61

Lexical Semantics for *escape*:

$\lambda x \lambda e^T \exists e^P, e^S [escape(e^T) \wedge act(e^P) \wedge confined(e^P) \wedge agent(e^P, x) \wedge \neg confined(e^S) \wedge object(e^S, x)]$

Using the representation in 60 above, I now trace part of the derivation of the projective conclusion space for *prisoner*. Inheritance structures are

defined for each qualia role of an element. In the case above, values are specified for only two roles.

For each role, R , we apply a projective transformation π onto the predicate Q that is the value of that role. For example, from the telic role of *prisoner* we can generalize (e.g. drop the conjunct) to the concept of *being confined*. From this concept, we can apply the negation operator, generating the predicate opposition of *not-confined* and *confined*. To this, we apply the two temporal operators, \leq and \geq , generating two states: *free before capture* and *free after capture*. Finally, to these concepts, if we apply the operator *act*, varying who is responsible for the resulting transition event, we generate the concepts: *turn in*, *capture*, *escape*, and *release*.

Example 62

Projecting on *Telic Role* of *prisoner*:

- a. $\leq_G: [confine(y,x) \wedge loc(x,prison)] \Rightarrow confine(y,x)$
- b. $\leq: \exists E_1[\neg confine(E_1,y,x)]$
- c. $\exists E_2[confine(E_2,y,x)]$
- d. $\leq: E_1 \leq E_2 = T_1$
- e. $\leq: E_1 \leq E_1 = T_2$
- f. *act*: $act(x,T_1) = \text{“turn in”}$
- g. *act*: $act(y,T_1) = \text{“capture”}$
- h. *act*: $act(x,T_2) = \text{“escape”}$
- i. *act*: $act(y,T_2) = \text{“release”}$

These relations constitute the *projective conclusion space* for the telic role of *prisoner* relative to the application of the transformations mentioned above. Similar operations on the formal role will generate concepts such as *die* and *kill*. Generating such structures for all items in a sentence during analysis, we can take those graphs that result in no contradictions to be the legitimate semantic interpretations of the entire sentence.

Let us now return to the sentences in Example 58. It is now clear why these two sentences differ in their prototypicality (or the relevance conditions on their predication). The predicate *eat* is not within the space of related concepts generated from the semantics of the NP *the prisoner*, *escape*, however, did fall within the projective conclusion space for the Telic role of *prisoner*, as shown in Example 63.

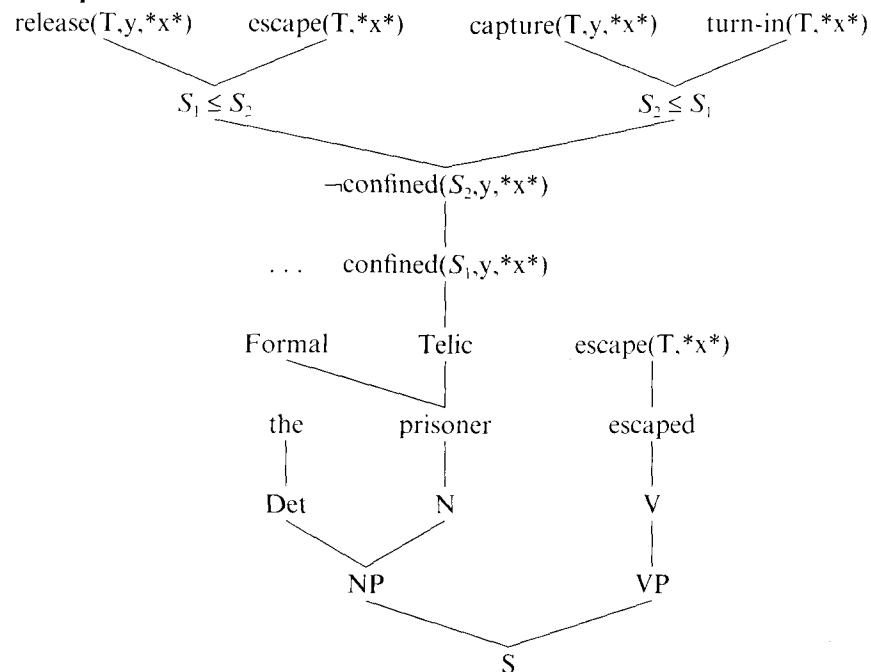
Example 63

Conclusion Space for (58):

- $escape \in P(\Phi_T(prisoner))$
- $eat \notin P(\Phi_T(prisoner))$

This is illustrated in Example 64 below.

Example 64



We can therefore use such a procedure as one metric for evaluating the “proximity” of a predication (Quillian 1968; Hobbs 1982). In the examples above, the difference in semanticity can now be seen as a structural distinction between the semantic representations for the elements in the sentence.

In this section, I have shown how the lexical inheritance structure of an item relates, in a generative fashion, the decompositional structure of a word to a much larger set of concepts that are related in obvious ways. What we have not addressed, however, is how the fixed inheritance information of a lexical item is formally derivable during composition. This issue is explicitly addressed in Briscoe *et al.* (1990) as well as Pustejovsky and Briscoe (1991).

8. Conclusion

In this paper I have outlined a framework for lexical semantic research that I believe can be useful for both computational linguists and theoretical linguists alike. I 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 verbs

and prepositions is eliminated because the semantic load is spread more evenly throughout the lexicon to the other lexical categories. I described a language for structuring the semantic information carried by nouns and adjectives, termed *Qualia structure*, as well as the rules of composition that allow this information to be incorporated into the semantic interpretation of larger expressions, including explicit methods for type coercion. Finally, I discussed how these richer lexical representations can be used to generate projective inheritance structures that connect the conceptual information associated with lexical items to the global conceptual lexicon. This suggests a way of accounting for relations such as coherence and the prototypicality of a predication. Although much of what I have presented here is incomplete and perhaps somewhat programmatic, I firmly believe this approach can help clarify the nature of word meaning and compositionality in natural language, and at the same time bring us closer to understanding the creative use of word senses.

Acknowledgments

I would like to thank the following for comments on earlier drafts of this paper: Peter Anick, Sabine Bergler, Bran Boguraev, Ted Briscoe, Noam Chomsky, Bob Ingria, George Miller, Sergei Nirenburg, and Rich Thomason.

Notes

- 1 This is still a contentious point and is an issue that is not at all resolved in the community. Hobbs (1987) and Wilensky (1990), for example, argue that there should be no distinction between commonsense knowledge and lexical knowledge. Nevertheless, I will suggest below that there are good reasons, both methodological and empirical, for establishing just such a division. Pustejovsky and Bergler (1991) contains a good survey on how this issue is addressed by the community.
- 2 Cf. Weinreich (1972) distinguishes between contrastive and complementary polysemy, essentially covering this same distinction. See Section 4 for discussion.
- 3 Both Klein and Sag (1985) and Chomsky (1981) assume, however, that there are reasons for relating these two forms structurally. See below and Pustejovsky (1989a) for details.
- 4 For further discussion on the advantages and disadvantages to both approaches, see Jackendoff (1983).
- 5 See Goodman (1951) and Chomsky (1955) for explanations of the method assumed here.
- 6 This approach is also better suited to the way people write systems in computational linguistics. Different people have distinct primitives for their own domains, and rather than committing a designer to a particular vocabulary of primitives, a lexical semantics should provide a method for the decomposition and composition of lexical items.
- 7 Some of these roles are reminiscent of descriptors used by various computational researchers, such as Wilks (1975b), Hayes (1979), and Hobbs *et al.* (1987). Within

the theory outlined here, these roles determine a minimal semantic description of a word that has both semantic and grammatical consequences.

- 8 This proposal is an extension of ideas explored by Bach (1986), Higginbotham (1985), and Allen (1984). For a full discussion, see Pustejovsky (1988, 1991). See Tenny (1987) for a proposal on how aspectual distinctions are mapped to the syntax.
- 9 For example, Dowty (1985) proposes multiple entries for verbs taking different subcategorizations. Gazdar *et al.* (1985), adopting the analysis in Klein and Sag (1985), propose a set of lexical type-shifting operations to capture sense relatedness. We return to this topic below.
- 10 I will be assuming a Davidsonian-style representation for the discussion below. Predicates in the language are typed for a particular event-sort, and thematic roles are treated as partial functions over the event (cf. Dowty 1989 and Chierchia 1989).
- 11 More precisely, the process e^p should reflect that it is the substance contained in the object x that is affected. See footnote 20 for explanation.
- 12 However, relative to the process of *growing*, the noun *potato* does denote an event:

1. Mary grew the potato.

- 13 Other examples of cospecifications are: a. *read a book*, b. *smoke a cigarette*, c. *mail a letter*, d. *deliver a lecture*, and e. *take a bath*.

There are several interesting things about such collocations. First, because the complement "selects" the verb that governs it (by virtue of knowledge of what is done to the object), the semantics of the phrase is changed. The semantic "connectedness," as it were, is tighter when cospecification obtains. In such cases, the verb is able to successfully drop the dative PP argument, as shown below in (1). When the complement does not select the verb governing it, dative-drop is ungrammatical as seen in (2) (although there are predicates selected by these nouns; e.g. *keep a secret*, *read a book*, and *play a record*).

- 1a. Romeo gave the lecture.
- 1b. Hamlet mailed a letter.
- 1c. Cordelia told a story.
- 1d. Gertrude showed a movie.
- 1e. Mary asked a question.
- 2a. *Bill told the secret.
- 2b. *Mary gave a book.
- 2c. *Cordelia showed the record.

For discussion see Pustejovsky (in press).

- 14 cf. Pustejovsky (forthcoming) for details.
- 15 As mentioned in footnote 11, this representation is incomplete. See footnote 20 for semantics of *bake*.
- 16 See Nunberg (1978) and Fauconnier (1985) for very clear discussions of the semantics of metonymy and the nature of reference shifts. See Wilks (1975) and Fass (1988) for computational models of metonymic resolution.
- 17 I am following Cardelli and Wegener (1985) and their characterization of polymorphic behavior.
- 18 See Verma and Mohanan (1991) for an extensive survey of experiencer subject constructions in different languages.

- 19 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.* (1987), and Croft (1991) for discussion of these qualitative aspects of meaning.
- 20 It should be pointed out that the lexical structure for the verb *bake* given above in 30 and 31 can more properly be characterized as a process acting on various qualia of the arguments.

$$1. \lambda_Y \lambda_X \lambda_e^p e^s \{bake(e^p) \wedge agent(e^p, x) \wedge object(e^p, Const(Y)) \wedge cake(e^s) \wedge object(e^s, Formal(x))\}$$

- 21 There are, of course, an indefinite number of interpretations, depending on pragmatic factors and various contextual influences. But I maintain that there are only a finite number of default interpretations available in such constructions. These form part of the lexical semantics of the noun. Additional evidence for this distinction is given in Pustejovsky and Anick (1988) and Briscoe *et al.* (1990).
- 22 Partee and Rooth (1983) suggest that all expressions in the language can be assigned a base type, while also being associated with a type ladder. Pustejovsky (1989a) extends this proposal, and argues that each expression *a* may have available to it, a set of shifting operators, which we call Σ_a , which operate over an expression, changing its type and denotation. By making reference to these operators directly in the rule of function application, we can treat the functor polymorphically, as illustrated below.

1. Function Application with Coercion (FA_c):

If *a* is of type $\langle b, a \rangle$, and β is of type *c*, then

- (a) if type $c = b$, then $\alpha(\beta)$ is of type *a*.
 (b) if there is a $\sigma \in \Sigma_\beta$ such that $\sigma(\beta)$ results in an expression of type *b*, then $\alpha(\sigma(\beta))$ is of type *a*.
 (c) otherwise a type error is produced.

- 23 There are interesting differences in complement types between *finish* and *complete*. The former takes both NP and a gerundive VP, while the latter takes only an NP (cf. for example, Freed [1979] for discussion).

- 1a. John finished the book.
 b. John finished writing the book.
- 2a. John completed the book.
 b. *John completed writing the book.

The difference would indicate that, contrary to some views (e.g. Wierzbicka [1988] and Dixon [1991]), lexical items need to carry both syntactic and semantic selectional information to determine the range of complements they may take. Notice here also that *complete* tends to select the agentive role value for its complement and not the telic role. The scope of semantic selection is explored at length in Pustejovsky (forthcoming).

- 24 There are many such classes of nominals, both two-dimensional such as those mentioned in the text, and three-dimensional, such as "room," "fireplace," and "pipe." They are interesting semantically, because they are logically ambiguous, referring to either the object or the aperture, but not both. Boguraev and Pustejovsky (forthcoming) show how these logical polysemies are in fact encoded in dictionary definitions for these words.

- 25 Anick and Pustejovsky (1990) explore how metrics such as *association ratios* can be used to statistically measure the notions of prototypicality mentioned here.
- 26 See, for example, Michalski (1983) and Smolka (1988) for a treatment making use of subsorts.
- 27 Such relations include not only hypernymy and hopyonymy, but also troponymy, which relates verbs by manner relations (cf. Miller 1985; Beckwith *et al.* 1989; Miller and Fellbaum 1991).

References

- Anick, P., and Pustejovsky, J. (1990). "An application of lexical semantics to knowledge acquisition from corpora." In *Proceedings, 13th International Conference on Computational Linguistics*, Helsinki, Finland.
- Atkins, Beryl T. (1987). "Semantic ID tags: corpus evidence for dictionary senses." In *Proceedings, 3rd Annual Conference at University of Waterloo*, Center for the New OED.
- Atkins, Beryl; Kegl, Judy; and Levin, Beth (1988). "Anatomy of a verb entry." *International Journal of Lexicography*, 1:84-126.
- Atkins, Beryl; Klavans, Judith; and Boguraev, Bran. "Semantic verb clusters from MRDs." Forthcoming.
- Bach, Emmon (1986). "The algebra of events." *Linguistics and Philosophy*, 9: 5-16.
- Barselou, Lawrence W. (1983) "Ad hoc categories." *Memory and Cognition*, 11:211-227.
- Beckwith, Richard; Fellbaum, C.; Gross, D.; and Miller, G. (1989). "WordNet: A lexical database organized on psycholinguistic principles." *Proceedings, First International Workshop on Lexical Acquisition, IJCAI*, Detroit.
- Bobrow, D. G., and Winograd, T. (1977). "An overview of KRL, a knowledge representation language." *Cognitive Science*, 1:3-46.
- Boguraev, Bran, and Briscoe, Ted (eds.) (1988). *Computational Lexicography for Natural Language Processing*. Harlow, Essex: Longman.
- Boguraev, Bran, and Pustejovsky, James. *Lexical Knowledge: Representation and Acquisition*. Forthcoming. Cambridge, MA: The MIT Press.
- Brachman, Ronald J. (1979). "On the epistemological status of semantic networks." In *Associative Networks: Representation and Use of Knowledge by Computers*, edited by N. Findler. New York: Academic Press.
- Brachman, R. J., and Schmolze, J. (1985). "An overview of the KL-ONE knowledge representation system." *Cognitive Science*, 9:171-216.
- Bresnan, Joan (ed.) (1982). *The Mental Representation of Grammatical Relations*. Cambridge, MA: The MIT Press.
- Briscoe, E., Copestake, A., and Boguraev, B. (1990). "Enjoy the paper: lexical semantics via lexicology." In *Proceedings, 13th International Conference on Computational Linguistics*, Helsinki, Finland.
- Cardelli, L., and Wegner, P. (1985). "On understanding types, data abstraction, and polymorphism." *ACM Computing Surveys*, 17(4): 471-522.
- Carlson, Gregory (1977). "Reference to kinds in English." Doctoral dissertation, University of Massachusetts.
- Carnap, Rudolf (1956). *Meaning and Necessity*. Chicago: University of Chicago Press.

- Charniak, Eugene, and Goldman, Robert (1988). "A logic for semantic interpretation." In *Proceedings, 26th Annual Meeting of the Association for Computational Linguistics*, Buffalo, NY.
- Chierchia, G. (1989). "Structured meanings, thematic roles, and control." In *Properties, Types, and Meaning*, Volume 2, edited by G. Chierchia, B. Partee, and R. Turner. Dordrecht: Kluwer Academic Publishers.
- Chomsky, Noam (1975). *The Logical Structure of Linguistic Theory*. Chicago: University of Chicago Press.
- Chomsky, Noam (1981). *Lectures on Government and Binding*. Dordrecht: Foris Publications.
- Collins, A., and Quillian, M. (1969). "Retrieval time from semantic memory." *Journal of Verbal Learning and Verbal Behavior*, 9: 240–247.
- Copestake, A., and Briscoe, T. (1991). "Lexical operations in a unification-based framework." In *Proceedings, SIGLEX Workshop on Lexical Semantics and Knowledge Representation*, edited by J. Pustejovsky and S. Bergler, Berkeley, CA.
- Croft, William (1991). *Categories and Relations in Syntax: The Clause-Level Organization of Information*. Chicago: University of Chicago Press.
- Cruse, D. A. (1986). *Lexical Semantics*. Cambridge, U.K.: Cambridge University Press.
- Dixon, R. M. W. (1991). *A New Approach to English Grammar, on Semantic Principles*. Oxford, U.K.: Oxford University Press.
- Dowty, David R. (1979). *Word Meaning and Montague Grammar*. Dordrecht: D. Reidel.
- Dowty, David R. (1985). "On some recent analyses of control." *Linguistics and Philosophy*, 8: 1–41.
- Dowty, David R. (1989). "On the semantic content of the notion 'thematic role'." In *Properties, Types, and Meaning Volume II*, edited by G. Chierchia, B. Partee, and R. Turner. Dordrecht: Kluwer Academic Publishers.
- Elliott, D. E. (1974). "Towards a grammar of exclamations." *Foundations of Language*, 11:231–246.
- Evans, Roger, and Gazdar, Gerald (1989). "Inference in DATR." In *Proceedings, Fourth European ACL Conference*. Manchester, England.
- Evans, Roger, and Gazdar, Gerald (1990). "The DATR papers: February 1990." Cognitive Science Research Paper CSRP 139, School of Cognitive and Computing Science, University of Sussex, Brighton, England.
- Fass, Dan (1988). "Collative semantics: A semantics for natural language processing." MCCS-99-118, Computing Research Laboratory, New Mexico State University.
- Fauconnier, G. (1985). *Mental Spaces*. Cambridge, MA: The MIT Press.
- Fillmore, Charles (1968). "The case for case." In *Universals in Linguistic Theory*, edited by E. Bach and R. Harms. New York: Holt, Rinehart, and Winston.
- Fillmore, Charles (1985). "Construction grammar." Ms.
- Flickinger, D.; Pollard, C.; and Wasow, T. (1985). "Structure-sharing in lexical representation." In *Proceedings, 23rd Annual Meeting of the ACL*. Chicago, IL.
- Fodor, Jerry (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.
- Freed, A. F. (1979). *The Semantics of English Aspectual Complementation*. Dordrecht: Reidel.

- Gazdar, G.; Klein, E.; Pullum, G.; and Sag, I. (1985). *Generalized Phrase Structure Grammar*. Cambridge, MA: Harvard University Press.
- Goodman, Nelson (1951). *The Structure of Appearance*. Dordrecht: Reidel.
- Grice, H. P. (1971). "Meaning." In *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics, and Psychology*, edited by D. Steinberg and L. Jacobovits. Cambridge: Cambridge University Press.
- Grimshaw, Jane (1979). "Complement selection and the lexicon." *Linguistic Inquiry*, 10:279–326.
- Grimshaw, Jane (1990). *Argument Structure*. Cambridge, MA: The MIT Press. in press.
- Gruber, Jeffrey (1965). "Studies in lexical relations." Doctoral dissertation, Massachusetts Institute of Technology.
- Hale, Ken, and Keyser, S. J. (1986). "Some transitivity alternations in English." Lexicon Project Working Papers 7, Center for Cognitive Science, MIT.
- Hale, Ken, and Keyser, S. J. (1987). "A view from the middle." Lexicon Project Working Papers 10, Center for Cognitive Science, MIT.
- Hayes, Patrick (1979). "Naive physics manifesto." In *Expert Systems in the Micro-Electronic Age*, edited by Donald Michie. Edinburgh: Edinburgh University Press.
- Higginbotham, James (1985). "On semantics." *Linguistic Inquiry*, 16:547–593.
- Hinrichs, Erhard W. (1985). "A compositional semantics for aktionarten and NP reference in English." Doctoral dissertation, Ohio State University.
- Hirst, Graeme (1987). *Semantic Interpretation and the Resolution of Ambiguity*. Cambridge: Cambridge University Press.
- Hobbs, Jerry (1982). "Towards an understanding of coherence in discourse." In *Strategies for Natural Language Processing*, edited by W. Lehnert and M. Ringle. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Hobbs, Jerry (1987). "World knowledge and word meaning." In *Proceedings, TINLAP-3*. Las Cruces, NM.
- Hobbs, Jerry; Croft, William; Davies, Todd; Edwards, Douglas; and Laws, Kenneth. (1987). "Commonsense metaphysics and lexical semantics." *Computational Linguistics*, 13:241–250.
- Hobbs, Jerry; Croft, William; Davies, Todd; Edwards, Douglas; and Laws, Kenneth, "The TACITUS commonsense knowledge base." Artificial Intelligence Center, SRI International.
- Hobbs, Jerry; Stickel, Mark; Martin, Paul; and Edwards, Douglas (1988). "Interpretation as abduction." In *Proceedings, 26th Annual Meeting of the Association for Computational Linguistics*. Buffalo, NY.
- Ingria, Robert, and Pustejovsky, James (1990). "Active objects in syntax, semantics, and parsing." In *The MIT Parsing Volume, 1989–1990*, edited by Carol Tenny. Center for Cognitive Science, MIT.
- Jackendoff, Ray (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, MA: The MIT Press.
- Jackendoff, Ray (1983). *Semantics and Cognition*. Cambridge, MA: The MIT Press.
- Katz, Jerrold J. (1972). *Semantic Theory*. New York: Harper and Row.
- Katz, Jerrold J., and Fodor, Jerry (1963). "The structure of a semantic theory." *Language*, 39(2): 170–210.
- Karttunen, Lauri (1971). "Implicative verbs." *Language*, 47: 340–58.
- Karttunen, Lauri (1974). "Presupposition and linguistic context." *Theoretical Linguistics* 1: 181–93.

- Katz, Jerrold, and Fodor, Jerry (1963). "The structure of a semantic theory." *Language*, 39: 170-210.
- Keenan, Edward, and Faltz, Leonard (1985). *Boolean Semantics for Natural Language*. Dordrecht: Reidel.
- Klein, E., and Sag, I. (1985). "Type-driven translation." *Linguistics and Philosophy*, 8: 163-202.
- Krifka, Manfred (1987). "Nominal reference and temporal constitution: Towards a semantics of quantity." FNS-Bericht 17. Forschungsstelle für natürlich-sprachliche System, Universität Tübingen.
- Lakoff, George (1970). *Irregularity in Syntax*. Holt, Rinehart, and Winston.
- Lakoff, George (1971). "On generative semantics." In *Semantics: An Interdisciplinary Reader*, edited by D. Steinberg and L. Jakobovits. Cambridge University Press.
- Lakoff, George (1987). *Women, Fire, and Dangerous Objects*. Chicago: University of Chicago Press.
- Levin, Beth (ed.) (1985). "Lexical semantics in review." Lexicon Project Working Papers Number 1, MIT.
- Levin, B. *Towards a Lexical Organization of English Verbs*. Chicago: University of Chicago Press. Forthcoming.
- Levin, Beth, and Rapoport, T. R. (1988). "Lexical subordination." *Proceedings of CLS 24*, 275-289.
- Levin, Beth, and Rappaport, Malka (1986). "The formation of adjectival passives." *Linguistic Inquiry*, 17:623-663.
- Levin, Beth, and Rappaport, Malka (1988). "On the nature of unaccusativity." In *Proceedings, NELS 1988*.
- Levin, Beth, and Rappaport, Malka. *Unaccusatives*. The MIT Press. forthcoming.
- Lloyd, G. E. R. (1968). *Aristotle: The Growth and Structure of his Thought*. Cambridge University Press.
- McKeon, Richard (1941). *The Basic Works of Aristotle*. Random House.
- Mel'čuk, I. (1988). *Dependency Syntax*. Albany, NY: SUNY Press.
- Michalski, R. S. (1983). "A theory and methodology of inductive learning." In *Machine Learning I*. R. S. Michalski, J. Carbonelli, and T. Mitchell, eds. Palo Alto, CA: Tioga Publishing.
- Miller, George (1985). "Dictionaries of the mind." In *Proceedings, 23rd Annual Meeting of the Association for Computational Linguistics*, Chicago, IL.
- Miller, G. A., and Fellbaum, C. (1991). "Semantic networks of English." *Cognition*, October 1991.
- Miller, George, and Johnson-Laird, Phillip (1976). *Language and Perception*. Cambridge, MA: Belknap. Harvard University Press.
- Moens, M., and Steedman, M. (1988). "Temporal ontology and temporal reference." *Computational Linguistics*, 14(2): 15-28.
- Montague, Richard (1974). *Formal Philosophy: The Collected Papers of Richard Montague*, edited by Richard Thomason. New Haven: Yale University Press.
- Moravcsik, J. M. (1975). "Aitia as generative factor in Aristotle's philosophy." *Dialogue*, 14:622-636.

- Mourelatos, Alexander (1981). "Events, processes, and states." In *Syntax and Semantics: Tense and Aspect*, edited by P. Tedeschi and A. Zaenen. New York: Academic Press.
- Nunberg, G. (1978). *The Pragmatics of Reference*. Bloomington, Indiana: Indiana University Linguistics Club.
- Quillian, M. Ross (1968). "Semantic memory." In *Semantic Information Processing*, edited by M. Minsky. Cambridge, MA: The MIT Press.
- Partee, Barbara, and Rooth, Mats (1983). "Generalized conjunction and type ambiguity." In *Meaning, Use, and Interpretation of Language*, edited by Bäuerle, Schwarze, and von Stechow. Walter de Gruyter.
- Passonneau, Rebecca J. (1988). "A computational model of the semantics of tense and aspect." *Computational Linguistics*, 14(2).
- Pustejovsky, James (1988). "The geometry of events." In *Studies in Generative Approaches to Aspect*, edited by Carol Tenny. Lexicon Project Working Papers 24, MIT.
- Pustejovsky, James (1989a). "Type coercion and selection." Paper presented at West Coast Conference on Formal Linguistics, Vancouver.
- Pustejovsky, James (1989b). "Issues in computational lexical semantics." In *Proceedings, Fourth European ACL Conference*, Manchester, England.
- Pustejovsky, James (1991). "The syntax of event structure." *Cognition*, 41.
- Pustejovsky, James (in press a). "Principles versus criteria: On Randall's catapult hypothesis." In *Theoretical Issues in Language Acquisition*, edited by J. Weissenborn, H. Goodluck, and T. Roeper. Dordrecht: Kluwer Academic Publishers.
- Pustejovsky, James (in press b). "Type coercion and lexical selection." In *Semantics and the Lexicon*, edited by J. Pustejovsky. Dordrecht: Kluwer Academic Publishers.
- Pustejovsky, James (forthcoming). *The Generative Lexicon: A Theory of Computational Lexical Semantics*. Cambridge, MA: The MIT Press.
- Pustejovsky, James, and Anick, Peter (1988). "On the semantic interpretation of nominals." In *Proceedings, COLING-1988*, Budapest.
- Pustejovsky, James, and Bergler, Sapine, eds. (in press). *Lexical Semantics and Commonsense Reasoning*. Berlin: Springer Verlag.
- Pustejovsky, James, and Boguraev, Bran (1991). "Lexical knowledge representation and natural language processing." In *IBM Journal of Research and Development*, 45:4.
- Roberts, R. B., and Goldstein. (1977). The FRL Manual. Technical Report AI Memo 409, MIT Artificial Intelligence Laboratory.
- Scha, Remko J. H. (1983). "Logical foundations for question answering." MS 12.331 Philips Research Laboratories, Eindhoven, the Netherlands.
- Schank, Roger (1975). *Conceptual Information Processing*. Amsterdam: North-Holland.
- Sueren, Pieter (1985). *Discourse Semantics*. Oxford: Blackwell Publishers.
- Smolka, G. (1988). "A feature logic with subsorts." Wissenschaftliches Zentrum der IBM Deutschland, LILOG-Report 33.
- Stump, G. (1981). "Frequency adjectives." *Linguistics and Philosophy*, 4:221-258.