

Increasing the coverage of a domain independent dialogue lexicon with VERBNET

Benoit Crabbé*, **Myroslava O. Dzikovska***, **William de Beaumont†**, **Mary Swift†**

* ICCS-HCRC, University of Edinburgh, 2 Buccleuch Place, EH8 9LW, Edinburgh, UK
{bcrabbe,mdzikovs}@inf.ed.ac.uk

†Department of Computer Science, University of Rochester, Rochester NY, USA
{wdebeaum,swift}@cs.rochester.edu

Abstract

This paper investigates how to extend coverage of a domain independent lexicon tailored for natural language understanding. We introduce two algorithms for adding lexical entries from VERBNET to the lexicon of the TRIPS spoken dialogue system. We report results on the efficiency of the method, discussing in particular precision versus coverage issues and implications for mapping to other lexical databases.

1 Introduction

This paper explores how different lexicons can be integrated with the goal of extending coverage of a deep parser and semantic interpreter. Lexical semantic databases (Kipper et al., 2000; Johnson and Fillmore, 2000; Dorr, 1997) use a frame-based model of lexical semantics. Each database groups words in classes where predicative words and their arguments are described. The classes are generally organised in an inheritance structure. Each such database can be used, among other things, to perform semantic interpretation. However, their actual structures are quite different, reflecting different underlying methodological approaches to lexical description, and this results in representation that are not directly compatible. Since no such database has full coverage of English, it is worth combining them in order to get a lexicon with better coverage and a unified representation for English.

We explore the issues related to merging verb descriptions from two lexical databases, which

have both syntactic and semantic incompatibilities, and compare two techniques for aligning semantic classes and the syntax-semantics mappings between them. The resulting lexicon is to be used in precise interpretation tasks, so its consistency and accuracy are a high priority. Thus, though it is possible to generate lexical entries automatically (Kwon and Hovy, 2006; Swift, 2005), we use a semi-automatic method in which an expert hand-checks the automatically generated entries before adding them to the lexicon. Therefore, our goal is to maximise the number of new useful entries added to the lexicon while minimising the number of entries that are discarded or hand-edited.

We take the mapping between the TRIPS lexicon and the VERBNET lexical database as a case study for our experiment. The TRIPS lexicon is used together with a parser to provide a natural language understanding component for several dialogue applications in different domains. It outputs highly detailed semantic representations suitable for complex dialogue tasks such as problem-solving and tutoring dialogue, *inter alia*. An essential feature of TRIPS is the integration of a detailed lexical semantic representation, semantic classes and theta role assignments in the parsing process.

Semantic types and role labelling are helpful in both deep (Tetreault, 2005) and shallow interpretation tasks (Narayanan and Harabagiu, 2004). TRIPS provides a convenient test case because its grammar is already equipped with the formal devices required to build up a frame-based semantic representation including this information.¹

¹While wide coverage grammars such as the *English Re-*

We chose VERBNET to extend the TRIPS lexicon because it includes a detailed syntax-semantic mappings, thus providing a more convenient interface to the syntactic component of the grammar than lexicons where this connection is left unclear, such as FRAMENET. However the methods described here are designed to be reusable for merging other lexical databases, in particular we intend to experiment with FRAMENET in the near future.

The plan of the paper is as follows: we first describe the target lexicon (Section 2) and the source lexicon (Section 3) for our experiment before describing the methodology for integration (Section 4). We finally present an evaluation of the techniques in Section 5.

2 The TRIPS Lexicon

The TRIPS lexicon (Dzikovska, 2004) is the target of the mapping procedure we describe in Section 4. It includes syntactic and semantic information necessary to build semantic representations usable in dialogue systems. The TRIPS parser is equipped with a fairly detailed grammar, but a major restriction on coverage in new domains is often lack of lexical information. The lexicon used in our evaluation comprised approximately 700 verb lemmas with 1010 senses (out of approximately 2500 total word senses, covering both open- and closed-class words). The lexicon is designed for incremental growth, since the lexical representation is domain-independent and the added words are then re-used in new domains.

A graphical representation of the information stored in the TRIPS lexicon and used in parsing is shown in Figure 1. The lexicon is a list of canonical word entries each of which is made of a set of sense definitions comprised of a LF type and a syntax-semantic template.

Semantic classes (LF types) in the TRIPS lexicon are organised in a domain-independent ontology (the LF ontology). The LF Ontology was originally based on a simplified version of FRAMENET

source Grammar (Copestake and Flickinger, 2000) build deep semantic representations which account for scoping and temporal structure, their lexicons do not provide information related to word senses and role labels, in part due to the additional difficulty involved building a wide coverage lexicon with the necessary lexical semantic information.

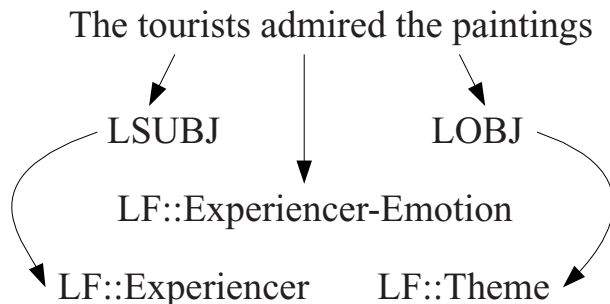


Figure 1: Information in the TRIPS word sense definition for mapping between syntactic and semantic roles.

(Baker et al., 1998; Dzikovska et al., 2004), with each LF type describing a particular situation, object or event and its participants. *Syntax-Semantics Templates* (or templates) capture the linking between the syntax and semantics (LF type and semantic roles) of a word. The semantic properties of an argument are described by means of a semantic role assigned to it and selectional restrictions.²

The TRIPS grammar contains a set of independently described lexical rules, such as the passive or dative shift rules, which are designed to create non-canonical lexical entries automatically, while preserving the linking properties defined in the canonical entry.

In this context adding an entry to the lexicon requires determining both the list of LF types and the list of templates for canonical contexts, that is, the list of mappings between a logical frame and a canonical subcategorization frame.

3 VERBNET

VERBNET (Kipper et al., 2000) provides an actual implementation of the descriptive work carried out by Levin (1993), which has been extended to cover prepositional constructions and corpus-based subcategorization frames (Kipper et al., 2004; Kipper et al., 2006).

VERBNET is a hierarchical verb lexicon in which verbs are organised in classes. The fundamental assumption underlying the classification is that the members of a given class share a similar syntactic

²The selectional restrictions are domain independent and specified using features derived from EuroWordNet (Vossen, 1997; Dzikovska et al., to appear).

behaviour, that is, they pattern in the same set of alternations, and are further assumed to share common semantic properties.³

VERBNET classes are organised in an inheritance hierarchy. Each class includes a set of members (verbs), a set of (subcategorization) frames and a set of semantic descriptions. Frames are descriptions of the linking between syntax and semantics for that class. Each frame argument contains a syntactic category augmented with syntactic features, and a corresponding thematic role. Each class also specifies a set of additional selectional restriction features. VERBNET further includes for each class a semantic description stated in terms of event semantics, that we ignore in this paper.

4 Methodology

The methodology used in the mapping process consists of two steps. First we translate the source, VERBNET, to an intermediate representation best suited for parsing purposes. Second this intermediate representation is translated to a specific target, here the TRIPS lexicon. At this stage of our work, the translation from VERBNET to the intermediate representation mainly concerns normalising syntactic information coded in VERBNET to make them easier to handle for parsing purposes, and the translation from the intermediate representation to the TRIPS lexicon focuses on translating semantic information. This architecture is best understood as a cross compilation scheme: we further expect to reuse this intermediate representation for producing outputs for different parsers and to accept inputs from other lexical databases such as FRAMENET.

4.1 The intermediate representation

The intermediate representation is a lexical representation scheme mainly tailored for parsing: in this context, a lexicon is thus made of a set of words, each of which consists of a lemma, a syntactic category and a list of sense definitions. Each sense definition has a name and a frame. The name of the sense definition is actually the name of the VERBNET class it derives from. The frame of the sense definition has a list of arguments, each of which con-

³In practice, it turns out that there are exceptions to that hypothesis (Kipper, 2005).

sists of a syntactic category, a syntactic function, a thematic role and possibly a set of prepositions and syntactic feature structures.

The content of the intermediate representation uses the following data categories. Syntactic categories, thematic roles and features are those used in VERBNET. We further add the syntactic functions described in (Carroll et al., 1998). Specifically, two categories left implicit in VERBNET by the use of feature structures are made explicit here: prepositional phrases (PP) and sentential arguments (S).

Each argument described in a sense definition frame is marked with respect to its coreness status. The coreness status aims to provide the lexicon with an operational account for common discrepancies between syntax and semantics descriptions. This status may be valued as *core*, *non-core* or *non-sem* and reflects the status of the argument with respect to the syntax-semantics interface.

Indeed, there is a methodological pitfall concerning the mapping between thematic roles and syntactic arguments: semantic arguments are not defined following criteria identical to those for syntactic arguments. The main criterion for describing semantic arguments is their participation in the event, situation, object described by the frame whereas the criterion for describing syntactic arguments is based on the obligatoriness or the specificity of the argument with respect to the verb. The following example illustrates such conflicts:

- (1) a. It is raining
- b. I am walking to the store

The *It* in example (1a) plays no role in the semantic representation, but is obligatory in syntax since it fills a subject position. The locative PP in example (1b) is traditionally not treated as an argument in syntax, rather as a modifier, hence it does not fill a complement position. Such phrases are, however, classified in VERBNET as part of the frames. Following this, we distinguish three kinds of arguments: *non-sem* as in (1a) are syntactic-only arguments with no semantic contribution. *non-core* as in (1b) contribute to the semantics but are not subcategorized.

4.2 From VERBNET to the intermediate representation

Given VERBNET as described in Section 3 and the intermediate representation we described above, the translation process requires mainly (1) to turn the class based representation of VERBNET into a list-of-word based representation (2) to mark arguments for coreness (3) to merge some arguments and (4) to annotate arguments with syntactic functions.

The first step is quite straightforward. Every member m of every VERBNET class C is associated with every frame of C yielding a new sense definition in the intermediate representation for m .

In the second step, each argument receives a coreness mark. Arguments marked as *non-core* are adverbs, and prepositional phrases introduced by a large class of prepositions (e.g. spatial prepositions). The arguments marked as *non-sem* are those with an impersonal *it*, typically members of the *weather* class. All other arguments listed in VERBNET frames are marked as *core*.

In the third step, syntactic arguments are merged to correspond better to phrase-based syntax.⁴ For example, the VERBNET encoding of subcategorization frames splits prepositional frames on two slots: one for the preposition and one for the noun phrase. We have merged the two arguments, to become a PP, also merging their syntactic and semantic features. Other merges at this stage include merging possessive arguments such as *John's brother* which are described with three argument slots in VERBNET frames. We merged them as a single NP.

The last step in the translation is the inference of syntactic functions. It is possible to reasonably infer syntactic functions from positional arguments and syntactic categories by (a) considering the following obliquity order over the set of syntactic functions used in the intermediate representation:⁵

$$(2) \text{ NCSubj} < \text{DOBJ} < \text{OBJ2} < \{\text{IOBJ}, \text{XCOMP}, \text{CCOMP}\}$$

⁴We also relabel some categories for convenience without affecting the process. For instance, VERBNET labels both clausal arguments and noun phrases with the category NP. The difference is made with syntactic features. We take advantage of the features to relabel clausal arguments with the category S.

⁵This order is partial, such that the 3 last functions are unordered wrt to each other. These functions are the subset of the functions described in (Carroll et al., 1998) relevant for handling VERBNET data.

and by (b) considering this problem as a transduction problem over two tapes. One tape being the tape of syntactic categories and the second the tape of syntactic functions. Given that, we designed a transducer that implements a category to function mapping. It implements the above obliquity order together with an additional mapping constraint: nouns can only map to NCSUBJ, DOBJ, prepositional phrases can only map to OBJ2, IOBJ, infinitival clauses can only map to XCOMP and finite clauses to CCOMP.

We further added refinements to account for frames that do not encode their arguments following the canonical obliquity order: for dealing with dative shift encoded in VERBNET with two different frames and for dealing with impersonal contexts, so that we eventually used the transducer in Figure 2. All states except 0 are meant to be final. The transduction operates only on *core* and *non-sem* arguments, *non-core* arguments are systematically associated with an adjunct function. This transducer is capable of correctly handling the majority of VERBNET frames, finding a functional assignment for more than 99% of the instances.

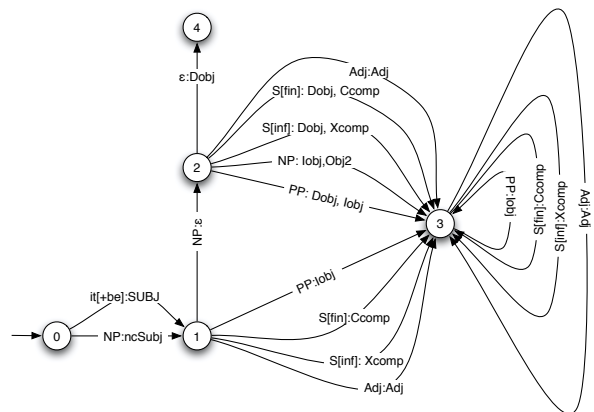


Figure 2: A transducer for assigning syntactic functions to ordered subcategorization frames

4.3 From Intermediate representation to TRIPS

Recall that a TRIPS lexical entry is comprised of an LF type with a set of semantic roles and a template representing the mappings from syntactic functions to semantic roles. Converting from our intermediate representation to the TRIPS format involves two steps:

- For every word sense, determine the appropriate TRIPS LF type
- Establish the correspondence between VERBNET and TRIPS syntactic and semantic arguments, and generate the appropriate mapping in the TRIPS format.

We investigated two strategies to align semantic classes (VERBNET classes and TRIPS LFs). Both use a class intersection algorithm as a basis for decision: two semantic classes are considered a match if they are associated with the same lexical items.

The intersection algorithm takes advantage of the fact that both VERBNET and TRIPS contain lexical sets. A lexical set for VERBNET is a class name and the set of its members, for TRIPS it is an LF type and the set of words that are associated with it in the lexicon. Our intersection algorithm computes the intersection between every VERBNET lexical set and every TRIPS lexical set. The sets which intersect are then considered as candidate mappings from a VERBNET class to a TRIPS class.

However, this technique produces many 1-word class intersections, and leads to spurious entries. We considered two ways of improving precision: first by requiring a significantly large intersection, second by using syntactic structure as a filter. We discuss them in turn.

4.4 Direct Mapping Between Semantic Representations

The first technique which we tried for mapping between TRIPS and VERBNET semantic representations is to map the classes directly. We consider all candidate mappings between the TRIPS and VERBNET classes, and take the match with the largest intersection. We then align the semantic roles between the two classes and produce all possible syntax-semantics mappings specified by VERBNET. This technique has the advantage of providing the most complete set of syntactic frames and syntax-semantics mappings which can be retrieved from VERBNET. However, since VERBNET lists many possible subcategorization frames for every word, guessing the class incorrectly is very expensive, resulting in many spurious senses generated. We use a class intersection threshold to improve reliability.

VERBNET ROLE	TRIPS ROLES
Theme	LF::THEME, LF::ADDRESSEE, LF::ALONG, LF::ENTITY
Cause	LF::CAUSE, LF::THEME
Experiencer	LF::EXPERIENCER, LF::COGNIZER
Source	LF::FROM-LOC, LF::SOURCE, LF::PATH
Destination	LF::GOAL, LF::TO-LOC
Recipient	LF::RECIPIENT, LF::ADDRESSEE, LF::GOAL
Instrument	LF::INSTRUMENT

Table 1: Sample VERBNET to TRIPS role mappings

At present, we count an LF type match as successfully guessed if there is an intersection in lexical entries above the threshold (we determined 3 words as a best value by finding an optimal balance of precision/recall figures over a small gold-standard mapping set). Since the classes contain closely related items, larger intersection means a more reliable mapping. If the VERBNET class is not successfully mapped to an LF type then no TRIPS lexical entry is generated.

Once the correspondence between the LF type and the VERBNET class has been established, semantic arguments have to be aligned between the two classes. We established a role mapping table (a sample is shown in Table 1), which is an extended version of the mapping from Swift (2005). The role mapping is one to many (each VERBNET role maps to 1 to 8 TRIPS roles), however, since the appropriate LF type has been identified prior to argument mapping, we usually have a unique mapping based on the roles defined by the LF type.⁶

Once the classes and semantic roles have been aligned, the mapping of syntactic functions between the intermediate representation and TRIPS syntax is quite straightforward. Functional and category mappings are one to one and do not raise specific problems. Syntactic features are also translated into TRIPS representation.

To illustrate the results obtained by the automatic mapping process, two of the sense definitions generated for the verb *relish* are shown in Figure 3. The TRIPS entries contain references to the class description in the TRIPS LF ontology (line introduced by

⁶In rare cases where more than 1 correspondence is possible, we are using the first value in the intersection as the default.

```

;; entries
(relish
  (SENSES
    ((EXAMPLE "The tourists admired the paintings")
      (LF-PARENT LF::EXPERIENCER-EMOTION)
      (TEMPL VN-EXPERIENCER-THEME-TEMPL-84))
    ((EXAMPLE "The children liked that the clown had a red nose")
      (LF-PARENT LF::EXPERIENCER-EMOTION)
      (TEMPL VN-EXPERIENCER-THEME-XP-TEMPL-87))
  ))
;;Templates
(VN-EXPERIENCER-THEME-TEMPL-84
  (ARGUMENTS
    (LSUBJ (% NP) LF::EXPERIENCER)
    (LOBJ (% NP) LF::THEME)
  ))
(VN-EXPERIENCER-THEME-XP-TEMPL-87
  (ARGUMENTS
    (LSUBJ (% NP) LF::EXPERIENCER)
    (LCOMP (% CP (vform fin) (ctype s-finite)) LF::THEME)
  ))

```

Figure 3: Sample TRIPS generated entries

LF-PARENT) and to a template (line introduced by TEMPL) generated on the fly by our syntactic conversion algorithm. The first sense definition and template in Figure 3 represent the same information shown graphically in Figure 1. Each argument in a template is assigned a syntactic function, a feature structure describing its syntactic properties, and a mapping to a semantic role defined in the LF type definition (not depicted here).

4.5 Filtering with syntactic structure

The approach described in the previous section provides a fairly complete set of subcategorization frames for each word, provided that the class correspondence has been established successfully. However, it misses classes with small intersections and classes for which some but not all members match (see Section 5 for discussion). To address these issues we tried another approach that automatically generates all possible class matches between TRIPS and VERBNET, again using class member intersection, but using the a TRIPS syntactic template as an additional filter on the class match. For each potential match, a human evaluator is presented with the following:

```

{confidence score
 {verbs in TRIPS-VN class intersection}/
 LF-type TRIPS-template
 => VN-class: {VN class members}}

```

The confidence score is based on the number of verbs in the intersection, weighted by taking into account the number of verbs remaining in the respective TRIPS and VERBNET classes. The template used for filtering is taken from all templates that oc-

cur with the TRIPS words in this intersection (one match per template is generated for inspection). For example:

```

93.271%
{clutch,grip,clasp,hold,wield,grasp}/
lf::body-manipulation agent-theme-xp-templ
=> hold-15.1-1: {handle}

```

This gives the evaluator additional syntactic information to make the judgement on class intersections. The evaluator can reject entire class matches, or just individual verbs from the VERBNET class which don't quite fit an otherwise good match. We only used the templates already in TRIPS (those corresponding to each of the word senses in the intersection) to avoid overwhelming the evaluator with a large number of possibly spurious template matches resulting from an incorrect class match. This technique allows us to pick up class matches based on a single member intersection, such as:

```

7.814%
{swallow}/
lf::consume agent-theme-xp-templ
=> gobble-39.3-2: {gulp,guzzle,quaff,swig}

```

However, the entries obtained are not guaranteed to cover all frames in VERBNET because if a given alternation is not already covered in TRIPS, it is not derived from VERBNET with this method.

5 Evaluation and discussion

Since our goal in this evaluation is to balance the coverage of VERBNET with precision, we correspondingly evaluate along those two dimensions. For both techniques, we evaluate how many word senses were added, and the number of different words defined and VERBNET classes covered. As a measure of precision we use, for those entries which were retrieved, the percentage of those which could be taken "as is" (good entries) and the percentage of entries which could be taken with minor edits (for example, changing an LF type to a more specific subclass, or changing a semantic role in a template). The results of evaluation are shown in Table 2.⁷

Since for mapping with syntax filtering we considered all possible TRIPS-VERBNET intersections, it in effect presents an upper bound the number of words shared between the two databases. Further

⁷"nocos" table rows exclude the other_cos VERBNET class, which is exceptionally broad and skews evaluation results.

Type	Class mapping					Mapping with syntax filtering				
	Total	Good	Edit	Bad	%usable	Total	Good	Edit	Bad	%usable
Sense	3075	1000	196	1879	0.39	11036	1688	87	9261	0.16
Word	744	274	98	372	0.5	2138	1211	153	714	0.64
Class	15	10	1	4	0.73	198	129	2	67	0.66
Sense-nocos	1136	654	196	286	0.75	7989	1493	87	6409	0.20
Word-nocos	422	218	98	106	0.75	1763	1059	153	491	0.69
Class-nocos	14	9	1	4	0.71	197	128	2	67	0.65

Table 2: Evaluation results for different acquisition techniques. %usable = (good + editable) / bad”.

extension would require extending the TRIPS LF Ontology with additional types to cover the missing classes. As can be seen from this table, 65% of VERBNET classes have an analogous class in TRIPS. At the same time, there is a very large number of class intersections possible, so if all possible intersections are generated, only a very small percentage of generated word senses (16%) is usable in the combined system. Thus developing techniques to filter out the irrelevant senses and class matches is important for successful hand-checking.

Our evaluation also shows that while class intersection with thresholding provides higher precision, it does not capture many words and verb senses. One reason for this is data sparsity. TRIPS is relatively small, and both TRIPS and VERBNET contain a number of 1-word classes, which cannot be reliably mapped without human intervention. This problem can be alleviated in part as the size of the database grows. We expect this technique to have better recall when the combined lexicon is used to merge with a different lexical database such as FRAMENET.

However, a more difficult issue to resolve is differences in class structure. VERBNET was built around the theory of syntactic alternations, while TRIPS used FRAMENET structure as a starting point, simplifying the role structure to make connection to parsing more straightforward (Dzikovska et al., 2004). Therefore TRIPS does not require that all words associated with the same LF type share syntactic behaviour, so there are a number of VERBNET classes with members which have to be split between different TRIPS classes based on additional semantic properties. 70% of all good matches in the filtering technique were such partial matches. This significantly disadvantages the thresholding tech-

nique, which provides the mappings on class level, not allowing for splitting word entries between the classes.

We believe that the best solution can be found by combining these two techniques. The thresholding technique could be used to establish reliable class mappings, providing classes where many entries could be transferred “as is”. The mapping can then be examined to determine incorrect class mappings as well as the cases where classes should be split based on individual words. For those entries judged reliable in the first pass, the syntactic structure can be transferred fully and quickly, while the syntactic filtering technique, which requires more manual checking, can be used to transfer other entries in the intersections where class mapping could not be established reliably.

Establishing class and member correspondence is a general problem with merging any two semantic lexicons. Similar issues have been noted in comparing FRAMENET and VERBNET (Baker and Ruppenhofer, 2002). A method recently proposed by Kwon and Hovy (2006) aligns words in different semantic lexicons to WordNet senses, and then aligns semantic roles based on those matches. Since we are designing a lexicon for semantic interpretation, it is important for us that all words should be associated with frames in a shared hierarchy, to be used in further interpretation tasks. We are considering using this alignment technique to further align semantic classes, in order to produce a shared database for interpretation covering words from multiple sources.

6 Conclusion

In this paper, we presented a methodology for merging lexicons including syntactic and lexical semantic

information. We developed a model based on cross-compilation ideas to provide an intermediate representation which could be used to generate entries for different parsing formalisms. Mapping semantic properties is the most difficult part of the process, and we evaluated two different techniques for establishing correspondence between classes and lexical entries, using TRIPS and VERBNET lexicons as a case study. We showed that a thresholding technique has a high precision, but low recall due to inconsistencies in semantic structure, and data sparsity. We can increase recall by partitioning class intersections more finely by filtering with syntactic structure. Further refining the mapping technique, and using it to add mappings to other lexical databases such as FRAMENET is part of our ongoing work.

Acknowledgements

We thank Karin Kipper for providing us useful documentation on the VERBNET feature system, and Charles Callaway for technical help with the final version. This material is based on work supported by grants from the Office of Naval Research under numbers N000140510048 and N000140510043, from NSF #IIS-0328811, DARPA #NBCHD030010 via subcontract to SRI #03-000223 and NSF #E1A-0080124.

References

- C. F. Baker and J. Ruppenhofer. 2002. Framenet's frames vs. Levin's verb classes. In *Proceedings of the 28th Annual Meeting of the Berkeley Linguistics Society*, pages 27–38.
- C. F. Baker, C. Fillmore, and J. B. Lowe. 1998. The Berkeley Framenet project. In *Proceedings of CoLing-ACL*, Montreal.
- J. Carroll, E. Briscoe, and A. Sanfilippo. 1998. Parser evaluation: A survey and a new proposal. In *Proceedings of LREC-98*.
- A. Copestake and D. Flickinger. 2000. An open source grammar development environment and broad-coverage English grammar using HPSG. In *Proceedings of LREC-2000*, Athens, Greece.
- B. Dorr. 1997. Large-scale dictionary construction for foreign language tutoring and interlingual machine translation. *Machine Translation*, 12(4):271–375.
- M. O. Dzikovska, M. D. Swift, and J. F. Allen. 2004. Building a computational lexicon and ontology with FrameNet. In *Proceedings of LREC workshop on Building Lexical Resources from Semantically Annotated Corpora*, Lisbon.
- M. O. Dzikovska, M. D. Swift, and J. F. Allen. to appear. Customizing meaning: Building domain-specific semantic representations from a generic lexicon. In H. Bunt, editor, *Computing Meaning, Volume 3*, Studies in Linguistics and Philosophy. Kluwer.
- M. O. Dzikovska. 2004. *A Practical Semantic Representation for Natural Language Parsing*. Ph.D. thesis, University of Rochester, Rochester NY.
- C. Johnson and C. J. Fillmore. 2000. The FrameNet tagset for frame-semantic and syntactic coding of predicate-argument structure. In *Proceedings ANLP-NAACL 2000*, Seattle, WA.
- K. Kipper, H. T. Dang, and M. Palmer. 2000. Class-based construction of a verb lexicon. In *Proceedings of AAAI*, Austin.
- K. Kipper, B. Snyder, and M. Palmer. 2004. Using prepositions to extend a verb lexicon. In *Proceedings of HLT-NAACL 2004 Workshop on Computational Lexical Semantics*, pages 23–29, Boston, MA.
- K. Kipper, A. Korhonen, N. Ryant, and M. Palmer. 2006. Extending Verbnet with novel verb classes. In *Proceedings of LREC-2006*.
- K. Kipper. 2005. *Verbnet: A broad coverage, comprehensive verb lexicon*. Ph.D. thesis, University of Pennsylvania.
- N. Kwon and E. H. Hovy. 2006. Integrating semantic frames from multiple sources. In A. F. Gelbukh, editor, *CICLing*, volume 3878 of *Lecture Notes in Computer Science*, pages 1–12. Springer.
- B. Levin. 1993. *English Verb Classes and Alternations*. The University of Chicago Press.
- S. Narayanan and S. Harabagiu. 2004. Question answering based on semantic structures. In *Proceedings of International Conference on Computational Linguistics (COLING 2004)*, Geneva, Switzerland.
- M. Swift. 2005. Towards automatic verb acquisition from Verbnet for spoken dialog processing. In *Proceedings of Interdisciplinary Workshop on the Identification and Representation of Verb Features and Verb Classes*, Saarbruecken, Germany.
- J. Tetreault. 2005. *Empirical Evaluations of Pronoun Resolution*. Ph.D. thesis, University of Rochester.
- P. Vossen. 1997. Eurowordnet: A multilingual database for information retrieval. In *Proceedings of the Delos workshop on Cross-language Information Retrieval*.