Proposed Next Steps for The Preposition Project

Ken Litkowski
CL Research
9208 Gue Road
Damascus, MD 20872 USA
ken@clres.com

Abstract

One task of SemEval-2007 was designed to disambiguate prepositions using a set of instances developed under The Preposition Project. Ostensibly, this was a completion of the project, since the results from the task were generally quite favorable, at a level consistent with other disambiguation tasks of lexical items. However, subsequent work suggests that many nuances about preposition behavior have not yet been captured, particularly in the absence of sufficient information in dictionary entries of prepositions to permit this disambiguation. This paper lays out the gaps and proposes next steps for the characterization of preposition behavior.

1 Introduction

The Preposition Project (TPP) is designed to provide a comprehensive account of the behavior of English prepositions (Litkowski & Hargraves (2005); Litkowski & Hargraves (2006)). TPP began with a comprehensive set of preposition definitions, developed a set of characterizing properties, assembled a set of sentences illustrating many of these senses, provided a lexicographical analysis of the senses, and provided an online version where the preposition properties could be examined and downloaded.[1] The next logical step for TPP was to make this information available for analysis and to hope that it would be adequate for disambiguation.

SemEval-2007 included a task for preposition disambiguation (Litkowski & Hargraves (2007)). While this task attracted only three participants, they were able to achieve results comparable to those of other word-sense disambiguation tasks and to provide useful approaches exploring preposition behavior. The latter included use of feature spaces with maximum entropy analysis, examination of preposition substitutability, and mutual disambiguation with other sentence elements.

While these results seemed to be complete, subsequent developments have revealed many opportunities for richer and more comprehensive characterization of preposition behavior. We explore these issues by describing further analyses of preposition data (section 2), advances in preposition disambiguation (section 3), opportunities for further advances, particularly from closer integration with FrameNet (section 4), the need for improvements in preposition lexicography, particularly for disambiguating data (section 5), and the need for representations of preposition meaning that can be used in NLP applications (section 6).

2 New Developments in TPP

The basic data in TPP consists of (1) a comprehensive sense inventory with definitions and examples, and for each sense, (2) identification of the preposition class, (3) semantic relation characterizations, (4) complement properties, (5) attachment properties, (6) permissible syntactic positions, (7) FrameNet frame and frame element characterizations, (8) synonymous prepositions, and (9) sense relations. For major prepositions, a lexicographic "treatment" is available to provide insights into the behavior of the preposition along with identification of idioms that use the preposition. These data, including a DIMAP dictionary and a MySQL database, can be downloaded from Online TPP.

The semantic relation for a sense has been placed into the preposition class identified for that sense. Since the preposition definitions admit of placement into a taxonomic hierarchy, it has been possible to perform a digraph analysis of the prep-

[1] Online TPP: http://www.clres.com/cgi-bin/onlineTPP/find_prep.cgi.

osition classes (Hargraves & Litkowski (2008)).[2] This digraph analysis followed the general principles described in Litkowski (2002) and is based on the observation that most preposition definitions end in a preposition that can serve as the basis for a hierarchy. During TPP, this terminal preposition (e.g., *with* in the definition of **into**, "to a position of contact with") was disambiguated by the lexicographer. The resulting analyses by preposition classes allows the construction of a digraph for each preposition class, as well as for the entire set.

Litkowski (2009) performed an initial analysis of these preposition classes, particularly considering the FrameNet (Baker et al. (1998)) frame elements that seemed to be associated with the classes. By jointly considering the digraphs and the FrameNet frame elements, this analysis suggested that perhaps six of the classes could likely be subsumed under the remaining classes. This result is examined below in considering proposed next steps for TPP.

In the SemEval-2007 task, 34 prepositions were included for disambiguation. Nearly 25,000 sentences were used in this task, two-thirds for training and one-third for the test run. The set of instances has now been further cleaned and extended, with the addition of 3,000 FrameNet sentences for prepositions tagged in TPP along with 7,500 instances from the Oxford Dictionary of English sentence dictionary. This latter set was designed to provide up to 20 example sentences for each sense. The resultant set provides examples for 632 senses of 260 prepositions. The Online TPP now provides links so that the examples for each sense, where available, can be examined. In addition, these data can be downloaded as MySQL database tables for the definitions and for the sentences.

3 Recent Advances in Preposition Disambiguation

Ye & Baldwin (2007) provided the basic analysis for preposition disambiguation. They used a maximum-entropy-based system that analyzed collocation features, syntactic features, and semantic-role features. They found that collocations (bag of words and bag of synsets) were the most important features, with little contribution from syntax and semantics. Their system established a baseline of 69 percent accuracy.

Tratz & Hovy (2009) delved deeper into a characterization of collocation features, taking into account higher order aspects of the context, including the governing phrase, part of speech type, and a WordNet-based semantic class. Their maximum entropy model improved the accuracy of the disambiguation to 75 percent.

Hovy et al. (2010) extended this work to improve fine-grained accuracy to 85 percent and coarse-grained accuracy to 92 percent. In this work, also using a maximum entropy model, contextual, WordNet-based and further miscellaneous features were used. Essentially, the feature sets were refined and elaborated. In this analysis, focus was also given to identifying the specific features that contributed to the disambiguation. The most dominant feature was the governor or point of attachment of the prepositional phrase, with important contributions from the head of the prepositional object and the word to the left of the preposition (which may frequently be the governor, but sometimes not). The general conclusions were that selective context is better than simple windows and that using a simple tagger to identify context was comparable to using a dependency parser.

The preceding studies were supervised disambiguation, relying on the training data in SemEval-2007 to build models for the test set. Hovy et al. (2011) used these findings to develop and examine models for unsupervised preposition disambiguation. These models essentially focused on the heads and the objects for disambiguating the preposition. They achieved an accuracy of 56 percent. The import of this result is that a much larger set of data can be examined (as compared to that available in the SemEval task) and characterized with increasing features and then used as an enhanced mechanism for disambiguation.

4 Opportunities for Further Advances in Preposition Disambiguation

The preceding section suggests that increasingly nuanced characterizations of context may bring about further progress in preposition disambiguation. The question is where are such improvements

[2] The classes are **Activity**, **Agent**, **Backdrop**, **Barrier**, **Cause**, **Doubles**, **Exception**, **Means/Medium**, **Membership**, **Party**, **Possession**, **Quantity**, **Scalar**, **Spatial**, **Substance**, **Tandem**, **Target**, **Temporal**, **Topic**, and **Void**.

to be found. Most of the SemEval-based studies did not actually make use of the TPP data. In the next sections, we examine particular TPP components (from among those listed earlier) and how they might be useful (while indicating some of their shortcomings).

4.1 Complement Properties

While it may seem obvious that the preposition object should be of paramount importance in disambiguation, particularly since the definitions characterize the object, this is not the case. The preceding studies showed that the object is of less importance than the governor. In addition, a preposition confers some additional meaning to its object. In "abdication from the throne", "throne" is a concrete object to which is added a meaning that it is also to be considered a role or responsibility.[3]

Notwithstanding, a large number of senses have objects which can be characterized quite specifically and used for disambiguation. The lexicographer's characterizations of these senses need to be examined in detail and labeled as such. Future studies may be able to exploit such information.

4.2 Attachment Properties

As suggested by the several systems, the governor or point of attachment of the preposition is of paramount importance. The main improvements in disambiguation have stemmed from the ability to pinpoint the governor, primarily from the use of dependency grammars. The attachment properties in TPP may provide some further guidance in this area. While these characterizations have not been written systematically, the lexicographer did reuse many of the descriptions. Further analysis of these descriptions may prove useful.

Casual examination of these descriptions, however, suggests that there might be some further difficulties. This may be an area where further nuanced refinements may be required. In the example above, "abdication from the throne", the applicable sense of **from** is "indicating separation or removal", with the attachment property described as "verbs and verbal nouns of separation." The problem is that there is no classification system in any lexical resource (dictionary or otherwise) in which this is a class.

[3] This may also be viewed as type coercion.

Several of the systems used WordNet file numbers as a surrogate for semantic classes. The single word "separate" has 13 verb senses with 6 different file numbers. FrameNet puts the verb "separate" into 4 distinct frames. A Roget-style thesaurus has 3 categories. Each of these resources can be used to find synonyms that might provide a coherent set. This would suggest that, rather than using seeds to provide a coherent class, it might be better to examine the definitions of governors for which this sense of **from** applies and identify commonalities of meaning that give rise to this sense.

4.3 Quirk Syntax and Comprehensive View of Preposition Behavior

The TPP data includes a characterization of the syntactic functions of prepositional phrases, identified as Quirk syntax. This is based on a classification scheme from Quirk et al. (1985): (1) noun postmodifier, (2) adverbial (a) adjunct, (b) subjunct, (c) disjunct, (d) conjunct, (3) (a) verb complement, and (b) adjective complement. Hovy et al. (2010) noted this behavior in discussing fronting ("*In May*, prices dropped by 5%."). They do not discuss the extent to which the fronting feature contributed to disambiguation.

The effect of fronting behavior or other types of syntactic behavior suggests that additional strategies might be necessary to identify governors. In the example above, a word-selection rule might be able to look for the first verb after a following comma when a prepositional phrase is fronted. Further examination of any constraints on syntactic function might be useful.

As mentioned above, Hovy et al. (2010) also performed their analysis at the coarse-grained level, achieving an accuracy of 92 percent. Since the SemEval data cover only the 34 most frequent and also most polysemous prepositions, it is questionable whether their behavior mirrors overall preposition behavior. With the addition of 7,500 instances for an additional 226 prepositions, many of which are monosemous, it is possible that the application of their methods by preposition classes may provide further insights into preposition behavior.

4.4 FrameNet Frames and Frame Elements

Conceptually, every prepositional phrase is a FrameNet frame element. This is reflected in the sentences used in TPP and SemEval-2007. An immediate question is why these were not used or usable in preposition disambiguation. In the example sense for **from**, there are 352 sentences, but the phrases use 13 frames and 14 distinct frame elements. A similar variety is to be found for all the preposition senses. Several steps have been taken to examine this diversity.

4.4.1 Frame Element Taxonomy

FrameNet 1.5 contains 1019 frames using 1170 distinct frame element names. While the frame element names capture distinct frame semantics, they are frequently reused (e.g., **Agent** is used in 180 frames). In addition, through the use of its frame-to-frame relations, by which a finer-grained frame is related to a coarser-grained frame and the conceptual similarity of frame element names is laid out specifically, it is possible to lay out a frame element taxonomy.

By analyzing the full set of frame-to-frame relations, it has been possible to create hypernymic links between frame elements. A frame element dictionary was created with each frame element as an entry and with a sense capturing the hypernymic relation. This permitted a digraph analysis of the frame elements. Initially, this digraph contained several inconsistencies and cycles. The lexicologists at CL Research carefully examined the definitions of the frame elements and the cycles to obtain a strict taxonomy of the frame elements.[4] The resultant taxonomy was able to map the entire set of frame elements into 12 primitives.[5,6]

The similarity between the set of frame element primitives and the preposition classes (see above) is striking. When viewed at this level, the possibility of collapsing the frame elements associated with each sense may be useful in preposition disambiguation. Furthermore, it should be noted that frame elements are very similar to the semantic relation names constructed by the lexicographer in TPP. This may prove useful in considering the fact that one of the biggest issues facing FrameNet is its coverage.

4.4.2 Prepositions in FrameNet

An important factor in considering next steps for TPP is how prepositions are treated in FrameNet. In general, prepositions are not targets in FrameNet. However, there are 94 lexical unit senses (out of a total of almost 12,000 lexical units) identified as prepositions, within 20 frames. Their presence raises several questions, some of which delve into the foundations of FrameNet. Litkowski (2007) explores steps for integrating TPP data into FrameNet.

Several of the lexical units identified as prepositions are phrases rather than bare prepositions (e.g., *at variance*, *under investigation*). The full phrase is a FrameNet target and gives rise to its own frame; these will not be further considered here.

Each lexical unit sense has its own definition, frequently from the *Concise Oxford Dictionary*. Since TPP uses definitions from the *Oxford Dictionary of English*, the definitions are frequently identical or slight variants of one another; this facilitates a mapping between TPP and FrameNet.

A major issue about any mapping is the status of a FrameNet preposition that is treated as a target and given its own frame, whereas TPP views a prepositional phrase as a frame element, rather than a frame. To understand this conflict, we consider the basic approach in FrameNet tagging, which focuses on individual sentences and individual words as targets. The question is how to treat a sentence that may have multiple targets. Baker et al. (2007) briefly described how frames might be viewed in a full-text representation with multiple targets. To deal with this issue, they viewed a full annotation of a sentence as roughly equivalent to a dependency parse, so that multiple frames within a sentence would fall into a dependency relationship. That is, one frame can be viewed as a slot filler in a higher frame. Based on this approach, then, a preposition that has its own frame is likely to be subsumed into a higher frame in a dependency parse.

This approach does not solve all the issues associated with preposition senses in FrameNet. An example is provided by the **Locative_relation**

[4] Each frame element in a frame has its own definition. Thus, there are 180 definitions for **Agent**. Many of these are identical, but the FrameNet lexicographers made no attempt to ensure consistency. This actually proved useful, since it allowed distinctions to be made where a frame element name was polysemous.

[5] These are **Reason**, **Degree**, **Purpose**, **Cause**, **Instrument**, **Entity**, **Time**, **Phenomenon**, **Role**, **Path**, **Topic**, and **State**.

[6] The Frame Element Taxonomy is at http://www.clres.com/db/feindex.html. This site includes a full description of how it was developed. The data for the taxonomy, including the DIMAP dictionary and a MySQL database, can be downloaded.

frame. This frame has four primary frame elements: **Ground** (the preposition object), **Figure** (the governor), **Distance**, and **Direction**. The first two frame elements are considered core, while the latter two are extra-thematic (i.e., not essential to the frame). FrameNet has about 33 preposition senses that evoke this frame. In examining the definitions of these lexical unit senses and other spatial prepositions (in either TPP or FrameNet), values (i.e., fillers for the frame elements) for **Distance** (**next to**: "a position immediately to one side of") or **Direction** (**below**: "at a lower level than") are present. However, FrameNet currently has no mechanism by which a frame element value is directly incorporated into the meaning of a target word. Many targets identify an "incorporated FE", but this doesn't quite go far enough.

Another problematic issue is the inheritance structure of FrameNet frames. The frame **Time_vector** is identified as using the **Direction** frame (thereby collapsing temporal characterizations into spatial ones). In the definition of this frame, it is noted that the **Domain** frame element of **Direction** should be specified as "Time". However, as is the case with frame element values, this mechanism is not overtly built into FrameNet inheritance. On the other hand, this operation might more properly be viewed as one that should be performed when attempting to build representations (discussed further in section 6).

4.4.3 Building Consistency Between TPP and FrameNet

This section has identified many areas of investigation that might be used in attempting to make use of FrameNet frame elements in preposition disambiguation. As described in Section 3, all disambiguation systems made use of WordNet properties, such as synsets and file numbers, to capture semantic properties of the governors and complements. We believe that FrameNet frame elements offer an intermediate level of semantic characterization that might be usefully explored.

The release of FrameNet 1.5 may facilitate this exploration. Each tagged sentence in FrameNet now has a full parse, along with some amount of chunking corresponding to the frame elements. This applies to each FrameNet sentence in TPP, which includes the FrameNet sentence number and a link back to the targeted lexical unit file where this detail is located.

At the same time, it is important to consider the limitations of FrameNet. In the FrameNet data included in SemEval, there were no examples of many senses. Thus, this set cannot be considered a balanced representation of behavior even for those prepositions that were included. In addition, since the coverage of FrameNet admittedly does not cover all possible frames, it is likely that many senses in TPP will not correspond to existing frames. Attempting to identify those senses may prove useful in identifying gaps in FrameNet coverage.

5 Building Disambiguating Lexical Entries for Prepositions

While much progress has been made in disambiguating preposition senses, these results have not yet influenced the construction of the lexical entries in ordinary dictionaries. In this section, we consider some possible desiderata for preposition senses.

Atkins & Rundell (2008) set a standard for lexical entries of content words (verbs, nouns, adjectives, and adverbs) by identifying a set of typical constructions in which these words are used. These are lists of lexicographically relevant co-occurrences for these types of head words, and are intended to be based on corpus evidence. A comparable set has not been developed for prepositions.

A first step in building lexical entries for prepositions is to extract the features that have been most important in the disambiguation of each sense. Using the maximum entropy models, these would be the features that provide the greatest information gain. This would provide the general feature space that is relevant for each sense and would provide the basis for organizing the features relevant to each preposition.

At this point, it would be necessary to move away from the individual prepositions to consider preposition classes. Thus, for example, there are many spatial and temporal preposition senses across many prepositions. By considering the sets of features of all spatial preposition senses, it becomes possible to build a more coherent decision tree.

This approach has been demonstrated in detail by Müller et al. (2010a) and Müller et al. (2010b) for German spatial and temporal prepositions, as well as the initial stages for other classes. The basis for doing something comparable for English prepositions may lie in the combination of data from the maximum entropy results, the digraphs of the TPP classes (section 2 above), and the consideration of preposition senses in FrameNet (section 4.4.2 above).

In addition to their development of an annotation schema for spatial and temporal prepositions, Müller et al. (2010a) and Müller et al. (2010b) have also implemented mechanisms for validating such annotations in corpus instances. This is an important step for ensuring an appropriate level of sense granularity in the sense inventory. With such a mechanism, it would be possible to ascertain the extent of inter-annotator agreement and to rework the sense characterizations as necessary.

OntoNotes (Hovy et al. (2006)) has also used this approach in working with the WordNet sense inventory. A similar approach is also proposed for prepositions in Hovy et al. (2011) using Amazon's Mechanical Turk.

An important factor in the success of these validation schemes is the representativeness of the corpus instances that are used. As pointed out above, it is questionable whether the FrameNet preposition instances provide complete coverage. While the addition of the Oxford sentence examples in TPP may provide a more complete coverage, this is not representative.

Corpus Pattern Analysis (CPA) (Hanks, 2004) may provide a more balanced approach. CPA is being applied to verbs and begins with a sample from the British National Corpus, thus providing representative coverage. Then, for each verb, a set of syntagmatic patterns is developed. The next step is to tag each corpus instance with one of these patterns. All instances in the sample (or the complete set for verbs appearing fewer than 250 times) must be tagged.

An important component of CPA is the comprehensive characterization of verb behavior, which is encapsulated in the use of various templates to be filled for each sense. A similar structure would be required for prepositions. The development of such a structure would require a careful analysis of the features that have been identified in the maximum entropy models for preposition disambiguation.

Procedures that have been taken to implement a CPA of prepositions is described in Litkowski (2012). This differs slightly from the CPA of verbs, where the analysis focuses on individual verbs, without considering how one verb may relate to another. For prepositions, using data from TPP and the various disambiguation methods, it is possible to take a more comprehensive analysis.

6 Representations of Preposition Meaning

While preposition disambiguation has been the focus of the discussion thus far, the purpose of identifying the appropriate sense is to obtain a representation that can be plugged into a representation of a larger text. Current lexicographic practices are not geared toward such an end, so any attempt to do so must be regarded as at an initial stage. The preceding discussion, however, suggests a number of possible areas that might be investigated in the development of representations of preposition meaning.

An obvious first possibility is the use of frame semantics as the core representation schema. As discussed above, however, a key issue is whether to use frames or frame elements. A prepositional phrase is, by definition, in some relation to another element of a sentence, suggesting that the role of the phrase is as a frame element. However, we saw that many prepositions in FrameNet have been characterized as giving rise to their own frames. In addition, despite the frame status of these prepositions, they are still largely subservient to other sentence elements, particularly in dependency grammars. This approach, therefore, of attempting to build frames for each preposition, seems to be appropriate. This will have to keep in mind that there are many gaps in FrameNet that will have to be filled and that, in filling them, a broader perspective on preposition classes is likely to be appropriate.

At the present time, the object of FrameNet is simply to build representations, for individual frames and for full texts, but not necessarily for further computational purposes. Ideally, we would want our prepositional meanings to be suitable for use in various computational applications such as question answering, information extraction, and text summarization. Up to now, use of prepositions in this way has mostly been a matter of lookup, where we see, for example, whether a given preposition phrase might answer a temporal question.

This suggests that another approach to representing preposition meaning might be appropriate.

Special markup languages are being developed to handle temporal and spatial meanings in text. These developments have largely viewed temporal and spatial recognition as the identification of so-called named entities.

TimeML[7] is a a robust specification language for events and temporal expressions in natural language. It is used for time-stamping events, ordering events with respect to one another, and reasoning with temporal expressions and about events. Derczynski and Gaizauskas (2010) provide a system for analysis of features specific to temporally-annotated corpora. As expected, TimeML frequently makes use of temporal prepositions in these analyses. In view of this close tie, it would be useful to make use of TimeML representations and characterizations in representing prepositions. It should be noted that an important characteristic of TimeML is that it is explicitly designed for reasoning with temporal expressions.

SpatialML[8] (Mani et al., 2008) is a similar annotation scheme for marking up references to places in natural language. Like TimeML, SpatialML has its origins in attempts to process named entities, and as such, it focuses on identifying geo-coordinates. It also includes a region calculus for characterizing relationships among places. SpatialML also provides various attributes for characterizing places; these can be used for characterizing spatial prepositions. However, with the focus on places, this markup language does not capture all the nuances of spatial prepositions.[9]

CausalML[10] is an XML-based file format developed to store information on railway accidents or incidents gathered in a Why-Because-Analysis (WB-Analysis or WBA). In this markup language, features include actors and a characterization of factor types (e.g., internal_event, internal_state, source_state). Each factor is assigned a node identifier, which allows the construction of graphs to represent the WBA. While this system has not been used in any linguistic analysis, it illustrates a possibility for the analysis of causal relationships.

These markup languages suggest the utility of characterizing preposition meanings in terms that first capture the essence of the primitives (in the TPP preposition classes or the FrameNet frame elements) and second allow their use in computational frameworks. While the three examples focus on single relationship types, such frameworks would ultimately need to be combined into an overall framework.

This discussion of how preposition meanings might be represented is clearly in a very early stage of development. However, it suggests that computational considerations should be taken into account in this development. Further, this perspective might be able to address the issue of prepositional polysemy. When considered abstractly, the difference between coarse- and fine-grained disambiguation as noted in Hovy et al. (2011) may appear overly pedantic. This would suggest that the appropriate level of granularity should be driven by computational needs.

7 Conclusions

We have seen that the success of the preposition disambiguation task in SemEval-2007 did not constitute an end to the needs of the preposition project. Significant further progress has been made in this task and we outlined some possibilities for still further progress.

Examination of one area for further progress, the relationship between TPP and FrameNet, suggested many issues that could use further investigation. These issues also suggested the need for the improvement of preposition lexical entries to include (1) more information that would facilitate disambiguation and (2) representations of preposition meaning that could be used in computational applications.

References

B. T. Sue Atkins and Michael Rundell. 2008. *The Oxford Guide to Practical Lexicography*. Oxford University Press, Oxford.

Collin Baker, Charles Fillmore, and John Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics*, pages 86–90,

[7] http://www.timeml.org/site/index.html

[8] http://sourceforge.net/projects/spatialml

[9] SemEval-2012 has a task Spatial Role Labeling, which focuses on spatial relationships rather than geographic characterizations.

[10] Oliver Lemke (2004), http://www.tu-braunschweig.de/ifev/forschung/projekte/causalml

Montreal, Quebec, Canada, August. Association for Computational Linguistics.

Collin Baker, Michael Ellsworth, and Katrin Erk. 2007. SemEval-2007 Task 19: Frame Semantic Structure Extraction. In *Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007),* Prague, Czech Republic.

Leon Derczynski and Robert Gaizauskas. 2010. Analysing Temporally Annotated Corpora with CAVaT. In *Proceedings of the Seventh conference on International Language Resources and Evaluation (LREC'10),* Valletta, Malta. European Language Resources Association (ELRA)

Patrick Hanks. 2004. Corpus Pattern Analysis. In *EURALEX Proceedings.* Vol. I, pp. 87-98. Lorient, France: Université de Bretagne-Sud.

Orin Hargraves and Ken Litkowski. 2008. Generic Classes of Prepositions. http://www.clres.com/PrepositionClasses.htm.

Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and and Ralph Weishedel. 2006. OntoNotes: The 90% Solution. In *Proceedings of Human Language Technologies: The 2011 Annual Conference of the North American Chapter of the ACL,* New York, June. Association for Computational Linguistics.

Dirk Hovy, Stephen Tratz, and Eduard Hovy. 2010. What's in a Preposition? Dimensions of Sense Disambiguation for an Interesting Word Class. In *Coling 2010: Posters*, pages 454–462, Beijing, China, August. Coling 2010 Organizing Committee.

Dirk Hovy, Ashish Vaswani, Stephen Tratz, David Chiang, and Eduard Hovy. 2011. Models and Training for Unsupervised Preposition Disambiguation. In *Proceedings of Human Language Technologies: The 2011 Annual Conference of the North American Chapter of the Association for Computational Linguistics,* Portland, Oregon, June. Association for Computational Linguistics.

Ken Litkowski. 2002. Digraph Analysis of Dictionary Preposition Definitions. In *Proceedings of the SIGLEX/SENSEVAL Workshop on Word Sense Disambiguation: Recent Successes and Future Directions*. Philadelphia, pp. 9-16. Association for Computational Linguistics.

Ken Litkowski. 2007. Integrating Data from The Preposition Project into FrameNet. *Technical Report 07-01,* Damascus, MD: CL Research.

Ken Litkowski. 2009. Analysis of Preposition Classes. *Technical Report 09-01*, Damascus, MD: CL Research.

Ken Litkowski. 2012. Corpus Pattern Analysis of Preposition. *Technical Report 12-02*, Damascus, MD: CL Research.

Ken Litkowski and Orin Hargraves. 2005. The preposition project. *ACL-SIGSEM Workshop on "The Linguistic Dimensions of Prepositions and Their Use in Computational Linguistic Formalisms and Applications"*, pages 171–179.

Ken Litkowski and Orin Hargraves. 2006. Coverage and Inheritance in The Preposition Project. In: *Proceedings of the Third ACL-SIGSEM Workshop on Prepositions*. Trento, Italy.ACL. 89-94.

Ken Litkowski and Orin Hargraves. 2007. SemEval-2007 Task 06: Word-Sense Disambiguation of Prepositions. In *Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007),* Prague, Czech Republic.

Inderjeet Mani, Janet Hitzeman, Justin Richer, Dave Harris, Rob Quimby and Ben Wellner. 2008. SpatialML: Annotation Scheme, Corpora, and Tools. In *Proceedings of the Sixth Conference on International Language Resources and Evaluation (LREC'08),* Marrakech, Morocco. European Language Resources Association (ELRA)

Antje Müller, Olaf Hülscher, Claudia Roch, Katja Keßelmeier, Tobias Stadtfeld, Jan Strunk and Tibor Kiss. (2010a). An Annotation Schema for Preposition Senses in German. In *Proceedings of the ACL Linguistic Annotation Workshop (LAW IV).* Uppsala, Sweden.

Antje Müller, Claudia Roch, Tobias Stadtfeld, Tibor Kiss. (2010b). The Annotation of Preposition Senses in German. Submitted to Stolterfoht, B. and S. Featherston (Eds.): *Proceedings of Linguistic Evidence 2010.* Tübingen, Germany. Mouton de Gruyter.

Randolph Quirk, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. 1985. A Comprehensive Grammar of the English Language. Longman, New York.

Stephen Tratz and Dirk Hovy. 2009. Disambiguation of Preposition Sense Using Linguistically Motivated Features. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Student Research Workshop and Doctoral Consortium*, pages 96–100, Boulder, Colorado, June. Association for Computational Linguistics.

Patrick Ye and Timothy Baldwin. 2007. MELB-YB: Preposition Sense Disambiguation Using Rich Semantic Features. In *Proceedings of the 4th International Workshop on Semantic Evaluations (SemEval-2007),* Prague, Czech Republic.