ELSEVIER

Cognition 61 (1996) 127–159

COGNITION

Selectional constraints: an information-theoretic model and its computational realization

Philip Resnik*

Sun Microsystems Laboratories, 2 Elizabeth Drive, Chelmsford, MA 01824-4195, USA

Abstract

A new, information-theoretic model of selectional constraints is proposed. The strategy adopted here is a minimalist one: how far can one get making as few assumptions as possible? In keeping with that strategy, the proposed model consists of only two components: first, a fairly generic taxonomic representation of concepts, and, second, a probabilistic formalization of selectional constraints defined in terms of that taxonomy, computed on the basis of simple, observable frequencies of co-occurrence between predicates and their arguments. Unlike traditional selection restrictions, the information-theoretic approach avoids empirical problems associated with definitional theories of word meaning, accommodates the observation that semantic anomaly often appears to be a matter of degree, and provides an account of how selectional constraints can be learned. A computational implementation of the model "learns" selectional constraints from collections of naturally occurring text; the predictions of the implemented model are evaluated against judgments elicited from adult subjects, and used to explore the way that arguments are syntactically realized for a class of English verbs. The paper concludes with a discussion of the role of selectional constraints in the acquisition of verb meaning.

1. Introduction

Selectional constraints are limitations on the applicability of natural language predicates to arguments. For example, the following exchange with a 5-year-old child makes it clear that, where a green cow is unlikely but nonetheless conceivable, a green idea is not only unlikely but downright unthinkable (Lila Gleitman, personal communication; see Landau and Gleitman, 1985).

(1) (a) Experimenter: Could a cow be green?
  (b) Subject: I think they're usually brown or white.

* Corresponding author: E-mail: philip.resnik@east.sun.com; fax: (508) 250-5067.

PII S0010-0277(96)00722-6

(2) (a) Experimenter: Could an idea be green?
(b) Subject: No, silly! They're only in your head.

The discussion of selectional constraints has a long history, but much of that discussion concerns the truth-theoretic status of sentences in which a selectional constraint is violated, like *The idea is not green* (Horn, 1989). The question of how selectional constraints might be acquired has received less attention, as has the question of why many predicates seem more flexible about their arguments than the examples in the semantics literature would suggest.

What follows is an attempt to answer these questions by presenting a new formal model of selectional constraints and a computational realization of that model. The strategy adopted here is a minimalist one: how far can one get making as few assumptions as possible? In keeping with that strategy, the proposed model consists of only two components: first, a fairly generic taxonomic representation of concepts, and, second, a probabilistic formalization of selectional constraints defined in terms of that taxonomy, computed on the basis of simple, observable frequencies of co-occurrence between predicates and their arguments. Minimizing the representational assumptions simplifies the problem of accounting for how selectional constraints are learned. Formalizing the model in information-theoretic terms leads to an illuminating interpretation of selectional constraints and their flexibility: *how strongly* a predicate selects for an argument is identified with the quantity of *information* it carries about that argument, where information is interpreted in a strict mathematical sense. The model is assessed by means of a computational implementation, comparing its predictions against judgments elected from adult subjects and against naturally occurring data.

The most familiar approach to characterizing selectional constraints is the notion of selection restrictions, introduced as part of Katz and Fodor's (Katz and Fodor, 1964) broader semantic theory based on the notion of defining features. They outlined a decompositional theory of word meaning in which lexical entries specified the features applicable to a particular lexical item – the classic example is the noun *bachelor*, which, among other things, can mean an unmarried man (semantic features HUMAN and MALE) or a young fur seal without a mate (semantic features ANIMAL and MALE). For words that denote predicates, Katz and Fodor proposed that the arguments in their lexical entries be annotated with restrictions identifying the necessary and sufficient conditions that a semantically acceptable argument must meet. Such conditions were represented as Boolean functions of semantic features; for example, (3) gives the selection restrictions on the arguments for the verb *hit* when used as in *The man hits the ground with a rock.*

(3)		
	SUBJECT	HUMAN or HIGHER ANIMAL
	OBJECT	PHYSICAL OBJECT
	INSTRUMENTAL	PHYSICAL OBJECT

As an attempt to explain mental objects, Katz and Fodor's theory was met with criticism of its details as well as of the definitional enterprise as a whole (e.g., McCawley, 1968; McCawley, 1971; Fodor et al., 1980; Armstrong et al., 1983; also see Katz, 1972 for later refinements of the theory and Fodor, 1977, for critical discussion). Identifying restrictions that are both necessary and sufficient, and choosing the primitives themselves, is viewed by many to be an insurmountable problem: Armstrong et al. (1983, p. 268) go so far as to comment, "Generally speaking, it is widely agreed today in philosophy, linguistics, and psychology, that the definitional program for everyday lexical categories has been defeated – at least in its pristine form."[1] Those criticisms notwithstanding, Katz and Fodor's theory – a small number of primitive semantic features, combined into Boolean expressions – offers a uniform explanation for a range of semantic phenomena. In a sentence like *The bachelor hit the baseball*, for example, only a human sense of *bachelor* is natural, since the "young fur seal" sense fails the test on the subject of *hit*. Semantic anomalies arise via the same mechanism when no sense of a word in an argument position can pass the Boolean test, as in *Sincerity hit the baseball*. It is perhaps for that reason that Boolean selection restrictions remain a dominant theory of constraints on predicate–argument combination (e.g., see Allen, 1995) despite widely acknowledged empirical limitations.

An alternative to the Katz–Fodor account is the view that selectional constraints are not a phenomenon of lexical semantics per se, but just a reflection of the more general inferential system underlying language understanding. Johnson-Laird (1983) presents one clear statement of this position, arguing that what appear to be semantic selectional constraints are actually inferences that have been conventionalized because of their frequency and predictability.[2] However, where traditional selection restrictions appear to be overly restrictive, assuming a representational vocabulary too impoverished and rigid to capture the apparent scope and flexibility of real-world mental categories, treating selectional constraints as part of a broader inferential system seems equally problematic for the opposite reason: a general theory of inference must assume the entire representational arsenal that people use in understanding language, ranging from social mores to naive physics. The principle of Occam's Razor suggests that before consigning selectional constraints to the vast, poorly understood territory of general reasoning, we first look for a more restricted model.

The approach to selectional constraints taken here starts with fewer assumptions than the definitional view, resolving many of its empirical difficulties, and at the same time is expressed in a formally precise way so as to avoid the theoretical

[1] See Wilks et al. (1996) for an up-to-date discussion and an argument for the continued use of semantic primitives in a less pristine fashion.

[2] Views consistent with Johnson-Laird's position are found elsewhere, though sometimes less explicitly. For example, in psycholinguistic work on the effects of argument plausibility in on-line processing, one can find traditionally semantic distinctions, particularly the distinction between animates and inanimates, grouped with pragmatic factors (Holmes et al., 1989; Tabossi et al., 1994; MacDonald, 1993). Similarly, computational work on language understanding as a variety of theorem proving (e.g., Hobbs et al., 1993; Alshawi and Carter, 1992) often makes no formal distinction between axioms representing selection restrictions and axioms encoding general factual knowledge.

open-endedness of the inferential approach. The model was designed with the following criteria in mind:

1. It should be possible but not necessary to assume that word meanings are decomposable into definitional features (Fodor et al., 1980; Johnson-Laird, 1983; Armstrong et al., 1983).
2. The model should allow for selectional constraints that make reference not just to a small, privileged semantic vocabulary, but to "any piece of semantic information which may figure in the semantic representation of an item"; for example, the constraints associated with the verbs *devein* (restricted to shrimp or prawns) and *diagonalize* (restricted to matrices) (McCawley, 1968).
3. The model should accommodate the observation that in many cases semantic anomaly – that is, the violation of selectional constraints – appears to be a matter of degree rather than an all-or-nothing phenomenon (Drange, 1966; Fodor, 1977).

An additional goal of the model is an account of how selectional constraints can be learned, subject to the criteria just outlined. The issue of learning can be interpreted in two ways. The first concerns the construction of a model of adult selectional constraints that can be implemented computationally in order to accommodate studies using a realistic range of data rather than a carefully chosen set of "toy" cases. If the basis of the implemented model is not just a small semantic vocabulary, but conceptual information on a larger scale, and especially if the selectional constraints are graded rather than categorical, then introspective methods for constructing it are surely impractical. The second sense of learning is that of language acquisition: how might selectional constraints emerge in a child acquiring language? Here even a practical methodology for model construction is not enough: the account must also require no more prior knowledge than that available to a language learner, and it must not assume the existence of input data to which the learner would not have access. This is true even if the model is thought of only as a starting point for further investigation, as is the case here, rather than as a detailed description of actual mental processes in children.

Because observed data is an integral part of the model, these two aspects of learning – constructing a model of mature mental representations, and simulating one component of the child's verb acquisition process – have requirements that are for the most part consistent with each other. Where they differ most is in their assumptions about the lexical representation of predicates at the time learning takes place. It seems obvious that, for adults, selectional constraints should be treated as part of rich meaning representations associated with distinct verb senses. For children, however, it is not clear which aspects of a verb's lexical representation are already in place at the time its selectional constraints are learned – recent proposals by Gropen (1993), Gleitman and Gillette (1995), and Grimshaw (1994), to be discussed in detail later, suggest that knowledge of selectional constraints may in fact be instrumental in acquiring those lexical representations. In an attempt to address the learning issue, therefore, the strategy adopted here has been to steer a middle course: model adults, but minimize assumptions about lexical representation.

SELECTIONAL CONSTRAINTS

2. Model

2.1. Informal description

Intuitively, selectional constraints specify what is or is not an appropriate argument for a particular predicate. The idea behind the model is a simple one. In effect, one asks: compared to predicates in general, how much does this particular predicate appear to influence the conceptual class of the words that appear as its argument, as measured by observed co-occurrence frequencies? If the predicate imposes a strong selectional constraint, then the observed frequencies of some conceptual classes of argument will be noticeably greater or less for this predicate than they would be on average. If the predicate selects only weakly, the frequency distribution of its arguments will differ less from what would be expected on average.

The remainder of this section develops this idea in greater detail. In particular, what constitutes a conceptual class of arguments? What is meant by *predicate*, and how does it relate to theories of lexical representation? How can one quantify the influence of a predicate on the frequency distribution of its arguments? And finally, how does learning take place in such a model?

2.2. Representation of arguments and predicates

The first component of the model is a conceptual taxonomy, or semantic network (Sowa, 1991; Lehmann, 1992), in which classes are related by subsumption. For example, such a network might identify BEVERAGE as a subclass of LIQUID and as a superclass of WINE. In contrast to definitional semantic features, this taxonomy is intended to capture conceptual information, and therefore it may also include classes that are specific to a particular language or culture.

In order to keep representational assumptions to a minimum, taxonomic classes will be thought of as collections of unanalyzed word meanings. That is, one can think of BEVERAGE not only as a word meaning itself, but also as a label identifying a set that contains WATER, WINE, COFFEE, etc., where a word in capital letters denotes some unspecified mental representation. The class LIQUID will then be a proper superset, containing not only the word-concept BEVERAGE and all members of the set identified by the label BEVERAGE, but also such non-beverages as OIL and ANTIFREEZE. A word meaning may belong to any number of conceptual classes – for example, COOKIE might be categorized both with other foods (BREAD, etc.) and with other small objects (CRAYON, etc.).

This minimal formalization of taxonomic classes is consistent with more elaborated lexical representations of arguments, such as those based on "qualia structure" (Pustejovsky, 1995). For example, Pustejovsky et al. (1993) give the following representations for *book* and *tape* (in the sense of a computer-readable magnetic tape):

$$(4)\quad \begin{bmatrix} \textbf{book(x,y)} \\ \text{CONST} = \textbf{information(y)} \\ \text{FORMAL} = \textbf{physobj(x)} \\ \text{TELIC} = \textbf{read(T,w,y)} \\ \text{AGENTIVE} = \textbf{write(T,z,y)} \end{bmatrix} \quad \begin{bmatrix} \textbf{tape(x,y)} \\ \text{CONST} = \textbf{information(y)} \\ \text{FORMAL} = \textbf{physiobj(x)} \text{ and } \textbf{2-dimensional(x)} \\ \text{TELIC} = \textbf{contain(S,x,y)} \\ \text{AGENTIVE} = \textbf{write(T,z,y)} \end{bmatrix}$$

According to the simpler taxonomic scheme assumed here, class membership would be used to account for most of the same information. For example, the classes illustrated in (5) might represent the sets of objects that are (a) created by writing information, (b) physical objects, and (c) containers of some kind.

(5) (a) {BOOK, DIARY, DISKETTE, NOVEL, TAPE,...}
(b) {ANCHOVY, BICYCLE, BOOK, DISKETTE, NOVEL, RECORD, TAPE, ...}
(c) {BOWL, BOX, CUP, DISKETTE, TAPE, ...}

Crucially, this taxonomic view of concepts permits but does not require semantic decomposition: the concept BOOK is implicitly identified as a physical object, intended to be read, and so forth, by virtue of its juxtaposition with other concepts, rather than by the choice of semantic components in its meaning representation. Miller (1990) distinguishes these two forms of representation using the terms *constructive* and *differential* – a constructive lexical theory must support accurate reconstruction of concepts by a person or machine not already in possession of those concepts, whereas in a differential theory, it is assumed that the goal is to differentiate among concepts that are already known. Formalizing arguments according to a differential style of representation leads to a model consistent with more constructive theories without relying on the details of how arguments are represented in the mental lexicon.[3]

In order to plausibly support an account of how selectional constraints are acquired, it should be evident that this first, taxonomic component of the model requires two relatively weak assumptions with regard to the mental representation of arguments. The first is that learners already have a grasp of the noun lexicon at the time selectional constraints are being learned. This is suggested by the observation that nouns precede verbs in acquisition (Nelson, 1973; Gentner, 1982), and it is also supported by evidence that learning to map noun forms to noun concepts is a relatively easy task (Gleitman and Gillette, 1995). The second assumption is that the learner has already organized concepts into discrete classes.

[3] Another advantage of the differential approach is methodological rather than theoretical: in contrast to computational lexicons built according to constructive theories, which cover at best a small subset of the English language, the differential framework has been used to construct a very broad coverage on-line lexicon for English (Miller, 1990), making it possible to adopt a research methodology involving computational simulation on a large scale.

Although there is less clear evidence for this, children at least as young as 3 years old can classify pictures of objects in the same manner as adults for basic level categories such as TABLE and FISH, and sorting objects into superordinate categories such as FURNITURE and ANIMAL reaches adult competence by about age 8 (Rosch et al., 1976).[4]

For predicates, as was the case for arguments, the formal model makes as few representational assumptions as possible. The literature on lexical representations for predicates has had little to say about selectional constraints or how they relate to the conceptual content of their arguments; rather, its focus tends to be on those aspects of lexical representation that determine how arguments are realized syntactically (e.g., Dowty, 1991; Levin, 1993; Pinker, 1989). To adopt the terminology of Grimshaw (1993, 1994), most work on the representation of predicates concerns *semantic structure*, whereas selectional constraints are a matter of *semantic content*. The model, therefore, treats predicates as very abstract formal objects, making no reference to their internal structure.

The issue here is not whether such internal structure exists, which it certainly does, but whether the details of that structure need to be explicitly represented in the formal model of selectional constraints in order to make reasonable predictions. A lexical theory such as Jackendoff's (Jackendoff, 1983) situates selectional constraints as information appearing in the context of a rich representation of the predicate's meaning, such as the annotation LIQUID appearing as a constraint on one argument of the verb *drink*.

$$(6)\quad \left[\begin{array}{l} \text{drink} \\ \text{V} \\ \underline{\qquad}\langle \text{NP}_j\rangle \\ [_{Event}\ \text{CAUSE}([_{Thing}\ ]_i,\ [_{Event}\ \text{GO}([_{Thing}\ \text{LIQUID}]_j, \\ \quad [_{Path}\ \text{TO}([_{Place}\ \text{IN}([_{Thing}\ \text{MOUTH OF}([_{Thing}\ ]_i)])])])])] \end{array}\right]$$

The model here is consistent with such theories: it plays the same role as the annotation LIQUID does in (6). Crucially, however, the model makes no assumptions about the nature of the lexical representation within which the constraint is situated, and by hypothesis it can make reasonable predictions without doing so.

As a final representational issue, there is the question of whether selectional constraints should be thought of as associated with semantic predicates, as is typically assumed, or with some more abstract representation that takes multiple senses into account. For purposes of modeling adult lexical knowledge, selectional constraints are obviously associated with word senses – for example, *Gielgud played Hamlet* is a fine sentence on the theatrical reading of *play*, but a selectional

[4] It is worth noting that a taxonomy of this kind can be viewed as implicitly encoding inferential relationships: one can define class membership in terms of shared entailments among the members of the class, and subsumption by inheritance of those entailments. In that respect, the approach taken here adopts the inferential view of selectional constraints, albeit in a much more constrained fashion. More discussion on this point can be found elsewhere (Resnik, 1993).

violation in the "play on a musical instrument" sense of the verb. However, once the element of learning is introduced, the situation becomes less clear, since learners observe not word senses but word forms. One would like to appeal to the observation that the arguments of a verb often resolve which sense of the verb is being used, but this would constitute a circular argument: in order to learn selectional constraints, look at the co-occurrence of arguments with senses of the verb; in order to decide which sense of a verb is correct for a given observed instance, choose the sense for which the argument best matches the selectional constraints. And even this presupposes that the learner has already discovered what the possible senses are for a given verb; however, that learning process may itself involve knowledge of selectional constraints, a possibility that will be taken up in the General Discussion.

For these reasons, the model to be described will be expressed not in terms of semantic predicates (i.e., distinct word senses) but more abstractly by conflating multiple word senses into a single representation. For example, the model will treat PLAY as if its lexical representation conflates the senses of playing a musical instrument, playing a theatrical role, playing a game, and so forth, with the result being a selectional constraint that is multimodal. It must be emphasized that nothing *prevents* the model from expressing distinct selectional constraints for distinct word senses – one need only think of those senses as constituting distinct predicates in the formal model, identifiable as PLAY_1, PLAY_2, etc. – but in order to deal sensibly with issues of learning it is impossible to assume that the learner has already succeeded in acquiring those distinctions.

2.3. Formalization of selectional constraints

Given a representation of argument concepts and predicates as just discussed, the second component of the model characterizes selectional constraints in terms of a probabilistic relationship between predicates and conceptual categories or classes. Intuitively, the idea is this: rather than obeying restrictions or hard constraints on applicability, a predicate preferentially associates with certain classes of arguments. To state this another way, preferences constitute the effect that the predicate has on what appears in an argument position. For example, the adjective *blue* does not *restrict* itself to arguments having a tangible surface – the sky is blue, and so is ocean water even below the surface. Rather, the effect of the predicate is that its arguments tend to be physical entities and to have surfaces. Similarly, the verb *admire* has an effect on what appears as its subject: these tend to be physical, animate, human, capable of the higher psychological functions, and so forth, though no Boolean combination of these properties need be both necessary and sufficient.

Formally, let P be a random variable ranging over the set $\{p_1, ..., p_m\}$ of predicates under consideration. Let C be a random variable ranging over the set $\{c_1, ..., c_n\}$ of classes in the taxonomy described above. Finally, let r denote the argument position of interest. Given this probabilistic framework, the intuitive notion of preference can now be phrased more precisely as the following question:

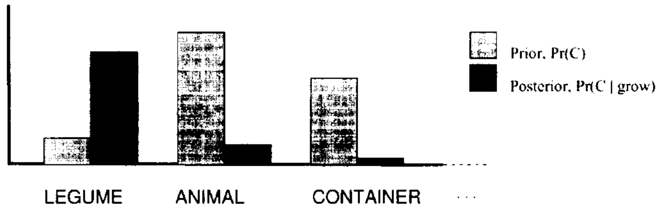

Fig. 1. Example of a prior distribution and a posterior distribution.

what effect does the choice of a particular predicate $P=p_i$ have on the distribution of C in argument position r?

Fig. 1 illustrates how this might work for a particular predicate, the verb *grow*, with respect to its direct object argument.[5] The light bars represent part of what the distribution of argument classes might be, regardless of the particular predicate. Independent of the verb, some classes are, a priori, simply more likely to be referred to in direct object position, and some less likely. For example, in the absence of any other information, animals might be more likely to be mentioned in direct object position than legumes. However, given the particular verb *grow*, this distribution changes to the one shown by the dark bars: some classes (e.g., animals) become much less likely, and others (e.g., legumes) become much more likely.

In the present model, it is this relationship, the difference between the *prior* distribution, $\Pr(c)$, and the *posterior* distribution, $\Pr(c|p_i)$, that constitutes selectional preference. On this account, the features, properties, or inferences that govern selectional constraints remain entirely hidden: selectional relationships are characterized entirely by the probabilistic relationship between a predicate and the classes of its arguments.

A difference between two probability distributions can be expressed in precise terms using an information-theoretic measure known as *relative entropy* (Kullback and Leibler, 1951; Cover and Thomas, 1991), which is defined as:

$$D(p||q) = \sum_x p(x) \log \frac{p(x)}{q(x)}.$$

Intuitively, if probability distribution p is interpreted as the "truth" and distribution q is interpreted as an approximation of the true distribution, then the relative entropy D(p||q) measures the amount of extra information one would need to add to the approximation in order to make it fit the truth perfectly.

[5] As noted earlier, the argument position r could just as well be stated with respect to thematic roles (or other predicate–argument relationships) rather than surface syntactic relationships such as subject and object; the choice of relationship has no bearing on the formal definition of the model. For the sake of readability, r is suppressed in the figure and all formal definitions. For example, the frequency with which a noun n occurs as the direct object of a verb v will be denoted freq(v, n) rather than $\text{freq}_{\text{obj}}(v, n)$. Also note that the three conceptual categories in the figure are just an illustrative sample: the model involves using the prior and posterior probabilities of *all* the categories in the argument taxonomy.

The prior distribution of classes, $\Pr(c)$, represents an uninformed approximation of what the distribution of arguments looks like, one that does not take the predicate into account at all. The posterior, $\Pr(c|p_i)$, is the true distribution of argument classes with respect to a particular predicate p_i. So, treating the former as q and the latter as p, the difference between the two distributions is quantified as:

$$\begin{aligned} S(p_i) &= D(\Pr(c|p_i) \| \Pr(c)) \\ &= \sum_c \Pr(c|p_i) \log \frac{\Pr(c|p_i)}{\Pr(c)} \end{aligned}$$

I will call this quantity *selectional preference strength*.[6]

Notice that, in this model, the selectional preference strength of a predicate is not just a number, but a number with a precisely specified meaning. Treating $\Pr(c|p_i)$ as truth and $\Pr(c)$ as approximation, the selectional preference strength of p_i translates as the cost, in information, of not taking the predicate into account. Therefore, in a very direct way, the selectional preference strength of a predicate can be understood as the amount of information it carries about its argument.

As illustrated in Fig. 1, selectional preference is characterized here as a relationship between a predicate and the entire conceptual space of arguments, and selectional preference strength reduces that relationship to a single quantifiable value. Neither of those succeeds in answering the most frequently asked question concerning selectional constraints, namely to what extent a *particular* conceptual class "fits" as the argument to a given predicate. For that purpose, it is useful to observe that each term in the sum that defines selectional preference strength represents the contribution of a single conceptual class. Classes that become more likely given the predicate contribute a positive amount, and classes that become less likely contribute a negative amount. This suggests that the selectional relationship between a *particular* class and a predicate can be expressed in terms of the relative contribution that class makes to the overall selectional preference strength.

To express this formally, let the *selectional association* between a predicate p_i and a class c be defined as:

$$A(p_i, c) = \frac{\Pr(c|p_i) \log \frac{\Pr(c|p_i)}{\Pr(c)}}{S(p_i)}$$

Unlike selectional preference strength, which is always greater than or equal to zero, the selectional association between a predicate and a class can also have a negative value, indicating the extent to which that class is *dispreferred* as an argument. It is selectional association that serves as the source of predictions regarding acceptability and anomaly: the scale from negative to positive values of selectional association is interpreted as the degree of acceptability. It is inherently quantitative in its definition, being based on probabilities, which reflects the

[6] For a discussion of related measures in a different setting, see Smyth and Goodman (1992).

observation that there is something graded rather than categorical about judgments of semantic anomaly (Drange, 1966).

3. Implementation

In realizing the model of selectional constraints as a computational implementation, the noun database from WordNet (version 1.2; Miller, 1990) was used as the computational representation of adults' taxonomic knowledge about the conceptual categories of arguments. WordNet was suitable for this task because of its organization (disambiguated word senses organized via class subsumption), its scale (dictionary-level coverage of English), and the principles underlying its construction (specifically, the differential lexical theory on which it is based, and its attempt to draw a reasonable line between lexical concepts and general knowledge).[7]

The probabilities underlying the information-theoretic model were approximated for the verb–object relationship by collecting frequencies of co-occurrence for verbs and their objects from three different sources, effectively yielding three different simulations of the model.[8] These sources included the following:

1. The Brown corpus of American English (Francis and Kučera, 1982), appearing in parsed form within the Penn Treebank (Marcus et al., 1993); direct objects were extracted automatically according to syntactic criteria.
2. Parental turns from transcribed speech in the CHILDES collection of parent–child interactions (MacWhinney and Snow, 1985); direct objects were identified automatically through heuristic procedures.[9]
3. Verb–object norms collected from human subjects in an unpublished study by Anne Lederer at the University of Pennsylvania.

The joint probabilities in the model were estimated from observed verbs and objects as follows:

$$\Pr(v,c) = \frac{1}{N} \sum_{n \in \mathrm{words}(c)} \frac{1}{|\mathrm{classes}(n)|} \mathrm{freq}(v, n)$$

[7] See Resnik (1993, Ch. 2) for a more detailed discussion.

[8] In all instances, the head of a noun phrase in direct object position was treated as if it were the direct object. As McCawley (1968) convincingly argues, selectional constraints concern not the lexical heads of argument constituents but the concepts denoted by the entire phrase – for example, *a toy soldier* should be treated as a kind of TOY – but such cases appeared relatively infrequently.

[9] All the parental data in English then available in CHILDES were merged; these included data gathered by the following researchers: Bates, Bernstein, Bloom, Bohannon, Braine, Brown, Clark, Evans, Garvey, Gathercole, Gleason, Hall, Higginson, Howe, Kuczaj, MacWhinney, Sachs, Snow, Suppes, Vanhouten, and Warren. MacWhinney and Snow (1985) discuss details of the CHILDES collection.

where v is a verb; c is a class, that is, a node in the WordNet noun taxonomy; words(c) denotes the set of nouns for which any sense is subsumed by c; classes(n) denotes the set of WordNet classes to which noun n belongs (in any of its senses); freq(v, n) is the number of times n appeared as the object of v; and N is the total number of instances in the observed sample ($= \Sigma_{v',n'}$ freq(v', n')). This estimation procedure can be interpreted as distributing the "credit" for an ambiguous noun uniformly across the set of meanings it might be used to express, a necessary step in the absence of a computationally feasible way to discriminate among the word senses of arguments in advance (Yarowsky, 1992; Resnik, 1993).

4. Qualitative behavior

Before turning to experiments using the implemented model, it is interesting to briefly explore its qualitative behavior. To begin with, Table 1 shows the set of verbs used in the norming study, selected because they occurred frequently in a collection of parent–child interactions (Anne Lederer, personal communication). Each verb appears together with the selectional preference strength for its direct object, as measured using probability estimates automatically learned from the three collections discussed above.[10]

Table 1
Strength of selectional preference for direct objects

Verb	Strength			Verb	Strength		
	Brown	CHILDES	Norms		Brown	CHILDES	Norms
pour	4.80	2.30	2.57	explain	2.39	4.41	2.20
drink	4.38	2.38	2.83	read	2.35	2.58	1.81
pack	4.12	3.71	1.75	watch	1.97	1.44	1.86
sing	3.58	3.15	2.63	do	1.84	–	2.21
steal	3.52	2.28	1.34	hear	1.70	1.67	1.71
eat	3.51	1.15	2.47	call	1.52	0.95	2.39
hang	3.35	2.03	1.96	want	1.52	0.70	1.71
wear	3.13	2.02	2.30	show	1.39	1.83	1.42
open	2.93	2.41	1.88	bring	1.33	0.88	1.04
push	2.87	1.77	1.98	put	1.24	0.40	1.34
say	2.82	0.94	2.56	see	1.06	0.48	1.54
pull	2.77	1.55	2.22	find	0.96	0.71	1.30
like	2.59	0.89	1.30	take	0.93	0.74	1.28
write	2.54	2.33	2.18	get	0.82	0.28	1.17
play	2.51	2.13	2.64	give	0.79	1.18	1.81
hit	2.49	1.31	1.91	make	0.72	0.77	1.58
catch	2.47	1.67	1.92	have	0.43	–	1.23

[10] The verbs *do* and *have* were excluded from the CHILDES sample because in that collection there was no way to automatically determine whether an observed instance represented use as a verb or use as an auxiliary.

The behavior in Table 1 suggests that reasonable selectional constraints can in fact be acquired from observed data with only limited representational assumptions. At least qualitatively, the model appears to have captured coarse-grained intuitions about strength of preference – for example, my own intuitions suggest that the top third contains a preponderance of verbs that clearly constrain their objects strongly (e.g., *drink, eat, sing*), the bottom third contains a preponderance of verbs that select quite weakly (e.g., *find, take, get*), and the middle third contains intermediate cases (e.g., *hit, watch, want*). At the same time, the simple ordering draws attention to the shift from categorical notions of semantic well-formedness toward a graded characterization of how predicates and arguments relate. For example, in principle practically any physical object can appear as the object of *push*, but the model has assigned the verb a strong selectional constraint, reflecting its sensitivity to an observed *tendency* for *push* to appear with one kind of object rather than another – in the Brown corpus data, things that are pushed are most likely to be buttons (27%), with cars (7%) and boats (6%) in second and third place.

Fig. 2 provides a qualitative illustration of the *content* of selectional preference for two of the verbs, *eat* and *find*, as computed using the Brown corpus as a learning sample. The figure shows the "selectional profile" for the direct object argument of each verb. WordNet classes are lined up along the x-axis, and the vertical bar for each indicates its selectional association as direct object of the verb. As one might expect, the selectional profile for *eat* illustrates a selectional pattern that is greater in overall magnitude than the profile for *find*, a visual impression that is confirmed by its higher value for selectional preference strength. Moreover, its profile is more specific: the class with the highest value describes the conceptual category of foods, and some other classes with particularly high values include substances and meals. In contrast, the selectional profile for *find* shows a far weaker and less specific pattern of preference. Notably, however, this is not equivalent to saying that *find* places no constraints at all on its direct object. While

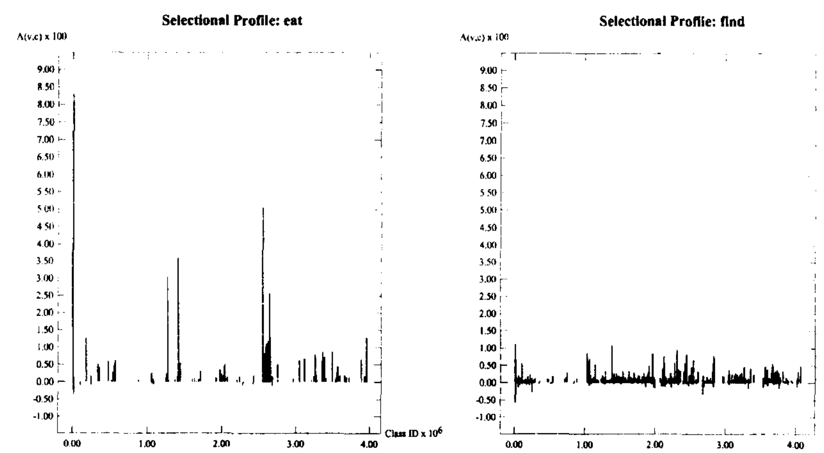

Fig. 2. Selectional profiles for *eat* and *find*.

no clear pattern is evident among classes with positive values of selectional association, the two WordNet classes most *dispreferred* as objects of *find* are those describing actions and states.

EXPERIMENTS

5. Comparison with plausibility judgments

Going beyond a qualitative impression of the model, a more systematic empirical test arises in the context of research into on-line processes during sentence comprehension. A number of researchers have explored the role that argument plausibility plays in processes such as local syntactic disambiguation.

(7) (a) The speaker proposed by the group would work perfectly for the program.
(b) The solution proposed by the group would work perfectly for the program.

Example (7) illustrates how argument plausibility can affect disambiguation decisions. In (7a), there is a tendency to interpret *speaker* as the agent of *propose*, leading to difficulty when subjects reach the disambiguating *by* phrase; (7b) causes no such problem because *solution* is implausible as the agent of *propose* (Trueswell et al., 1994).

In general, the plausibility data for such experiments are obtained by pretesting. For example, Holmes et al. (1989) evaluate argument plausibility by asking subjects to rate sentences like (8) on a 1-to-5 scale.

(8) (a) The mechanic warned the driver.
(b) The mechanic warned the engine.

Although selectional constraints are certainly not the only factor involved in assessing plausibility, they can be expected to play a role in subjects' plausibility judgments. Therefore one way to assess the performance of a computational model like the one proposed here is to compare its "judgments" of selectional fit against the plausibility ratings elicited from human subjects.

EXPERIMENT 1

The first experiment tested the model's ability to distinguish plausible from implausible direct objects of verbs, using data from the study by Holmes et al. (1989). As part of their study, Holmes et al. selected a set of 16 verbs having a bias for NP complements, and constructed for each verb a pair of sentences, one involving a plausible object and the other an implausible object, where plausibility was initially judged according to the experimenters' intuitions. For example, plausible and implausible objects for the verb *warn* are given above in example (8). The intuitive judgments were then confirmed by having a set of subjects rate the sentences for plausibility on a scale from 1 (low plausibility) to 5 (high plausibility). Holmes et al. report a mean rating of 4.5 for the sentences containing plausible objects, and a mean rating of 2.2 for implausible objects, a statistically significant difference.

The goal of Experiment 1 was simply to verify that the model's ratings of selectional fit – as measured by selectional association – also yields a significant difference between the two groups.

6. Materials

Verb–object pairings were extracted by hand from the 16 pairs of sentences in Holmes et al. (1989, Appendix 2). Verbs and nouns were manually reduced to their root forms, where necessary.

7. Method

The Brown corpus was used as a learning sample, as described above. For each pairing of a verb v and direct object n, the selectional association $A(v, c)$ was calculated for each WordNet class c to which n belongs. The greatest such value of selectional association, denoted A_{max}, was assigned as the model's rating for the pairing (v, n).

8. Results and discussion

Table 2 shows the verb–object pairings. Each pairing is shown with its rating according to the model (Assoc $= A_{max} \times 100$), together with a short description of the class that maximized the value of selectional association (Class). For example, the verb–object combination *read article* is assigned a rating of 6.80, and the class responsible for that value is ⟨writing⟩, a WordNet class comprising "anything expressed in letters; reading matter."[11] The mean ratings for plausible and implausible objects are respectively 2.69 and 1.45, a statistically significant difference (Mann–Whitney $U = 77.5$, $p < .05$).

Holmes et al. report that subjects produced a lower average rating for the implausible than the plausible objects in 15 of the 16 cases; the model arrives at the correct ordering in 11 of the 16 cases, though two of those (*teach* and *understand*) are of doubtful reliability. The shortcomings of the model arise from several sources. One obvious problem is word sense ambiguity: unlike human

[11] Phrases in angle brackets constitute brief descriptions of the content of WordNet classes, used rather than numerical identifiers for convenience of exposition.

Table 2
Selectional ratings for plausible and implausible direct objects

Verb	Plausible			Implausible		
	Object	Assoc.	Class	Object	Assoc.	Class
see	friend	5.79	⟨entity⟩	method	−0.01	⟨method⟩
read	article	6.80	⟨writing⟩	fashion	−0.20	⟨activity⟩
find	label	1.10	⟨abstraction⟩	fever	0.22	⟨psych. feature⟩
hear	story	1.89	⟨communication⟩	issue	1.89	⟨communication⟩
write	letter	7.26	⟨writing⟩	market	0.00	⟨commerce⟩
urge	daughter	1.14	⟨life form⟩	contrast	1.86	⟨act⟩
warn	driver	4.73	⟨person⟩	engine	3.61	⟨entity⟩
judge	contest	1.30	⟨contest⟩	climate	0.28	⟨state⟩
teach	language	1.87	⟨cognition⟩	distance	1.86	⟨psych. feature⟩
show	sample	1.44	⟨psych. feature⟩	travel	0.41	⟨happening⟩
expect	visit	0.59	⟨act⟩	mouth	5.93	⟨entity⟩
answer	request	4.49	⟨speech act⟩	tragedy	3.88	⟨communication⟩
recognize	author	0.50	⟨entity⟩	pocket	0.50	⟨entity⟩
repeat	comment	1.23	⟨communication⟩	journal	1.23	⟨communication⟩
understand	concept	1.52	⟨cognition⟩	session	1.51	⟨social relation⟩
remember	reply	1.31	⟨statement⟩	smoke	0.20	⟨article of commerce⟩

subjects, the computer program had access neither to the full sentence context of the verb–object pairings, nor to more general background knowledge, either of which could indicate that a noun was being used in a particular sense. As a result, the most highly rated class in some cases represents a thorough misconstrual of the argument. For example, *answer tragedy* is rated as highly as it is because, given the verb *answer*, the sense of *tragedy* most strongly selected for is that of a dramatic composition, which WordNet classifies as a form of written communication and therefore communication.[12] When the model is presented with manually disambiguated items, using the full sentence as context to select a WordNet sense for the argument, and the selectional association then recalculated accordingly, the difference between the mean ratings is much clearer (means of 2.53 for plausible items vs. 0.96 for implausible items, Mann–Whitney $U=49$, $p<.0025$) and the plausible item is assigned a higher score than the implausible item in 13 of the 16 cases.

In addition, some of the counterintuitive items may reveal limitations of association measure itself. For example, engine is rated as a rather plausible object of *warn*, as will be any other physical entity, because (a) *warn* tends to co-occur with people, which are physical entities, (b) *warn* does not tend to co-occur with

[12] Interestingly, the selectional constraints alone do perform a limited form of sense disambiguation: *article* is taken to denote a piece of written text (rather than a word like *the* or *an*, or a man-made object), *story* is taken to denote a report or narrative (rather than a floor of a building), *driver* is interpreted as a person rather than as a kind of golf club, and so forth. This is a statistical realization of the same phenomenon that Katz and Fodor (1964) captured by having Boolean selection restrictions rule out inappropriate combinations of senses, though, notably, in this experiment the operative constraints were learned automatically from naturally occurring text.

direct objects that are not physical entities, and therefore (c) the posterior probability Pr(⟨entity⟩|*warn*) is much higher than the prior probability Pr(⟨entity⟩) (see equation 1), leading to a high value of selectional association for class ⟨entity⟩. As defined, the model has an insufficient basis for deciding that the property of being a person is more important than the property of being an entity, where objects of *warn* are concerned. This may be a case where assuming too little, in particular ignoring the verb's semantics, could be limiting the model's accuracy.

Finally, it is worth noting that the ratings assigned by the model are not simply a matter of co-occurrence frequency. Of the 16 items in the group of plausible objects, 10 never occurred with the corresponding verb in the learning sample at all. As an example, the object in the pair *warn driver* receives its rating because, given the objects of *warn* that *did* occur, the selectional profile of the verb assigned ⟨person⟩ a higher value of selectional association than any other conceptual category to which *driver* belongs. Thus the conceptual organization of arguments – in the implementation, the WordNet noun taxonomy – is playing a crucial role. Nor is simple co-occurrence frequency by conceptual category responsible for the ratings: the class ⟨entity⟩ ("something having concrete existence; living or nonliving") has a greater co-occurrence frequency with *warn* than does ⟨person⟩, since the former is a more generic concept than the latter, yet the values of selectional association predict, correctly, that people are still a better semantic fit as objects of warn than are physical entities in general.

EXPERIMENT 2

The second experiment focused not on a hard distinction between plausible and implausible arguments, but on a comparison of the magnitudes of selectional association and typicality ratings elicited from human subjects. The quantitative nature of the model is based on the premise that human judgments of the "fit" between predicates and arguments are graded rather than categorical; therefore one would predict that, to the extent that the model is accurate and the simulation faithful, quantitative "judgments" of fit (selectional association) should correlate with the quantitative typicality judgments elicited from adult subjects.

9. Materials

Twenty-eight sentences from a study by Trueswell et al. (1994, Appendix 3, Experiment 1) all had the same form as example (7).[13] The verb in the embedded clause and the modified noun were extracted by hand from the sentences and

[13] Since Holmes et al. (1989) provide only mean plausibility ratings for the two groups, rather than mean ratings by item, it was not possible to look for a correlation between selectional association and the human ratings using the data from Experiment 1.

treated as a verb–object pairing. Verbs and nouns were manually reduced to their root forms, where necessary. For example, (7b) yielded *solution* as a candidate object for *propose*. One pairing was excluded because the verb, *grade*, never appeared with a surface direct object in the Brown corpus.

As Trueswell et al. (1994, Section 3.2.3) report, typicality ratings for the pairs were obtained as part of a large norming project conducted at the University of Southern California, involving a total of 107 subjects. Subjects were asked to rate verb–noun pairings by rating questions like "How typical is it for evidence to be examined by someone" on a 7-point scale, with a rating of 1 indicating "not typical at all" and 7 indicating "very typical."

10. Method

Same as Experiment 1.

11. Results and discussion

Table 3 shows each verb–object pairing, together with the mean typicality ratings assigned by subjects in the norming study and the ratings assigned by the model. There is a statistically significant correlation between the two sets of ratings ($r=0.46$, $F(1, 25)=6.78$, $p(F)<.02$); the scatterplot in Fig. 3 shows the relationship between the two sets. The correlation remains significant if the two points at the lower left corner of the plot are eliminated as outliers ($r=0.44$, $F(1, 23)=5.60$, $p(F)<.03$), though it degrades if the point at the extreme upper right is removed also ($r=0.37$, $F(1, 22)=3.57$, $p(F)<.08$).

The correlation between the model's ratings, based on selectional association, and human subject ratings, based on judgments of typicality, provides some support for the model as a quantitative approximation of selectional constraints, though the match is far from perfect. As in Experiment 1, the model is limited by the unavailability of word sense information; however, unlike in Experiment 1, manual disambiguation of the arguments for this set of items had no appreciable effect on the result ($r=0.45$, $F(1, 25)=6.43$, $p(F)<.02$). In part, this reflects some choices made in constructing WordNet that left certain senses simply unavailable to the model – for example, in this version of WordNet *poultry* appears in the taxonomy only as a kind of animal, not as a food, which has a marked effect on its rating. However, as in Experiment 1, shortcomings of the model may also reflect the inherent limits that result from assuming as little as possible, or the reliance on a limited training sample, or both; for example, subjects' ratings of typicality undoubtedly reflect implicit inferences made using the mental model they construct of a hypothetical situation, given a sentence to be rated (Johnson-Laird, 1983), and that information is unavailable to the model presented here.

That said, Experiments 1 and 2 provide an empirical assessment of the model in

Table 3
Typicality ratings data on items from Trueswell et al. (1994)

Verb	Object	Rating	Assoc.	Class
eat	poultry	6.6	3.54	⟨entity⟩
do	work	6.4	12.03	⟨entity⟩
examine	evidence	6.3	0.78	⟨content⟩
throw	ball	6.2	4.52	⟨entity⟩
see	truck	6.1	5.79	⟨entity⟩
write	letter	6.1	7.26	⟨writing⟩
grow	crop	6.0	2.78	⟨object⟩
break	vase	5.9	1.02	⟨artifact⟩
describe	necklace	5.9	0.00	⟨necklace⟩
draw	poster	5.9	0.20	⟨horse⟩
take	money	5.6	0.02	⟨money⟩
expect	package	5.5	5.93	⟨entity⟩
transport	gold	5.5	8.62	⟨substance⟩
lift	bricks	5.4	5.10	⟨entity⟩
study	painting	5.4	0.79	⟨art⟩
show	wallpaper	5.3	0.00	⟨wallpaper⟩
steal	van	5.2	1.86	⟨object⟩
request	equipment	5.1	3.72	⟨entity⟩
scratch	sofa	4.8	1.83	⟨object⟩
want	account	4.6	0.83	⟨communication⟩
identify	jewelry	4.4	0.46	⟨decoration⟩
select	recipe	4.4	0.00	⟨message⟩
choose	computer	4.3	2.84	⟨entity⟩
recognize	van	4.1	0.50	⟨entity⟩
attack	power_plant	3.2	0.00	–
love	textbook	1.9	0.00	⟨publication⟩
capture	valley	1.8	0.00	⟨valley⟩

a way that previous proposals have not. Given an abstract formal characterization of the model, a computational simulation was derived using linguistic input (the Brown corpus) and a conceptual organization (WordNet) not designed for this specific task, and the simulation was then used to provide ratings evaluated objectively against human judgments, where those judgments were elicited experimentally for an entirely different purpose. The resulting predictions appear generally consistent with adult knowledge of selectional constraints.

12. Analysis of implicit object alternations

In a second set of experiments, the model of selectional constraints was employed in a rather different way. Unlike Experiments 1 and 2, which sought to evaluate the output of the model as directly as possible against human judgments, here the model of selectional constraints was used to investigate a linguistic problem concerning the way in which certain verbs in English express their arguments. The goal was not a simple quantitative result, such as the correlation

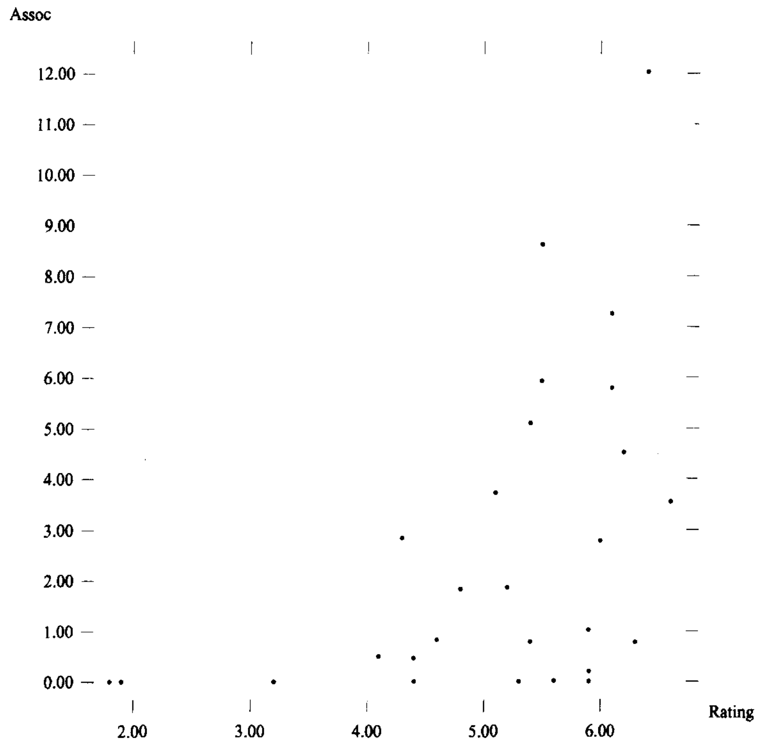

Fig. 3. Selectional association ratings versus typicality ratings.

with human judgments in Experiment 2, but rather an improved understanding of the role selectional constraints might play in linguistic behavior, using the formal model (and its computational realization) as an exploratory tool.

A well- known phenomenon of English is the ability of some verbs to appear without their direct objects, as in (9b) and (10c):

(9) (a) John ate his dinner before Mary arrived.
(b) John ate before Mary arrived.

(10) (a) Remember the game we were watching last night?
(b) Well, the Mets won it.
(c) Well, the Mets won.

Examples of this kind illustrate *diathesis alternations*, so called because they concern the ability of verbs to alternate between different ways of expressing their arguments. In the *unspecified object* alternation (Levin, 1993), illustrated in example (9), the omitted object must receive an indefinite or existential interpretation. In the *specified object* alternation (Cote, 1992), illustrated in example (10), an omitted object must refer back to something already under discussion. I will group both these alternations together under the heading of *implicit object* alternations.[14]

Diathesis alternations are a central focus in research on the representation and acquisition of the lexicon (Levin, 1993; Pinker, 1989) because the syntactic alternations in which a verb participates often reflect some aspect of its semantics, and to a great extent the acquisition of a verb consists in relating its surface syntactic behavior to its underlying semantic representation. Unlike many alternations, however, the implicit object alternations appear to involve not only the verb itself, but also the relationship between the verb and its argument. In particular, it has been suggested that verbs allowing implicit objects satisfy aspectual criteria, together with a requirement that the omitted object be typical of the verb or that its properties be in some sense inferable (Lehrer, 1970; Browne, 1971; Mittwoch, 1971, 1982; Rice, 1988; Fellbaum and Kegl, 1989; Levin, 1993; Brisson, 1994). The aspectual criteria have been discussed formally within the framework of Vendler's (Vendler, 1967) aspectual classes, and for the most part consist in the requirement that verbs participating in this alternation denote activities when used intransitively (though see Mittwoch, 1971, Mittwoch, 1982 for fuller discussion). The requirement of inferability or typicality, however, has been expressed only informally until now.

The model of selectional constraints proposed here provides an opportunity to explore the relationship between implicit objects and the notions of inferability or typicality, in a more formal setting. As discussed earlier, the selectional preference strength of a verb for an argument – defined in terms of relative entropy – can be viewed quite literally as the quantity of *information* that the verb carries about that argument. Given that characterization, one should expect that, given the verb, the more strongly an argument is selected for, the easier it should be to infer, or equivalently, the more "typical" that argument should seem. This, in turn, leads to the following prediction: the ability of a verb to participate in implicit object alternations is largely predicted by its selectional preference strength.

Ideally, this prediction should not be tested on its own. As noted, the aspectual class of a verb plays an important role in determining whether or not it participates in the indefinite object alternation (though the influence of aspect on the specified object alternation is less well studied); thus aspectual class should be one of the variables in this experiment. Unfortunately, aspectual class is a difficult concept to pin down, and getting reliable judgments concerning the aspectual class of a verb is quite difficult. As a result, therefore, the relationship between selectional preference strength and implicit object alternations was investigated without explicitly taking aspectual class into account.

[14] Also see Fillmore's (Fillmore, 1986) discussion of "indefinite null complements" and "definite null complements." There are, of course, other reasons an English verb might appear without its object. For example, objects can be omitted in descriptions of habitual or characteristic activities – "Pussycats eat, but tigers devour!" This study concerned only "lexically conditioned" omission of direct objects; see discussion by Fellbaum and Kegl (1989) and Resnik (1993).

EXPERIMENT 3

The predicted connection between selectional preference strength and implicit objects was tested by means of a computational experiment.

13. Materials

The set of verbs investigated is given in Table 4; verbs selected for the set were chosen because they occurred frequently in a collection of parent–child interactions (Anne Lederer, personal communication). Before computing selectional preference strength, the verbs were manually classified according to linguistic diagnostics (Cote, 1992; Resnik, 1993) into two groups, designated Alternating (those verbs that participate in one of the implicit object alternations) and Non-Alternating.

14. Method

Selectional preference strength was computed separately for the three corpora, resulting in the values given earlier in Table 1.

15. Results and discussion

The three corpora can be treated as yielding three different replications of the same experiment. In all of them, verbs participating in one of the implicit object alternations have a significantly higher strength of selectional preference for the direct object than verbs that do not. For the model constructed using the Brown corpus data, the means for the Alternating and Non-Alternating groups were respectively 2.97 and 1.73 (Mann–Whitney $U=55$, $p=.001$). For the model using data from the CHILDES collection, the means were respectively 2.25 and 1.13 (Mann–Whitney $U=37$, $p<.0005$). Finally, for the model using the human subject norms the means were respectively 2.17 and 1.66 (Mann–Whitney $U=57$, $p<.0025$).

These results confirm the hypothesis that verbs participating in implicit object

Table 4
Classification of verbs with respect to implicit object alternations

Group	Verbs
ALTERNATING	call drink eat explain hear pack play pour pull push read sing steal watch write
NON-ALTERNATING	bring catch do find get give hang have hit like make open put say see show take want wear

alternations select more strongly for the direct objects than verbs that do not. Replications using different corpora confirm that the difference is not the result of quirky statistical behavior in a particular corpus. There is, however, no clear threshold separating the two groups of verbs. For example, using the Brown Corpus data, the three "weakest" alternating verbs are *call*, *hear*, and *watch*, with selectional preference strengths ranging from 1.52 to 1.97. Likewise, the three "strongest" non-alternating verbs are *hang*, *wear*, and *open*, with selectional preference strengths ranging from 2.93 to 3.35.

EXPERIMENT 4

Experiment 3 confirmed the hypothesis that optionality of the direct object is connected to selectional preference, the rationale being that strength of selectional preference is, as formalized here, a measure of how easy it is to infer or reconstruct necessary properties of the omitted object. Given that selection is relevant to lexical-syntactic properties – that is, lexical knowledge bearing on syntactic competence – a natural question to ask is whether selectional preference affects performance, as well. If selectional preference strength measures how much information a verb carries about its object, then on-line inferences about omitted objects should be easier for verbs that select strongly rather than weakly, and this should be reflected in on-line syntactic behavior.

Ease of inference is a subject for investigation by psycholinguistic rather than computational methods. However, in performance, a speaker or writer is likely to be influenced by how easy it will be for the listener or reader to arrive at the correct interpretation. In particular, one would expect that verbs for which the object is readily inferable will omit that argument correspondingly more frequently than verbs for which the object is not easily inferred. In another experiment, therefore, I have again identified ease of inference with strength of selectional preference, this time predicting a correlation between selectional preference and the omission of direct objects in naturally occurring text.

16. Materials

In order to determine the frequency with which verbs omit their objects, I extracted from the Brown Corpus a random sample of 100 instances of each verb used in the preceding experiment (or as many instances as were available, if fewer than 100). For each instance. I used the full sentence in which the verb appeared, together with the full preceding sentence, to decide whether or not this instance was an example of an implicit object construction. The judgments were made using the same linguistic diagnostics as in Experiment 3.[15]

[15] For practical reasons the verb *have* was excluded from this procedure.

17. Method

This experiment used the same values of selectional preference strength computed in Experiment 3.

18. Results and discussion

A correlation between selectional preference strength and percentage of implicit objects emerged in each of the three versions of the experiment: Brown Corpus ($N=33$, $r=0.48$, $F(1, 31)=9.53$, $p(F)<.005$), CHILDES ($N=32$, $r=0.36$, $F(1, 30)=4.33$, $p(F)<.05$), human subject norms ($N=33$, $r=0.58$, $F(1, 31)=15.74$, $p(F)<.0005$). Table 5 shows the percentage of objects omitted for each verb together with the selectional preference strength calculated in each of the three experiments (repeated from Table 1).

As the table shows, some verbs deviate by failing to omit their objects despite very strong selection for the direct object (e.g., *wear, say, pour*). In many cases these verbs fail to meet the aspectual requirements for participation in the indefinite object alternation, namely that the event denoted by the verb be interpretable as an activity (in the sense of Vendler, 1967) if the verb appears without an object. It is important to notice, however, that deviations of the opposite kind do not appear to occur: verbs do not omit their objects frequently unless they possess a high selectional preference strength. I would argue that this

Table 5
Strength of selectional preference and frequency of implicit direct objects

Verb	% Implicit	Strength			Verb	%Implicit	Strength		
		Brown	CHILDES	Norms			Brown	CHILDES	Norms
drink	45.1	4.38	2.38	2.83	do	0.0	1.84	–	2.21
sing	38.3	3.58	3.15	2.63	find	0.0	0.96	0.71	1.30
eat	31.8	3.51	1.15	2.47	get	0.0	0.82	0.28	1.17
write	25.2	2.54	2.33	2.18	give	0.0	0.79	1.18	1.81
play	19.6	2.51	2.13	2.64	hang	0.9	3.35	2.03	1.96
read	12.7	2.35	2.58	1.81	have	–	0.43	–	1.23
hit	9.2	2.49	1.31	1.91	like	0.0	2.59	0.89	1.30
call	7.3	1.52	0.95	2.39	make	0.0	0.72	0.77	1.58
steal	7.9	3.52	2.28	1.34	pour	0.0	4.80	2.30	2.57
pack	4.9	4.12	3.71	1.75	put	0.0	1.24	0.40	1.34
open	3.7	2.93	2.41	1.88	say	0.0	2.82	0.94	2.56
explain	2.7	2.39	4.41	2.20	see	0.9	1.06	0.48	1.54
hear	2.8	1.70	1.67	1.71	show	0.0	1.39	1.83	1.42
catch	1.8	2.47	1.67	1.92	take	0.0	0.93	0.74	1.28
pull	1.9	2.77	1.55	2.22	want	0.0	1.52	0.70	1.71
push	1.0	2.87	1.77	1.98	watch	0.0	1.97	1.44	1.86
bring	0.0	1.33	0.88	1.04	wear	0.0	3.13	2.02	2.30

pattern reflects an underlying hard requirement, namely that strong selection is a necessary condition for object omission.

Although previous work argues convincingly that the syntactic realization of arguments for this group of verbs is connected with aspectual category,[16] this study goes beyond the assumption that the syntactic behavior of these verbs arises *only* by virtue of their semantic structure, pursuing the possibility that in this case argument realization is also a function of semantic content (where semantic structure and semantic content are interpreted in the sense of Grimshaw, 1993, as discussed earlier). Experiment 3 has as its starting point the intuition, previously stated informally, that unexpressed objects tend to be restricted to "typical" or "probable" arguments of the verb (Rice, 1988; Levin, 1993; Brisson, 1994), together with the familiar idea that the inferred semantic properties of the omitted object constitute information carried by the verb in the form of a selectional constraint (e.g., Jackendoff, 1990). Taken together, these suggest that in order for a verb to participate in these alternations, the selectional constraint it imposes on its object must be relatively specific. In the present model, selectional constraints are formalized in terms of a measure of informativeness, which leads to the testable prediction that participation in the alternation should be connected with high information content. This prediction is borne out by the significant difference in selectional preference strength between the alternating and non-alternating verbs.

Experiment 4 goes further by considering the performance aspects of an alternation. Previous work on diathesis alternations understandably has had little to say about this, focusing on semantic representations; in contrast, the present model makes quantitative predictions, thus making it possible to ask whether some quantitative aspects of lexical representation might be reflected in such quantifiable aspects of performance as on-line processing, frequency, and the like. The two perspectives are rather complementary: a great body of work on diathesis alternations supports the view that semantic structure imposes categorical constraints on the way verbs *can* realize their arguments; the correlations in Experiment 4 support the view that semantic content provides quantitative constraints on the way verbs *do* realize their arguments.

GENERAL DISCUSSION

19. Properties of the model

The model attempts to address many of the issues raised in the literature on selectional constraints. Returning to the criteria outlined earlier:

1. Representational assumptions about word meaning are kept to a minimum. In

[16] And perhaps other elements of verb meaning, as well: some subclasses of activity verbs that participate in these alternations appear to be semantically coherent (e.g., *dust, iron, sweep, crochet, sew, knit, draw, sing, paint*).

particular, it is not necessary to make the assumption, now widely viewed as untenable, that a word meaning can he expressed definitionally, as a set of necessary and sufficient Boolean conditions. Nor do the predictions of the model require general inferential mechanisms. Instead, the model imposes a minimal requirement about the representation of arguments, namely that it be possible to organize word meanings into discrete classes. More elaborated representations, notably those based on the notion of semantic decomposition, are allowed but not required.

2. The formal model permits, and the computational implementation realizes, selectional constraints that can refer to arbitrarily specific properties rather than a small set of semantic primitives. Although properties are not an explicit part of the knowledge representation, they are expressed implicitly in the way the taxonomy organizes word meanings. By using an on-line lexicon, WordNet, that attempts to capture "all the concepts that are lexicalized in English" (George Miller, personal communication), the computational implementation of the model reflects in a fairly direct way McCawley's (McCawley, 1968) observation that "on any page of a large dictionary one finds words with incredibly specific selectional restrictions."
3. A notion akin to semantic anomaly arises within the model as a matter of degree, in the form of a quantitative measure of selectional association (cf. Wilks, 1975; Wilks and Fass, 1992). Traditional examples (green ideas, sincerity sleeping) still fall out naturally as selectional "violations" because when all the classes to which the argument belongs are considered, no class has a positive selectional association with the predicate.

The experiments presented here provide empirical support for this model as an approximation of adult knowledge of selectional constraints. Experiments 1 and 2 make a direct comparison between "judgments" of selectional fit, as computed by the model, and ratings of argument plausibility or typicality elicited from adult subjects. Although the adult ratings are influenced by factors other than selectional constraints per se, they represent the closest experimental approximation to selectional fit that one is likely to find readily available. Experiments 3 and 4 establish a connection between selectional constraints, as formalized in the model, and the syntactic realization of arguments for a class of verbs in English. Although the behavior of these verbs has been linked informally to some notion of inferability or typicality in the past, the interpretation of selectional preference strength as a quantity of information makes a new, formal account possible. The empirical success of this account lends further credibility to the underlying model, though the interaction between aspectual class and argument inferability for these verbs clearly requires further study.

20. Learning

A principal reason for adopting the strategy in this paper – seeking a model of selectional constraints that does not assume much prior knowledge of verbs' semantic representations – is that selectional constraints may themselves have a role in learning which verb forms go with which representations (the *mapping problem*: Fisher et al., 1994; Pinker, 1994). In one relevant line of research, Gleitman and colleagues explore the information potentially provided by different sources of evidence about verb meaning (Landau and Gleitman, 1985; Gleitman et al., 1988; Gleitman, 1990; Fisher et al., 1991; Lederer et al., 1991). Gleitman and Gillette (1995) show that selectional constraints provide adult subjects with significant constraints on the possible meanings of unknown verbs, well over and above the constraints provided by observation of visual context, the syntactic frame in which the verb is used, or even knowledge of the participants in an event without their specific relationship to the verb. The relevant manipulation in their experiments is one in which subjects were given lists of the nouns appearing in a mother's utterances to her child, alphabetized to obscure the predicate–argument structure and show only noun co-occurrences within each sentence. Asked to guess the verb, subjects in this condition guessed correctly only 13% of the time. Given the syntactic frames from the same utterances (with target verb and all nouns converted to corresponding nonsense words, as in *Rom GORPS that the rivenflak is grum*), subjects identified 52%, of the verbs correctly. But given the syntactic frames together with the nouns in their correct positions, subjects identified the correct verb 80% of the time. Gillette and Gleitman write:

> It doesn't much help to know that one of the words in an utterance was *hamburger*. But if this word is known to surface as direct object, the meaning of the verb might well be *eat* ... the structural information converts co-occurrence information to selectional information.

Correspondingly, Pinker (1994) writes:

> [Children] surely can infer much about what a verb means from the meanings of the other words in the sentence and from however much of the sentence's structure they are able to parse. If someone were to hear *I filped the delicious sandwich and now I'm full*, presumably he or she could figure out that ... *filp* means something like "eat". But ... no thanks are due to the verbs' syntactic frames (in this case, transitive). Rather, we know what those verbs mean because of the semantics of ... *sandwich, delicious, full*, and the partial syntactic analysis that links them together.

Pinker is alluding primarily to real-world contingencies, such as the fact that "hearing a verb used with sandwich suggests that it involves eating." However, the partial syntactic analysis he refers to provides critical information not available from the semantics of *sandwich* alone, nor any unordered combination of the known participants in the event. Notice that, given the same set of real-world contingencies but the sentence *The delicious sandwich filped me and now I'm full*, the meaning EAT would be an implausible hypothesis for the meaning of *filp*, because *sandwich* is an unlikely subject/agent, but a meaning such as SATISFY would be perfectly plausible. It is, therefore, neither the presence of *sandwich* as a participant in the event nor the syntactic structure alone that provides the necessary constraint, but rather the semantic content of *sandwich* as a particular argument of the verb. Pinker's theory of verb acquisition contrasts notably with Gleitman's in its emphasis on observation of real-world contingencies as opposed to the influence of grammatical context. However, knowledge about the semantic content of arguments – that is, selectional constraints – appears to have a role in both approaches, providing the learner with constraints on hypotheses about verb meaning.

Gropen (1993) proposes a specific mechanism for the use of selectional constraints, extending Pinker's (Pinker, 1989) learning theory in order to address the problem of verb polysemy. The core of the theory is a process of "semantic structure hypothesis testing" across observations of a verb being used in context. Presented with a new word, the child initially hypothesizes a meaning that may be overly specific (e.g., guessing that *open* means something like "act on a door in a pulling manner"); subsequent observations of the word are then used to refine the hypothesis, and the child retains only those elements of meaning that are consistent across situations of use (e.g, seeing a door opened by pushing it, the child expunges the "pulling" requirement from the verb's meaning). Gropen points out that verb polysemy complicates this basic account: if multiple senses of a verb are not individuated, this refinement of the hypothesis can blindly rule out elements of meaning that are crucial for one sense but not another. He suggests that children solve this problem by being sensitive to the correlation between verb meanings and the types of the participants in the observed situation labeled by the verb stem. Inconsistencies in the semantic properties of the verb's arguments across observed situations provide evidence that the learner should be considering distinct semantic structures, corresponding to different senses of the verb.[17]

Grimshaw (1994) attempts to reconcile the roles of syntactic and situational context in lexical learning, proposing a learning procedure in which selectional constraints play a key role. On Grimshaw's model, the learner (1) uses the observed situation to posit a sensible semantic relationship among the participants in that situation, then (2) constructs a lexical conceptual structure consistent with that relationship, and finally (3) uses general principles mapping from semantic form to syntactic structure in order to check whether the proposed conceptual structure is consistent with the surface syntactic structure of the observed utterance. Before moving from the first step to the second, however, the learning procedure verifies that the participants in the hypothesized semantic relationship are consistent with the candidate argument expressions in the sentence. It is this consistency check that requires the application of selectional constraints – as

[17] Siskind (1994) discusses an alternative approach to individuating word senses, using a notion of inconsistency more general than that of distinct participant types.

Grimshaw points out, "if the sentence contains two NPs, and one is *the ball*, then a verb meaning 'say' is not a candidate, since it is not a possible relationship between a ball and another entity" (p. 423). As another example, consider a situation in which I am observed making music by doing something to a guitar with my hands, and the noun phrases in the sentence are *Daddy* and *the guitar*. Although there are many possible construals of the observed situation that involve me and the guitar – for example, what my right hand is doing might look a lot like tickling – TICKLE is a poor candidate for the semantic relationship between the participants in the utterance because things that get tickled are generally animate.[18] Given multiple interpretations of the same observed event, then, selectional constraints play an early role in ruling out possible interpretations of the observed scene, by eliminating conceptual interpretations not consistent with the arguments in the sentence that was uttered.

All these theories of lexical acquisition make use of selectional constraints, and therefore each must ultimately provide an explicit account of how they arise. On the most commonly held view, selectional constraints emerge from some combination of lexical representation, factual knowledge, and inference about consistent patterns in the real world; however, the process by which this takes place has never been made explicit. Acquisition of selectional constraints in the model proposed here can be viewed as a first attempt to make such a process explicit, drastically simplified by reducing general factual knowledge to knowledge of taxonomic relationships, and by considering only patterns of linguistic behavior rather than patterns of more general real-world observation.

Viewing the model in this way requires two assumptions about child language learners. The first is that their representations of arguments are organized into conceptual classes. One cannot assume that children's organizational criteria match the adult representations realized by the implementation, but the formal model would seem to be stating the very minimum that one must assume in order to say anything at all about the semantic or conceptual properties of arguments. The second assumption is the availability of co-occurrence frequencies for predicates and arguments. This is a harder assumption to make, in the absence of further assumptions about knowledge of verb semantics; however, at least for surface syntactic relationships, there is evidence suggesting that children may be able to construct a skeletal parse on the basis of known words and prosodic information (e.g., Kemler Nelson et al., 1989; Lederer and Kelly, 1991; also see discussion in Gleitman et al., 1988, and Fisher et al., 1994). In the experiments that were conducted here using parental utterances from the CHILDES collection, direct objects were identified by a strategy not much more sophisticated than "find the nearest noun to the right of the verb," so the input provided to the learning

[18] As it happens, given the learning sample from CHILDES, the current model proposes ⟨relative⟩ as the most strongly associated class of direct objects for tickle, based on frequent observation of objects like *brother, daddy, and baby*.

algorithm was no richer than the "partial sentence representations" discussed by Fisher et al. (1994).[19]

To the extent that its assumptions are plausible, the information-theoretic model represents a starting point for modeling the role of selectional constraints in a more comprehensive theory of verb acquisition. More generally, the approach taken here has methodological implications: it provides an example of how computational methods and large-scale collections of naturalistic data can be used to explore issues that traditionally fall within the realms of linguistics and experimental psychology. Combining the resources of all these disciplines may be the key to a formal theory of language acquisition that takes into account the noisy, complex data faced by the language learner.

Acknowledgments

Much of this work was conducted as part of the author's doctoral dissertation at the University of Pennsylvania. The author gratefully acknowledges the discussions and criticism of Steve Abney, Michael Brent, Henry Gleitman, Lila Gleitman, Ellen Hays, Marti Hearst, Aravind Joshi, Bob Kuhns, Mark Liberman, Mitch Marcus, Jeff Siskind, Jim Waldo, and three anonymous reviewers, as well as the financial support of grants ARO DAAL 03-89-C-0031, DARPA N00014-90-J-1863, and Ben Franklin 91S.3078C-1, and an IBM Graduate Fellowship.

References

Allen, J. (1995). *Natural language understanding.* Redwood City, CA: Benjamin/Cummings Publishing Company.

Alshawi, H., & Carter, D. (1992). Sortal restrictions. In H. Alshawi (Ed.), *The core language engine.* Cambridge, MA: MIT Press.

Armstrong, S., Gleitman, L., & Gleitman, H. (1983). What some concepts might not be. *Cognition, 13,* 263–308.

[19] Jill Fain has pointed out that the algorithm as implemented makes assumptions that are unrealistic from a developmental point of view: some mechanism would have to be provided to enable incremental updating of probability estimates on the basis of new data, retention over time of the information that would require, and subsequent re-computation of selectional constraints. In addition, an anonymous reviewer suggests that a learner may be able to predict some selectional behavior of a new verb by transfer from a known verb with a related meaning, rather than relying on verb–argument co-occurrences involving the new verb. For example, the learner might at some point discover that the unknown verb *slurp* is related in meaning to the known verb *drink*, and therefore "transfer" the selectional constraints of *drink* to *slurp* along with other components of its lexical representation. (At which point, presumably the selectional constraints as well as other elements of meaning could be further modified by observation of how the verb is used in context.) Both of these are valid points; however, the model has not yet been extended to address either of these issues.

Brisson, C. (1994). The licensing of unexpressed objects in English verbs. In *Proceedings of the Chicago Linguistic Society's 30th Regional Meeting with a Parasession on Variation and Linguistic Theory*. Chicago: Chicago Linguistic Society.

Browne, W. (1971). Verbs and unspecified NP deletion. *Linguistic Inquiry, 2*, 259–260,

Cote, S. (1992). *Discourse functions of two types of null objects in English*. Presented at the 66th Annual Meeting of the Linguistic Society of America, Philadelphia, PA.

Cover, T.M., & Thomas, J.A. (1991). *Elements of information theory*. New York: Wiley.

Dowty, D. (1991). Thematic proto-roles and argument selection. *Language, 67*, 547–615.

Drange, T. (1966). *Type crossings: Sentential meaninglessness in the border area of linguistics and philosophy*. The Hague: Mouton.

Fellbaum, C., & Kegl, J. (1989). Taxonomic structures and cross-category linking in the lexicon. In *Proceedings of the Eastern States Conference on Linguistics, 28–30 September 1989* (pp. 93–104). Columbus: Ohio State University.

Fillmore, C. (1986). Pragmatically controlled zero anaphora. In *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* (pp. 95–107). Berkeley: Berkeley Linguistics Society.

Fisher, G., Gleitman, L., & Gleitman, H. (1991). On the semantic content of subcategorization frames. *Cognitive Psychology, 23*, 331–392.

Fisher, C., Hall, G., Rakowitz, S., & Gleitman, L. (1994). When it is better to receive than to give: syntactic and conceptual constraints on vocabulary growth. *Lingua, 92*, 333–376.

Fodor, J.A., Garrett, M.F., Walker, E.T., & Parkes, C. (1980). Against definitions. *Cognition, 8*, 1–105.

Fodor, J.D. (1977). *Semantics: Theories of meaning in generative grammar*. Cambridge, MA: Harvard University Press.

Francis, W., & Kučera, H. (1982). *Frequency analysis of English usage*. New York: Houghton Mifflin.

Gentner, D. (1982). Why nouns are learned before verbs: Linguistic relativity vs. natural partitioning. In S.A. Kuczaj (Ed.), *Language development: Language, culture, and cognition* (pp. 301–334). Hillsdale, NJ: Erlbaum.

Gleitman, L. (1990). The structural sources of verb meanings. *Language Acquisition, 1*, 3–55.

Gleitman, L., & Gillette, J. (1995). The role of syntax in verb learning. In P. Fletcher and B. MacWhinney (Eds.), *Handbook of child language*. Oxford: Blackwell.

Gleitman, L., Gleitman, H., Landau, B., & Wanner, E. (1988). Where learning begins: Initial representations for language learning. In F.J. Newmeyer (Ed.), *Linguistics: The Cambridge survey. Vol. III. Language: Psychological and biological aspects*. New York: Cambridge University Press.

Grimshaw, J. (1993). *Semantic structure and semantic content in lexical representation*. Unpublished manuscript, Center for Cognitive Science, Rutgers University, New Brunswick, NJ.

Grimshaw, J. (1994). Lexical Reconciliation. In L. Gleitman & B. Landau (Eds.), *The acquisition of the lexicon*. Cambridge, MA: MIT Press.

Gropen, J. (1993). *Participant types and the acquisition of verb polysemy*. Unpublished manuscript.

Hobbs, J., Stickel, M., Appelt, D., & Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence, 63*, 69–142.

Holmes, V.M., Stowe, L., & Cupples, L. (1989). Lexical expectations in parsing complement–verb sentences. *Journal of Memory and Language, 28*, 668–689.

Horn, L. (1989). *A natural history of negation*. Chicago: University of Chicago Press.

Jackendoff, R. (1983). *Semantics and cognition*. Cambridge, MA: MIT Press.

Jackendoff, R. (1990). *Semantic structures*. Cambridge, MA: MIT Press.

Johnson-Laird, P.N. (1983). *Mental models: Towards a cognitive science of language, inference, and consciousness*. Cambridge, MA: Harvard University Press.

Katz, J.J. (1972). *Semantic theory*. New York: Harper.

Katz, J.J. & Fodor, J.A. (1964). The structure of a semantic theory. In J.A. Fodor & J.J. Katz (Eds.), *The structure of language* (pp. 479–518). Englewood Cliffs, NJ: Prentice-Hall.

Kemler Nelson, D., Hirsh-Pasek, K., Jusezyk, R.L., & Cassidy, K. (1989). How the prosodic cues in motherese might assist language learning. *Journal of Child Language, 16*, 55–68.

Kullback, S., & Leibler, R.A. (1951). On information and sufficiency. *Annals of Mathematical Statistics, 22*, 79–86.

Landau, B., & Gleitman, L. (1985). *Language and experience*. Cambridge, MA: Harvard University Press.

Lederer, A., Gleitman, H., Gleitman, L., & Ariel, S. (1991). *An empirical investigation of the information sources for language acquisition: Cross-situational vs. cross-sentential observation*. Paper presented at the 16th Boston University Conference on Language Development.

Lederer, A., & Kelly, M. (1991). Prosodic correlates to the adjunct/complement distinction in Motherese. *Papers and Reports in Child Language Development, 30*.

Lehmann, F. (Ed.) (1992). *Semantic networks in artificial intelligence*. Oxford: Pergamon Press.

Lehrer, A. (1970). Verbs and deletable objects. *Lingua, 25*, 227–254.

Levin, B. (1993). *English verb classes and alternations: A preliminary investigation*. Chicago: University of Chicago Press.

MacDonald, M. (1993). The interaction of lexical and syntactic ambiguity. *Journal of Memory and Language, 32*, 692–715.

MacWhinney, B., & Snow, C. (1985). The child language data exchange system. *Journal of Child Language, 12*.

Marcus, M.P., Santorini, B., & Marcinkiewicz, M. (1993). Building a large annotated corpus of English: The Penn Treebank. *Computational Linguistics, 19*, 313–330.

McCawley, J. (1968). The role of semantics in a grammar. In E. Bach & R. Harms (Eds.), *Universals in linguistic theory* (pp. 124–169). New York: Holt, Rinehart and Winston.

McCawley, J.D. (1971). Interpretative semantics meets Frankenstein. *Foundations of Language, 7*, 285–296.

Miller, G. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography, 3*(4) (Special Issue).

Mittwoch, A. (1971). Idioms and Unspecified NP Deletion. *Linguistic Inquiry, 2*, 255–259.

Mittwoch, A. (1982). On the difference between eating and eating something: activities versus accomplishments. *Linguistic Inquiry, 13*, 113–122.

Nelson, K. (1973). Structure and strategy in learning to talk. *Monographs of the Society for Research in Child Development, 38*(1–2), Serial no. 149.

Pinker, S. (1989). *Learnability and cognition*. Cambridge, MA: MIT Press.

Pinker, S. (1994). How could a child use verb syntax to learn verb semantics?. *Lingua, 92*, 377–410.

Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.

Pustejovsky, J., Bergler, S., & Anick, P. (1993). Lexical semantic techniques for corpus analysis. *Computational Linguistics, 19*, 291–358.

Resnik, P. (1993). *Selection and information: A class-based approach to lexical relationships*. Ph.D. thesis, University of Pennsylvania.

Rice, S. (1988). Unlikely lexical entries. In *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* (pp. 202–212). Berkeley: Berkeley Linguistics Society.

Rosch, E., Mervis, C., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.

Siskind, J.M. (1994). Lexical acquisition in the presence of noise and homonymy. In *Proceedings of the National Conference on Artificial Intelligence, AAAI-94*. Cambridge: American Association for Artificial Intelligence/Cambridge: MIT Press.

Smyth, R., & Goodman, R. (1992). An information theoretic approach to rule induction from databases. *IEEE Transactions on Knowledge and Data Engineering, 4*, 301–316.

Sowa, J.F. (Ed.) (1991). *Principles of semantic networks*. Sam Mateo, CA: Morgan Kaufmann

Tabossi, P, Spivey-Knowlton, M.J., McRae, K., & Tanenhaus, M.K. (1994). Semantic effects on syntactic ambiguity resolution: Evidence for a constraint-based resolution process. In C. Umilta & M. Moscovitch (Eds.), *Attention and Performance* (vol. XV, pp.598–615). Hillside, NJ: Erlbaum.

Trueswell, J., Tanenhaus, M., & Garnsey, S. (1994). Semantic influences on parsing: Use of thematic role information in syntactic ambiguity resolution. *Journal of Memory and Language, 33*, 285–318.

Vendler, Z. (1967). *Linguistics in philosophy* (pp. 97–121). Ithaca: Cornell University Press.

Wilks, Y. (1975). An intelligent analyzer and understander of English. *Communications of the ACM, 18*, 264–274.

Wilks, Y., & Fass, D. (1992). The preference semantics family. *Computers and Mathematics with Applications, 23*, 205–221.

Wilks, Y.A., Slator, B.M., & Guthrie, L.M. (1996). *Electric words: Dictionaries, computers, and meanings*. Cambridge, MA: MIT Press.

Yarowsky, D. (1992). Word-sense disambiguation using statistical models of Roget's categories trained on large corpora. In *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92)* (pp. 454–460). Nantes, France.