

# Role of Word Sense Disambiguation in Lexical Acquisition: Predicting Semantics from Syntactic Cues

Bonnie J. Dorr and Doug Jones  
Department of Computer Science and  
Institute for Advanced Computer Studies  
University of Maryland  
A.V. Williams Building  
College Park, MD 20742  
{bonnie,jones}@umiacs.umd.edu

## Abstract

This paper addresses the issue of word-sense ambiguity in extraction from machine-readable resources for the construction of large-scale knowledge sources. We describe two experiments: one which ignored word-sense distinctions, resulting in 6.3% accuracy for semantic classification of verbs based on (Levin, 1993); and one which exploited word-sense distinctions, resulting in 97.9% accuracy. These experiments were dual purpose: (1) to validate the central thesis of the work of (Levin, 1993), i.e., that verb semantics and syntactic behavior are predictably related; (2) to demonstrate that a 15-fold improvement can be achieved in deriving semantic information from syntactic cues if we first divide the syntactic cues into distinct groupings that correlate with different word senses. Finally, we show that we can provide effective acquisition techniques for novel word senses using a combination of online sources.

## 1 Introduction

This paper addresses the issue of word-sense ambiguity in extraction from machine-readable resources for the construction of large-scale knowledge sources. We describe two experiments: one which ignored word-sense distinctions, resulting in 6.3% accuracy for semantic classification of verbs based on (Levin, 1993); and one which exploited word-sense distinctions, resulting in 97.9% accuracy. These experiments were dual purpose: (1) to validate the central thesis of the work of (Levin, 1993), i.e., that verb semantics and syntactic behavior are predictably related; (2) to demonstrate that a 15-fold improvement can be achieved in deriving semantic information from syntactic cues if we first divide the syntactic cues into distinct groupings that correlate with different word senses. Finally, we show that we can provide effective acquisition techniques for novel word senses using a combination of online sources, in particular, Longman's Dictionary of Contemporary English (LDOCE) (Procter, 1978), Levin's verb classification scheme (Levin, 1993), and WordNet (Miller, 1985). We have used these techniques to build a database of 10,000 English verb entries containing semantic information that we are currently porting into languages such as Arabic, Spanish, and Korean for multilingual NLP tasks such as foreign language tutoring and machine translation.

## 2 Automatic Lexical Acquisition for NLP Tasks

As machine-readable resources (i.e., online dictionaries, thesauri, and other knowledge sources) become readily available to NLP researchers, automated acquisition has become increasingly more attractive. Several researchers have noted that the average time needed to construct a lexical entry can be as much as 30 minutes (see, e.g., (Neff and McCord, 1990; Copestake et al., 1995; Walker and Amsler, 1986)). Given that we are aiming for large-scale lexicons of 20-60,000 words, automation of the acquisition process has become a necessity.

Previous research in automatic acquisition focuses primarily on the use of statistical techniques, such as bilingual alignment (Church and Hanks, 1990; Klavans and Tzoukermann, 1996; Wu and Xia, 1995), or extraction of syntactic constructions from online dictionaries and corpora (Brent, 1993; Dorr, Garman, and Weinberg, 1995). Others who have taken a more knowledge-based (interlingual) approach (Lonsdale, Mitamura, and Nyberg, 1996) do not provide a means for systematically deriving the relation between surface syntactic structures and their underlying semantic representations. Those who have taken more argument structures into account, e.g., (Copestake et al., 1995), do not take full advantage of the systematic relation between syntax and semantics during lexical acquisition.

We adopt the central thesis of Levin (1993), i.e., that the semantic class of a verb and its syntactic behavior are predictably related. We base our work on a correlation between semantic classes and patterns of grammar codes in the Longman's Dictionary of Contemporary English (LDOCE) (Procter, 1978). While the LDOCE has been used previously in automatic extraction tasks (Alshawi, 1989; Farwell, Guthrie, and Wilks, 1993; Boguraev and Briscoe, 1989; Wilks et al., 1989; Wilks et al., 1990) these tasks are primarily concerned with the extraction of other types of information including syntactic phrase structure and broad argument restrictions or with the derivation of semantic structures from definition analyses. The work of Sanfilippo and Poznanski (1992) is more closely related to our approach in that it attempts to recover a syntactic-semantic relation from machine-readable dictionaries. However, they claim that the semantic classification of

verbs based on standard machine-readable dictionaries (e.g., the LDOCE) is “a hopeless pursuit [since] standard dictionaries are simply not equipped to offer this kind of information with consistency and exhaustiveness.”

Others have also argued that the task of simplifying lexical entries on the basis of broad semantic class membership is complex and, perhaps, infeasible (see, e.g., Boguraev and Briscoe (1989)). However, a number of researchers (Fillmore, 1968; Grimshaw, 1990; Gruber, 1965; Guthrie et al., 1991; Hearst, 1991; Jackendoff, 1983; Jackendoff, 1990; Levin, 1993; Pinker, 1989; Yarowsky, 1992) have demonstrated conclusively that there is a clear relationship between syntactic context and word senses; it is our aim to exploit this relationship for the acquisition of semantic lexicons.

### 3 Syntax-Semantics Relation: Verb Classification Based on Syntactic Behavior

The central thesis of (Levin, 1993) is that the semantics of a verb and its syntactic behavior are predictably related. As a demonstration that such predictable relationships are not confined to an insignificant portion of the vocabulary, Levin surveys 4183 verbs, grouped into 191 semantic classes in Part Two of her book. The syntactic behavior of these classes is illustrated with 1668 example sentences, an average of 8 sentences per class.

Given the scope of Levin’s work, it is not easy to verify the central thesis. To this end, we created a database of Levin’s verb classes and example sentences from each class, and wrote a parser to extract basic syntactic patterns from the sentences.<sup>1</sup> We then characterized each semantic class by a set of syntactic patterns, which we call a *syntactic signature*, and used the resulting database as the basis of two experiments, both designed to discover whether the syntactic signatures tell us anything about the meaning of the verbs.<sup>2</sup> The first experiment, which we label *Verb-Based*, ignores word-sense distinctions by assigning one syntactic signature to each verb, regardless of whether it occurred in multiple classes. The second experiment, which we label *Class-Based*, implicitly takes word-sense distinctions into account by considering each occurrence of a verb individually and assigning it a single syntactic signature according to class membership.

The remainder of this section describes the assignment of signatures to semantic classes and the two experiments for determining the relation of syntactic information to semantic classes. We will see that our classification technique shows a 15-fold improvement in the experiment where we implicitly account for word-sense distinctions.

<sup>1</sup>Both the database and the parser are encoded in Quintus Prolog.

<sup>2</sup>The design of this experiment is inspired by the work of (Dubois and Saint-Dizier, 1995). In particular, we depart from the alternation-based data in (Levin, 1993), which is primarily binary in that sentences are presented in pairs which constitute an alternation. Following Saint-Dizier’s work, we construct N-ary syntactic characterizations. The choice is of no empirical consequence, but it simplifies the experiment by eliminating the problem of naming the syntactic patterns.

**Verbs:** break, chip, crack, crash, crush, fracture, rip, shatter, smash, snap, splinter, split, tear

#### Example Sentences:

Crystal vases break easily.  
 The hammer broke the window.  
 The window broke.  
 Tony broke her arm.  
 Tony broke his finger.  
 Tony broke the crystal vase.  
 Tony broke the cup against the wall.  
 Tony broke the glass to pieces.  
 Tony broke the piggy bank open.  
 Tony broke the window with a hammer.  
 Tony broke the window.  
 \*Tony broke at the window.  
 \*Tony broke herself on the arm.  
 \*Tony broke himself.  
 \*Tony broke the wall with the cup.  
 A break.

#### Derived Syntactic Signature:

1-[np,v] 1-[np,v,np] 1-[np,v,np,adjective]  
 1-[np,v,np,pp(against)] 1-[np,v,np,pp(to)]  
 1-[np,v,np,pp(with)] 1-[np,v,poss,np]  
 1-[np,v,adv(easily)] 1-[n]  
 0-[np,v,np,pp(with)] 0-[np,v,self]  
 0-[np,v,self,pp(on)] 0-[np,v,pp(at)]

Table 1: Syntactic Signature for Change of State – *break* subclass

#### 3.1 Assignment of Signatures

For the first experiment below, we construct a verb-based syntactic signature, while for the second experiment, we constructed a class-based signature.

The first step for constructing a signature is to decide what syntactic information to extract for the basic syntactic patterns that make up the signature. It turns out that a very simple strategy works well, namely, flat parses that contain lists of the major categories in the sentence, the verb, and a handful of other elements. The “parse”, then, for the sentence **Tony broke the crystal vase** is simply the syntactic pattern [np,v,np]. For **Tony broke the vase to pieces** we get [np,v,np,pp(to)]. Note that the pp node is marked with its head preposition. Table 1 shows an example class, the *break* subclass of the Change of State verbs (45.1), along with example sentences and the derived syntactic signature based on sentence patterns. Positive example sentences are denoted by the number 1 in the sentence patterns and negative example sentences are denoted by the number 0 (corresponding to sentences marked with a \*).

#### 3.2 Experiment 1: Verb-based Approach

In the first experiment, we ignored word sense distinctions and considered each verb only once, regardless of whether it occurred in multiple classes. In fact, 46% of the verbs appear more than once. In some cases, the verb appears to have a related sense even though it appears in different classes. For example, the verb *roll* appears in two subclasses of Manner of Motion Verbs that are distinguished on the basis of whether the grammatical subject is animate or inanimate. In other cases, the verb may have (largely) unrelated senses. For example, the verb *move* is both a Manner of Motion verb

and verb of Psychological State.

To compose the syntactic signatures for each verb, we collect all of the syntactic patterns associated with every class a particular verb appears in, regardless of the different classes are semantically related. A syntactic signature for a verb, by definition, is the union of the frames extracted from every example sentence for each verb. The outline of the verb-based experiment is as follows:

1. Automatically extract syntactic information from the example sentences.
2. Group the verbs according to their syntactic signature.
3. Determine where the two ways of grouping verbs overlap:
  - (a) the semantic classification given by Levin.
  - (b) the syntactic classification based on the derived syntactic signatures.

To return to the Change of State verbs, we now consider the syntactic signature of the verb *break*, rather than the signature of the semantic class as a unit. The verb *break* belongs not only to the Change of State class, but also four other classes: 10.6 *Cheat*, 23.2 *Split*, 40.8.3 *Hurt*, and 48.1.1 *Appear*. Each of these classes is characterized syntactically with a set of sentences. The union of the syntactic patterns corresponding to these sentences forms the syntactic signature for the verb. So although the signature for the Change of State class has 13 frames, the verb *break* has 39 frames from the other classes it appears in.

Conceptually, it is helpful to consider the difference between the *intension* of a function versus its *extension*. In this case, we are interested in the functions that group the verbs syntactically and semantically. Intensionally speaking, the definition of the function that groups verbs semantically would have something to do with the actual meaning of the verbs.<sup>3</sup> Likewise, the intension of the function that groups verbs syntactically would be defined in terms of something strictly syntactic, such as subcategorization frames. But the intensions of these functions are matters of significant theoretical investigation, and although much has been accomplished in this area, the question of mapping syntax to semantics and vice versa is an open research topic. Therefore, we can turn to the *extensions* of the functions: the actual groupings of verbs, based on these two separate criteria. The semantic extensions are sets of verb tokens, and likewise, the syntactic extensions are sets of verb tokens. To the extent that these functions map between syntax and semantics intensionally, they will pick out the same verbs extensionally.

So for the verb-based experiment, our technique for establishing the relatedness between the syntactic signatures and the semantic classes, is mediated by the verbs themselves. We compare the two orthogonal groupings of the inventory of verbs: the semantic classes defined by Levin and the sets of verbs that correspond to each of the derived syntactic signatures. When these two groupings overlap, we have discovered a mapping from the syntax of the verbs to their semantics, via the verb tokens. More specifically, we define the overlap index

<sup>3</sup>An example of the intensional characterization of the Levin classes are the definitions of Lexical Conceptual Structures which correspond to each of Levin’s semantic classes. See (Dorr and Voss, to appear).

as the number of overlapping verbs divided by the average of the number of verbs in the semantic class and the number of verbs in the syntactic signature. Thus an overlap index of 1.00 is a complete overlap and an overlap of 0 is completely disjoint. In this experiment, the sets of verbs with a high overlap index are of interest.

When we parsed the 1668 example sentences in Part Two of Levin’s book (including the negative examples), these sentences reduce to 282 unique patterns. The 191 sets of sentences listed with each of the 191 semantic classes in turn reduces to 748 distinct syntactic signatures. Since there are far more syntactic signatures than the 191 semantic classes, it is clear that the mapping between signatures and semantic classes is not direct. Only 12 mappings have complete overlaps. That means 6.3% of the 191 semantic classes have a complete overlap with a syntactic signature.

The results of this experiment are shown in Table 2. Three values are shown for each of the six variations in the experiment: the mean overlap, the median overlap, and the percentage of perfect overlaps (overlaps of value 1.00). In every case, the median is higher than the mean. Put another way, there is always a cluster of good overlaps, but the general tendency is to have fairly poor overlaps.

The six variations of the experiment are as follows. The first distinction is whether or not to count the negative evidence. We note that the use of negative examples, i.e., plausible uses of the verb in contexts which are disallowed, was a key component of this experiment. There are 1082 positive examples and 586 negative examples. Although this evidence is useful, it is not available in dictionaries, corpora, or other convenient resources that could be used to extend Levin’s classification. Thus, to extend our approach to novel word senses (i.e., words not occurring in Levin), we would not be able to use negative evidence. For this reason, we felt it necessary to determine the importance of negative evidence for building uniquely identifying syntactic signatures. As one might expect, throwing out the negative evidence degrades the usefulness of the signatures across the board. The results which had the negative evidence are shown in the left-hand column of numbers in Table 2, and the results which had only positive evidence are shown in the right-hand side.

The second, three-way, distinction involves prepositions, and breaks the two previous distinctions involving negative evidence into three sub-cases. Because we were interested in the role of prepositions in the signatures, we also ran the experiment with two different parse types: ones that ignored the actual prepositions in the pp’s, and ones that ignored all information except for the values of the prepositions. Interestingly, we still got useful results with these impoverished parses, although fewer semantic classes had uniquely-identifying syntactic signatures under these conditions. These results are shown in the three major rows of Table 2.

The best result, using both positive and negative evidence to identify semantic classes, gives 6.3% of the verbs having perfect overlaps relating semantic classes to syntactic signatures. See Table 2 for the full results.

### 3.3 Experiment 2: Class-based Approach

In this experiment, we attempt to discover whether each class-based syntactic signature uniquely identifies a sin-

Verb-based Experiment (No Disambiguation)

|                      | Overlap | With Negative Evidence | No Negative Evidence |
|----------------------|---------|------------------------|----------------------|
| Marked Prepositions  | Median  | 0.10                   | 0.09                 |
|                      | Mean    | 0.17                   | 0.17                 |
|                      | Perfect | 6.3%                   | 5.2%                 |
| Ignored Prepositions | Median  | 0.10                   | 0.09                 |
|                      | Mean    | 0.17                   | 0.16                 |
|                      | Perfect | 6.3%                   | 4.2%                 |
| Only Prepositions    | Median  | 0.10                   | 0.09                 |
|                      | Mean    | 0.16                   | 0.15                 |
|                      | Perfect | 3.1%                   | 3.1%                 |

Table 2: Verb-Based Results

gle semantic class. By focusing on the classes, the verbs are implicitly disambiguated: the word sense is by definition the sense of the verb as a member of a given class. To compare these signatures with the previous verb-based signatures, it may be helpful to note that a verb-based signature is the union of all of the class-based signatures of the semantic classes that the verb appears in.

The outline for this class-based experiment is as follows:

1. Automatically extract syntactic information from the example sentences to yield the syntactic signature for the class.
2. Determine which semantic classes have uniquely-identifying syntactic signatures.

If we use the class-based syntactic signatures containing preposition-marked pp’s and both positive and negative evidence, the 1668 example sentences reduce to 282 syntactic patterns, just as before. But now there are 189 class-based syntactic signatures, as compared with 748 verb-based signatures from before. 187 of them uniquely identify a semantic class, meaning that 97.9% of the classes have uniquely identifying syntactic signatures. Four of the semantic classes do not have enough syntactic information to distinguish them uniquely.<sup>4</sup>

Although the effects of the various distinctions were present in the verb-based experiment, these effects are much clearer in the class-based experiments. The effects of negative and positive evidence, as well as the three ways of handling prepositions show up much clearer here, as is clear in Table 4.

In the class-based experiment, we counted the percentage of semantic classes that had uniquely identifying signatures. In the verb-based experiment, we counted the number of perfect overlaps (i.e., index of 1.00) between the verbs as grouped in the semantic classes and grouped by syntactic signature. The overall results of the suite of experiments, illustrating the role of disambiguation, negative evidence, and prepositions, is shown in Table 4. There were three ways of treating prepositions: (i) mark the pp with the preposition, (ii) ignore the preposition, and (iii) keep only the prepositions. For these different strategies, we see the percentage of perfect overlaps, as well as both the

<sup>4</sup>Two of these classes correspond to one of the two non-unique signatures, and two correspond to the other non-unique signature.

Class-based Experiment (Disambiguated Verbs)

|                      | Overlap | With Negative Evidence | No Negative Evidence |
|----------------------|---------|------------------------|----------------------|
| Marked Prepositions  | Median  | 1.00                   | 1.00                 |
|                      | Mean    | 0.99                   | 0.93                 |
|                      | Perfect | 97.9%                  | 88.0%                |
| Ignored Prepositions | Median  | 1.00                   | 1.00                 |
|                      | Mean    | 0.96                   | 0.69                 |
|                      | Perfect | 87.4%                  | 52.4%                |
| Only Prepositions    | Median  | 1.00                   | 0.54                 |
|                      | Mean    | 0.82                   | 0.57                 |
|                      | Perfect | 66.5%                  | 42.9%                |

Table 3: Class-Based Results

median and mean overlap ratios for each experiment. These data show that the most important factor in the experiments is word-sense disambiguation.

|                            | With Negative Evidence | No Negative Evidence |
|----------------------------|------------------------|----------------------|
| <b>No Disambiguation</b>   |                        |                      |
| Marked Prepositions        | 6.3%                   | 5.2%                 |
| Ignored Prepositions       | 6.3%                   | 4.2%                 |
| Only Prepositions          | 3.1%                   | 3.1%                 |
| <b>With Disambiguation</b> |                        |                      |
| Marked Prepositions        | 97.9%                  | 88.0%                |
| Ignored Prepositions       | 87.4%                  | 52.4%                |
| Only Prepositions          | 66.5%                  | 42.9%                |

Table 4: Overall Results

## 4 Semantic Classification of Novel Words

As we saw above, word sense disambiguation is critical to the success of any lexical acquisition algorithm. The Levin-based verbs are already disambiguated by virtue of their membership in different classes. The difficulty, then, is to disambiguate and classify verbs that do not occur in Levin. Our current direction is to make use of the results of the first two experiments, i.e., the relation between syntactic patterns and semantic classes, but to use two additional techniques for disambiguation and classification of non-Levin verbs: (1) extraction of synonym sets provided in WordNet (Miller, 1985), an online lexical database containing thesaurus-like relations such as synonymy; and (2) selection of appropriate synonyms based on correlations between syntactic information in Longman’s Dictionary of Contemporary English (LDOCE) (Procter, 1978) and semantic classes in Levin. The basic idea is to first determine the most likely candidates for semantic classification of a verb by examining the verb’s synonym sets, many of which intersect directly with the verbs classified by Levin. The “closest” synonyms are then selected from these sets by comparing the LDOCE grammar codes of the unknown word with those associated with each synonym candidate. The use of LDOCE as a syntactic filter on the semantics derived from WordNet is the key to resolving word-sense ambiguity during the acquisition process. The full acquisition algorithm is as follows:

Given a verb, check Levin class.

1. If in Levin, classify directly.
2. If not in Levin, find synonym set from WordNet.
  - (a) If synonym in Levin, select the class that has the closest match with canonical LDOCE codes.
  - (b) If no synonyms in Levin or canonical LDOCE codes are completely mismatched, hypothesize new class.

Note that this algorithm assumes that there is a “canonical” set of LDOCE codes for each of Levin’s semantic classes. Table 5 describes the significance of a subset of the syntactic codes in LDOCE. (The total number of codes is 174.) We have developed a relation between LDOCE codes and Levin classes, in much the same way that we associated syntactic signatures with the semantic classes in the earlier experiments. These canonical codes are for syntactic filtering (checking for the closest match) in the classification algorithm.

As an example of how the word-sense disambiguation process and classification, consider the non-Levin verb *attempt*. The LDOCE specification for this verb is: T1 T3 T4 WV5 N. Using the synonymy feature of WordNet, the algorithm automatically extracts five candidate classes associated with the synonyms of this word: (1) Class 29.6 “Masquerade Verbs” (*act*), (2) Class 29.8 “Captain Verbs” (*pioneer*), (3) Class 31.1 “Amuse Verbs” (*try*), (4) Class 35.6 “Ferret Verbs” (*seek*), and (5) Class 55.2 “Complete Verbs” (*initiate*). The synonyms for each of these classes have the following LDOCE encodings, respectively: (1) I I-FOR I-ON I-UPON L1 L9 T1 N; (2) L9 T1 N; (3) I T1 T3 T4 WV4 N; (4) I I-AFTER I-FOR T1 T3; and (5) T1 T1-INTO N. The largest intersection with the syntactic codes for *attempt* occurs with the verb *try* (T1 T3 T4 N). However, Levin’s class 31.1 is not the correct class for *attempt* since this sense of *try* has a “negative amuse” meaning (e.g., *John’s behavior tried my patience*). In fact, the codes T1 T3 T4 are not part of the canonical class-code mapping associated with class 31.1. Thus, *attempt* falls under case 2(b) of the algorithm, and a new class is hypothesized. This is a case where word-sense disambiguation has allowed us to classify a new word and to enhance Levin’s verb classification by adding a new class to the word *try* as well. In our experiments, our algorithm found several additional non-Levin verbs that fell into this newly hypothesized class, including *aspire*, *attempt*, *dare*, *decide*, *desire*, *elect*, *need*, and *swear*.

We have automatically classified 10,000 “unknown” verbs, i.e., those not occurring in the Levin classification, using this technique. These verbs are taken from English “glosses” (i.e., translations) provided in bilingual dictionaries for Spanish and Arabic.<sup>5</sup> As a preliminary measure of success, we picked out 84 LDOCE *control vocabulary* verbs, (i.e., primitive words used for defining dictionary entries) and hand-checked our results. We found that 69 verbs were classified correctly,

<sup>5</sup>The Spanish-English dictionary was built at the University of Maryland; The Arabic-English dictionary was produced by Alpnet, a company in Utah that develops translation aids. We are also in the process of developing bilingual dictionaries for Korean and French, and we will be porting our LCS acquisition technology to these languages in the near future.

i.e., 82% accuracy.

## 5 Summary

We have conducted two experiments with the intent of addressing the issue of word-sense ambiguity in extraction from machine-readable resources for the construction of large-scale knowledge sources. In the first experiment, verbs that appeared in different classes collected the syntactic information from each class it appeared in. Therefore, the syntactic signature was composed from all of the example sentences from every class the verb appeared in. In some cases, the verbs were semantically unrelated and consequently the mapping from syntax to semantics was muddled. The second experiment attempted to determine a relationship between a semantic class and the syntactic information associated with each class. Not surprisingly, but not insignificantly, this relationship was very clear, since this experiment avoided the problem of word sense ambiguity. These experiments served to validate Levin’s claim that verb semantics and syntactic behavior are predictably related and also demonstrated that a significant component of any lexical acquisition program is the ability to perform word-sense disambiguation.

We have used the results of our first two experiments to help in constructing and augmenting online dictionaries for novel verb senses. We have used the same syntactic signatures to categorize new verbs into Levin’s classes on the basis of WordNet and LDOCE. We are currently porting these results to new languages using online bilingual lexicons.

## Acknowledgements

The research reported herein was supported, in part, by Army Research Office contract DAAL03-91-C-0034 through Battelle Corporation, NSF NYI IRI-9357731, Alfred P. Sloan Research Fellow Award BR3336, and a General Research Board Semester Award.

## References

- Alshawi, H. 1989. Analysing the Dictionary Definitions. In B. Boguraev and T. Briscoe, editor, *Computational Lexicography for Natural Language Processing*. Longman, London, pages 153–169.
- Boguraev, B. and T. Briscoe. 1989. Utilising the LDOCE Grammar Codes. In B. Boguraev and T. Briscoe, editor, *Computational Lexicography for Natural Language Processing*. Longman, London, pages 85–116.
- Brent, M. 1993. Unsupervised Learning of Lexical Syntax. *Computational Linguistics*, 19:243–262.
- Church, K. and P. Hanks. 1990. Word Association Norms, Mutual Information and Lexicography. *Computational Linguistics*, 16:22–29.
- Copestake, A., T. Briscoe, P. Vossen, A. Ageno, I. Castellon, F. Ribas, G. Rigau, H. Rodríguez, and A. Samiotou. 1995. Acquisition of Lexical Translation Relations from MRDS. *Machine Translation*, 9.
- Dorr, B., J. Garman, and A. Weinberg. 1995. From Syntactic Encodings to Thematic Roles: Building Lexical Entries for Interlingual MT. *Machine Translation*, 9.

| LDOCE Code         | Arguments       | Adjuncts  | Example                                      |
|--------------------|-----------------|-----------|--|
| I                  | —               | —         | Olivier is <b>acting</b> tonight             |
| I-AFTER            | —               | PP[after] | She <b>sought</b> after the truth            |
| I-FOR              | —               | PP[for]   | They <b>sought</b> for the right one         |
| I-ON               | —               | PP[on]    | He <b>acted</b> on our suggestion            |
| I-UPON             | —               | PP[upon]  | The drug <b>acted</b> upon the pain          |
| L1                 | NP              | —         | He <b>acts</b> the experienced man           |
| L9                 | ADV/PP          | —         | The play <b>acts</b> well                    |
| T1                 | NP              | —         | I <b>pioneered</b> the new land              |
| T1-INTO            | NP              | PP[into]  | We <b>initiated</b> him into the group       |
| T3                 | VP[to+inf]      | —         | He <b>tried</b> to do it                     |
| T4                 | VP[+prog]       | —         | She <b>tried</b> eating the new food         |
| WV4                | -ing adjectival | —         | I've had a <b>trying</b> day                 |
| WV5                | -ed adjectival  | —         | He was convicted for <b>attempted</b> murder |
| N (denominal verb) | —               | —         | <b>pioneer</b> (noun)                        |

Table 5: Sample Syntactic Codes used in LDOCE

- Dorr, B. and C. Voss. to appear. A Multi-Level Approach to Interlingual MT: Defining the Interface between Representational Languages. *International Journal of Expert Systems*.
- Dubois, D. and P. Saint-Dizier. 1995. Construction et représentation de classes sémantiques de verbes: une coopération entre syntaxe et cognition. manuscript, IRIT- CNRS, Toulouse, France.
- Farwell, D., L. Guthrie, and Y. Wilks. 1993. Automatically Creating Lexical Entries for ULTRA, a Multilingual MT System. *Machine Translation*, 8(3).
- Fillmore, C.J. 1968. The Case for Case. In E. Bach and R.T. Harms, editor, *Universals in Linguistic Theory*. Holt, Rinehart, and Winston, pages 1–88.
- Grimshaw, J. 1990. *Argument Structure*. MIT Press, Cambridge, MA.
- Gruber, J.S. 1965. *Studies in Lexical Relations*. Ph.D. thesis, MIT, Cambridge, MA.
- Guthrie, J., L. Guthrie, Y. Wilks, and H. Aidinejad. 1991. Subject-Dependent Co-occurrence and Word Sense Disambiguation. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, pages 146–152, University of California, Berkeley, CA.
- Hearst, M. 1991. Noun Homograph Disambiguation Using Local Context in Large Text Corpora. In *Using Corpora*, University of Waterloo, Waterloo, Ontario.
- Jackendoff, R. 1983. *Semantics and Cognition*. MIT Press, Cambridge, MA.
- Jackendoff, R. 1990. *Semantic Structures*. MIT Press, Cambridge, MA.
- Klavans, J.L. and E. Tzoukermann. 1996. Dictionaries and Corpora: Combining Corpus and Machine-readable Dictionary Data for Building Bilingual Lexicons. *Machine Translation*, 10.
- Levin, B. 1993. *English Verb Classes and Alternations: A Preliminary Investigation*. Chicago, IL.
- Lonsdale, D., T. Mitamura, and E. Nyberg. 1996. Acquisition of Large Lexicons for Practical Knowledge-Based MT. *Machine Translation*, 9.
- Miller, G. 1985. WORDNET: A Dictionary Browser. In *Proceedings of the First International Conference on Information in Data*, University of Waterloo Centre for the New OED, Waterloo, Ontario.
- Neff, M. and M. McCord. 1990. Acquiring Lexical Data from Machine-Readable Dictionary Resources for Machine Translation. In *Third International Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages (TMI-90)*, Austin, Texas.
- Pinker, S. 1989. *Learnability and Cognition: The Acquisition of Argument Structure*. MIT Press, Cambridge, MA.
- Procter, P. 1978. *Longman Dictionary of Contemporary English*. Longman, London.
- Sanfilippo, A. and V. Poznanski. 1992. The Acquisition of Lexical Knowledge from Combined Machine-Readable Dictionary Resources. In *Proceedings of the Applied Natural Language Processing Conference*, pages 80–87, Trento, Italy.
- Walker, D. and R. Amsler. 1986. The Use of Machine-readable Dictionaries in Sublanguage Analysis. In R. Grishman and R. Kittredge, editors, *Analyzing Language in Restricted Domains*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, pages 69–83.
- Wilks, Y., D. Fass, C.M. Guo, J.E. McDonald, and T. Plate. 1990. Providing Machine Tractable Dictionary Tools. *Machine Translation*, 5(2):99–154.
- Wilks, Y., D. Fass, C.M. Guo, J.E. McDonald, T. Plate, and B.M. Slator. 1989. A Tractable Machine Dictionary as a Resource for Computational Semantics. In B. Boguraev and T. Briscoe, editor, *Computational Lexicography for Natural Language Processing*. Longman, London, pages 85–116.
- Wu, D. and X. Xia. 1995. Large-Scale Automatic Extraction of an English-Chinese Translation Lexicon. *Machine Translation*, 9.
- Yarowsky, D. 1992. Word-Sense Disambiguation: Using Statistical Models of Roget's Categories Trained on Large Corpora. In *Proceedings of the Fourteenth International Conference on Computational Linguistics*, pages 454–460, Nantes, France.