Preposition Semantic Classification via PENN TREEBANK and FRAMENET

Tom O'Hara
Department of Computer Science
New Mexico State University
Las Cruces, NM 88003
tomohara@cs.nmsu.edu

Janyce Wiebe
Department of Computer Science
University of Pittsburgh
Pittsburgh, PA 15260
wiebe@cs.pitt.edu

Abstract

This paper reports on experiments in classifying the semantic role annotations assigned to prepositional phrases in both the PENN TREEBANK and FRAMENET. In both cases, experiments are done to see how the prepositions can be classified given the dataset's role inventory, using standard word-sense disambiguation features. In addition to using traditional word collocations, the experiments incorporate class-based collocations in the form of WordNet hypernyms. For Treebank, the word collocations achieve slightly better performance: 78.5% versus 77.4% when separate classifiers are used per preposition. When using a single classifier for all of the prepositions together, the combined approach yields a significant gain at 85.8% accuracy versus 81.3% for word-only collocations. For FrameNet, the combined use of both collocation types achieves better performance for the individual classifiers: 70.3% versus 68.5%. However, classification using a single classifier is not effective due to confusion among the fine-grained roles.

1 Introduction

English prepositions convey important relations in text. When used as verbal adjuncts, they are the principle means of conveying semantic roles for the supporting entities described by the predicate. Prepositions are highly ambiguous. A typical collegiate dictionary has dozens of senses for each of the common prepositions. These senses tend to be closely related, in contrast to the other parts of speech where there might be a variety of distinct senses.

Given the recent advances in word-sense disambiguation, due in part to SENSEVAL (Edmonds and Cotton, 2001), it would seem natural to apply the same basic approach to handling the disambiguation of prepositions. Of course, it is difficult to disambiguate prepositions at the granularity present in collegiate dictionaries, as illustrated later. Nonetheless, in certain cases this is feasible.

We provide results for disambiguating prepositions at two different levels of granularity. The coarse granularity is more typical of earlier work in computational linguistics, such as the role inventory proposed by Fillmore (1968), including high-level roles such as *instrument* and *location*. Recently, systems have incorporated fine-grained roles, often specific to particular domains. For example, in the Cyc KB there are close to 200 different types of semantic roles. These range from high-level roles (e.g., *beneficiaries*) through medium-level roles (e.g., *exchanges*) to highly specialized roles (e.g., *catalyst*).[1]

Preposition classification using two different semantic role inventories are investigated in this paper, taking advantage of large annotated corpora. After providing background to the work in Section 2, experiments over the semantic role annotations are discussed in Section 3. The results over TREEBANK (Marcus et al., 1994) are covered first. Treebank include about a dozen high-level roles similar to Fillmore's. Next, experiments using the finer-grained semantic role annotations in FRAMENET version 0.75 (Fillmore et al., 2001) are

[1] Part of the Cyc KB is freely available at www.opencyc.org.

presented. FrameNet includes over 140 roles, approaching but not quite as specialized as Cyc's inventory. Section 4 follows with a comparison to related work, emphasizing work in broad-coverage preposition disambiguation.

2 Background

2.1 Semantic roles in the PENN TREEBANK

The second version of the Penn Treebank (Marcus et al., 1994) added additional clause usage information to the parse tree annotations that are popular for natural language learning. This includes a few case-style relation annotations, which prove useful for disambiguating prepositions. For example, here is a simple parse tree with the new annotation format:

```
(S (NP-TPC-5 This)
    (NP-SBJ every man)
    (VP contains
        (NP *T*-5)
        (PP-LOC within
            (NP him))))
```

This shows that the prepositional phrase (PP) is providing the location for the state described by the verb phrase. Treating this as the preposition sense would yield the following annotation:

This every man contains within$_{LOC}$ him

The main semantic relations in TREEBANK are *beneficiary*, *direction*, *spatial extent*, *manner*, *location*, *purpose/reason*, and *temporal*. These tags can be applied to any verb complement but normally occur with clauses, adverbs, and prepositions. Frequency counts for the prepositional phrase (PP) case role annotations are shown in Table 1.

The frequencies for the most frequent prepositions that have occurred in the prepositional phrase annotations are shown later in Table 7. The table is ordered by entropy, which measures the inherent ambiguity in the classes as given by the annotations. Note that the *Baseline* column is the probability of the most frequent sense, which is a common estimate of the lower bound for classification experiments.

2.2 Semantic roles in FRAMENET

Berkeley's FRAMENET (Fillmore et al., 2001) project provides the most recent large-scale annotation of semantic roles. These are at a much finer granularity than those in TREEBANK, so they should prove quite useful for applications that learn detailed semantics from corpora. Table 2 shows the top semantic roles by frequency of annotation. This illustrates that the semantic roles in Framenet can be quite specific, as in the roles *cognizer*, *judge*, and *addressee*. In all, there are over 140 roles annotated with over 117,000 tagged instances.

FRAMENET annotations occur at the phrase level instead of the grammatical constituent level as in TREEBANK. The cases that involve prepositional phrases can be determined by the phrase-type attribute of the annotation. For example, consider the following annotation.

⟨S TPOS="56879338"⟩
⟨T TYPE="sense2"⟩⟨/T⟩
It$_{pnp}$ had$_{vhd}$ a$_{at0}$ sharp$_{aj0}$
,$_{pun}$ pointed$_{aj0}$ face$_{nn1}$ and$_{cjc}$
⟨C FE="**BodP**" PT="NP" GF="Ext"⟩
a$_{at0}$ feathery$_{aj0}$ tail$_{nn1}$ that$_{cjt}$
⟨/C⟩ ⟨C TARGET="y"⟩ arched$_{vvd}$⟨/C⟩
⟨C FE="**Path**" PT="*PP*" GF="Comp"⟩
over$_{avp-prp}$ its$_{dps}$ back$_{nn1}$
⟨/C⟩ .$_{pun}$⟨/S⟩

The constituent (C) tags identify the phrases that have been annotated. The target attribute indicates the predicating word for the overall frame. The frame element (FE) attribute indicates one of the semantic roles for the frame, and the phrase type (PT) attribute indicates the grammatical function of the phrase. We isolate the prepositional phrase annotation and treat it as the sense of the preposition. This yields the following annotation:

It had a sharp, pointed face and a feathery tail that arched over$_{Path}$ its back.

The annotation frequencies for the most frequent prepositions are shown later in Table 8, again ordered by entropy. This illustrates that the role distributions are more complicated, yielding higher entropy values on average. In all, there are over 100 prepositions with annotations, 65 with ten or more instances each.

Tag	Freq	Description
pp-loc	17220	locative
pp-tmp	10572	temporal
pp-dir	5453	direction
pp-mnr	1811	manner
pp-prp	1096	purpose/reason
pp-ext	280	spatial extent
pp-bnf	44	beneficiary

Table 1: TREEBANK semantic roles for PP's. Tag is the label for the role in the annotations. $Freq$ is frequency of the role occurrences.

Tag	Freq	Description
Spkr	8310	speaker
Msg	7103	message
SMov	6778	self-mover
Thm	6403	theme
Agt	5887	agent
Goal	5560	goal
Path	5422	path
Cog	4585	cognizer
Manr	4474	manner
Src	3706	source
Cont	3662	content
Exp	3567	experiencer
Eval	3108	evaluee
Judge	3107	judge
Top	3074	topic
Other	2531	undefined
Cause	2306	cause
Add	2266	addressee
Src-p	2179	perceptual source
Phen	1969	phenomenon
Reas	1789	reason
Area	1328	area
Degr	1320	degree
BodP	1230	body part
Prot	1106	protagonist

Table 2: Common FRAMENET semantic roles. The top 25 of 141 roles are shown.

3 Classification experiments

The task of selecting the semantic roles for the prepositions can be framed as an instance of word-sense disambiguation (WSD), where the semantic roles serve as the senses for the prepositions.

A straightforward approach for preposition disambiguation would be to use standard WSD features, such as the parts-of-speech of surrounding words and, more importantly, collocations (e.g., lexical associations). Although this can be highly accurate, it will likely overfit the data and generalize poorly. To overcome these problems, a class-based approach is used for the collocations, with WordNet high-level synsets as the source of the word classes. Therefore, in addition to using collocations in the form of other words, this uses collocations in the form of semantic categories.

A supervised approach for word-sense disambiguation is used following Bruce and Wiebe (1999). The results described here were obtained using the settings in Figure 1. These are similar to the settings used by O'Hara et al. (2000) in the first SENSEVAL competition, with the exception of the hypernym collocations. This shows that for the hypernym associations, only those words that occur within 5 words of the target prepositions are considered.[2]

The main difference from that of a standard WSD approach is that, during the determination of the class-based collocations, each word token is replaced by synset tokens for its hypernyms in WordNet, several of which might occur more than once. This introduces noise due to ambiguity, but given the conditional-independence selection scheme, the preference for hypernym synsets that occur for different words will compensate somewhat. O'Hara and Wiebe (2003) provide more details on the extraction of these hypernym collocations. The feature settings in Figure 1 are used in two different configurations: word-based collocations alone, and a combination of word-based and hypernym-based collocations. The combination generally produces

[2]This window size was chosen after estimating that on average the prepositional objects occur within 2.35 $+/-$ 1.26 words of the preposition and that the average attachment site is within 3.0 $+/-$ 2.98 words. These figures were produced by analyzing the parse trees for the semantic role annotations in the PENN TREEBANK.

Features:

POS-2	part-of-speech 2 words to left
POS-1:	part-of-speech 1 word to left
POS$+1$:	part-of-speech 1 word to right
POS$+2$:	part-of-speech 2 words to right
Prep	preposition being classified
WordColl$_i$:	word collocation for role i
HypernymColl$_i$:	hypernym collocation for role i

Collocation Context:

Word:	anywhere in the sentence
Hypernym:	within 5 words of target preposition

Collocation selection:

Frequency:	$f(word) > 1$
CI threshold:	$\frac{p(c\|coll)-p(c)}{p(c)} >= 0.2$
Organization:	per-class-binary

Model selection:

overall classifier: Decision tree
individual classifiers: Naive Bayes
10-fold cross-validation

Figure 1: Feature settings used in the preposition classification experiments. *CI* refers to conditional independence; the *per-class-binary* organization uses a separate binary feature per role (Wiebe et al., 1998).

the best results. This exploits the specific clues provided by the word collocations while generalizing to unseen cases via the hypernym collocations.

3.1 PENN TREEBANK

To see how these conceptual associations are derived, consider the differences in the prior versus class-based conditional probabilities for the semantic roles of the preposition 'at' in TREEBANK. Table 3 shows the global probabilities for the roles assigned to 'at'. Table 4 shows the conditional prob-

Relation	P(R)	Example
locative	.732	workers *at* a factory
temporal	.239	expired *at* midnight Tuesday
manner	.020	has grown *at* a sluggish pace
direction	.006	CDs aimed *at* individual investors

Table 3: Prior probabilities of semantic relations for 'at' in TREEBANK. $P(R)$ is the relative frequency. $Example$ usages are taken from the corpus.

Category	Relation	P(R|C)
ENTITY#1	locative	0.86
ENTITY#1	temporal	0.12
ENTITY#1	other	0.02
ABSTRACTION#6	locative	0.51
ABSTRACTION#6	temporal	0.46
ABSTRACTION#6	other	0.03

Table 4: Sample conditional probabilities of semantic relations for 'at' in TREEBANK. $Category$ is WordNet synset defining the category. $P(R|C)$ is probability of the relation given that the synset category occurs in the context.

Relation	P(R)	Example
addressee	.315	growled *at* the attendant
other	.092	chuckled heartily *at* this admission
phenomenon	.086	gazed *at* him with disgust
goal	.079	stationed a policeman *at* the gate
content	.051	angry *at* her stubbornness

Table 5: Prior probabilities of semantic relations for 'at' in FRAMENET for the top 5 of 40 applicable roles.

Category	Relation	P(R|C)
ENTITY#1	addressee	0.28
ENTITY#1	goal	0.11
ENTITY#1	phenomenon	0.10
ENTITY#1	other	0.09
ENTITY#1	content	0.03
ABSTRACTION#6	addressee	0.22
ABSTRACTION#6	other	0.14
ABSTRACTION#6	goal	0.12
ABSTRACTION#6	phenomenon	0.08
ABSTRACTION#6	content	0.05

Table 6: Sample conditional probabilities of semantic relations for 'at' in FRAMENET

abilities for these roles given that certain high-level WordNet categories occur in the context. These category probability estimates were derived by tabulating the occurrences of the hypernym synsets for the words occurring within a 5-word window of the target preposition. In a context with a concrete concept (ENTITY#1), the difference in the probability distributions shows that the *locative* interpretation becomes even more likely. In contrast, in a context with an abstract concept (ABSTRACTION#6), the difference in the probability distributions shows that the *temporal* interpretation becomes more likely. Therefore, these class-based lexical associations reflect the intuitive use of the prepositions.

The classification results for these prepositions in the PENN TREEBANK show that this approach is very effective. Table 9 shows the results when all of the prepositions are classified together. Unlike the general case for WSD, the sense inventory is the same for all the words here; therefore, a single classifier can be produced rather than individual classifiers. This has the advantage of allowing more training data to be used in the derivation of the clues indicative of each semantic role. Good accuracy is achieved when just using standard word collocations. Table 9 also shows that significant improvements are achieved using a combination of both types of collocations. For the combined case, the accuracy is 86.1%, using Weka's J48 classifier (Witten and Frank, 1999), which is an implementation of Quinlan's (1993) C4.5 decision tree learner. For comparison, Table 7 shows the results for individual classifiers created for each preposition (using Naive Bayes). In this case, the word-only collocations perform slightly better: 78.5% versus 77.8% accuracy.

3.2 FRAMENET

It is illustrative to compare the prior probabilities (i.e., P(R)) for FRAMENET to those seen earlier for 'at' in TREEBANK. See Table 5 for the most frequent roles out of the 40 cases that were assigned to it. This highlights a difference between the two sets of annotations. The common *temporal* role from TREEBANK is not directly represented in FRAMENET, and it is not subsumed by another specific role. Similarly, there is no direct role corresponding to *locative*, but it is partly subsumed by

Dataset Statistics	
Instances	26616
Classes	7
Entropy	1.917
Baseline	0.480

Experiment	Accuracy	STDEV
Word Only	81.1	.996
Combined	*86.1*	.491

Table 9: Overall results for preposition disambiguation with TREEBANK semantic roles. *Instances* is the number of role annotations. *Classes* is the number of distinct roles. *Entropy* measures non-uniformity of the role distributions. *Baseline* selects the most-frequent role. The *Word Only* experiment just uses word collocations, whereas *Combined* uses both word and hypernym collocations. *Accuracy* is average for percent correct over ten trials in cross validation. STDEV is the standard deviation over the trails. The difference in the two experiments is statistically significant at $p < 0.01$.

Dataset Statistics	
Instances	27300
Classes	129
Entropy	5.127
Baseline	0.149

Experiment	Accuracy	STDEV
Word Only	49.0	0.90
Combined	49.4	0.44

Table 10: Overall results for preposition disambiguation with FRAMENET semantic roles. See Table 9 for the legend.

Preposition	Freq	Entropy	Baseline	Word Only	Combined
through	332	1.668	0.438	0.598	0.634
as	224	1.647	0.399	0.820	0.879
by	1043	1.551	0.501	0.867	0.860
between	83	1.506	0.483	0.733	0.751
of	30	1.325	0.567	0.800	0.814
out	76	1.247	0.711	0.788	0.764
for	1406	1.223	0.655	0.805	0.796
on	1927	1.184	0.699	0.856	0.855
throughout	61	0.998	0.525	0.603	0.584
across	78	0.706	0.808	0.858	0.748
from	1521	0.517	0.917	0.912	0.882
Total	6781	1.233	0.609	0.785	0.778

Table 7: Per-word results for preposition disambiguation with TREEBANK semantic roles. *Freq* gives the frequency for the prepositions. *Entropy* measures non-uniformity of the role distributions. The *Baseline* experiment selects the most-frequent role. The *Word Only* experiment just uses word collocations, whereas *Combined* uses both word and hypernym collocations. Both columns show averages for percent correct over ten trials. *Total* averages the values of the individual experiments (except for $Freq$).

Prep	Freq	Entropy	Baseline	Word Only	Combined
between	286	3.258	0.490	0.325	0.537
against	210	2.998	0.481	0.310	0.586
under	125	2.977	0.385	0.448	0.440
as	593	2.827	0.521	0.388	0.598
over	620	2.802	0.505	0.408	0.526
behind	144	2.400	0.520	0.340	0.473
back	540	1.814	0.544	0.465	0.567
around	489	1.813	0.596	0.607	0.560
round	273	1.770	0.464	0.513	0.533
into	844	1.747	0.722	0.759	0.754
about	1359	1.720	0.682	0.706	0.778
through	673	1.571	0.755	0.780	0.779
up	488	1.462	0.736	0.736	0.713
towards	308	1.324	0.758	0.786	0.740
away	346	1.231	0.786	0.803	0.824
like	219	1.136	0.777	0.694	0.803
down	592	1.131	0.764	0.764	0.746
across	544	1.128	0.824	0.820	0.827
off	435	0.763	0.892	0.904	0.899
along	469	0.538	0.912	0.932	0.915
onto	107	0.393	0.926	0.944	0.939
past	166	0.357	0.925	0.940	0.938
Total	10432	1.684	0.657	0.685	0.703

Table 8: Per-word results for preposition disambiguation with FRAMENET semantic roles. See Table 7 for the legend.

goal. This reflects the bias of FRAMENET towards roles that are an integral part of the frame under consideration: location and time apply to all frames, so these cases are not generally annotated.

Table 9 shows the results of classification when all of the prepositions are classified together. The overall results are not that high due to the very large number of roles. However, the combined collocation approach still shows slight improvement (49.4% versus 49.0%). Table 8 shows the results when using individual classifiers. This shows that the combined collocations produce better results: 70.3% versus 68.5%. Unlike the case with Treebank, the performance is below that of the individual classifiers. This is due to the fine-grained nature of the role inventory. When all the roles are considered together, prepositions are prone to being misclassified with roles that they might not have occurred with in the training data, such as whenever other contextual clues are strong for that role. This is not a problem with Treebank given its small role inventory.

4 Related work

Until recently, there has not been much work specifically on preposition classification, especially with respect to general applicability in contrast to special purpose usages. Halliday (1956) did some early work on this in the context of machine translation. Later work in that area addressed the classification indirectly during translation. In some cases, the issue is avoided by translating the preposition into a corresponding foreign function word without regard to the preposition's underlying meaning (i.e., direct transfer). Other times an internal representation is helpful (Trujillo, 1992). Taylor (1993) discusses general strategies for preposition disambiguation using a cognitive linguistics framework and illustrates them for 'over'. There has been quite a bit of work in this area but mainly for spatial prepositions (Japkowicz and Wiebe, 1991; Zelinsky-Wibbelt, 1993).

There is currently more interest in this type of classification. Litkowski (2002) presents manually-derived rules for disambiguating prepositions, in particular for 'of'. Srihari et al. (2001) present manually-derived rules for disambiguating prepositions used in named entities.

Gildea and Jurafsky (2002) classify semantic role assignments using all the annotations in FRAMENET, for example, covering all types of verbal arguments. They use several features derived from the output of a parser, such as the constituent type of the phrase (e.g., NP) and the grammatical function (e.g., subject). They include lexical features for the headword of the phrase and the predicating word for the entire annotated frame. They report an accuracy of 76.9% with a baseline of 40.6% over the FRAMENET semantic roles. However, due to the conditioning of the classification on the predicating word for the frame, the range of roles for a particular classification is more limited than in our case.

Blaheta and Charniak (2000) classify semantic role assignments using all the annotations in TREEBANK. They use a few parser-derived features, such as the constituent labels for nearby nodes and part-of-speech for parent and grandparent nodes. They also include lexical features for the head and alternative head (since prepositions are considered as the head by their parser). They report an accuracy of 77.6% over the form/function tags from the PENN TREEBANK with a baseline of 37.8%,[3] Their task is somewhat different, since they address all adjuncts, not just prepositions, hence their lower baseline. In addition, they include the *nominal* and *adverbial* roles, which are syntactic and presumably more predictable than the others in this group. Van den Bosch and Bucholz (2002) also use the Treebank data to address the more general task of assigning function tags to arbitrary phrases. For features, they use parts of speech, words, and morphological clues. Chunking is done along with the tagging, but they only present results for the evaluation of both tasks taken together; their best approach achieves 78.9% accuracy.

5 Conclusion

Our approach to classifying prepositions according to the PENN TREEBANK annotations is fairly accurate (78.5% individually and 86.1% together), while retaining ability to generalize via class-based lexical associations. These annotations are suitable for

[3]They target all of the TREEBANK function tags but give performance figures broken down by the groupings defined in the Treebank tagging guidelines. The baseline figure shown above is their recall figure for the 'baseline 2' performance.

default classification of prepositions in case more fine-grained semantic role information cannot be determined. For the fine-grained FRAMENET roles, the performance is less accurate (70.3% individually and 49.4% together). In both cases, the best accuracy is achieved using a combination of standard word collocations along with class collocations in the form of WordNet hypernyms.

Future work will address cross-dataset experiments. In particular, we will see whether the word and hypernym associations learned over FrameNet can be carried over into Treebank, given a mapping of the fine-grained FrameNet roles into the coarse-grained Treebank ones. Such a mapping would be similar to the one developed by Gildea and Jurafsky (2002).

Acknowledgements

The first author is supported by a generous GAANN fellowship from the Department of Education. Some of the work used computing resources at NMSU made possible through MII Grants EIA-9810732 and EIA-0220590.

References

Don Blaheta and Eugene Charniak. 2000. Assigning function tags to parsed text. In *Proc. NAACL-00*.

Rebecca Bruce and Janyce Wiebe. 1999. Decomposable modeling in natural language processing. *Computational Linguistics*, 25 (2):195–208.

A. Van den Bosch and S. Buchholz. 2002. Shallow parsing on the basis of words only: A case study. In *Proceedings of the 40th Meeting of the Association for Computational Linguistics (ACL'02)*, pages 433–440. Philadelphia, PA, USA.

P. Edmonds and S. Cotton, editors. 2001. *Proceedings of the* SENSEVAL *2 Workshop*. Association for Computational Linguistics.

Charles J. Fillmore, Charles Wooters, and Collin F. Baker. 2001. Building a large lexical databank which provides deep semantics. In *Proceedings of the Pacific Asian Conference on Language, Information and Computation*. Hong Kong.

C. Fillmore. 1968. The case for case. In Emmon Bach and Rovert T. Harms, editors, *Universals in Linguistic Theory*. Holt, Rinehart and Winston, New York.

Daniel Gildea and Daniel Jurafsky. 2002. Automatic labeling of semantic roles. *Computational Linguistics*, 28(3):245–288.

M.A.K. Halliday. 1956. The linguistic basis of a mechanical thesaurus, and its application to English preposition classification. *Mechanical Translation*, 3(2):81–88.

Nathalie Japkowicz and Janyce Wiebe. 1991. Translating spatial prepositions using conceptual information. In *Proc. 29th Annual Meeting of the Assoc. for Computational Linguistics (ACL-91)*, pages 153–160.

K. C. Litkowski. 2002. Digraph analysis of dictionary preposition definitions. In *Proceedings of the Association for Computational Linguistics Special Interest Group on the Lexicon*. July 11, Philadelphia, PA.

Mitchell Marcus, Grace Kim, Mary Ann Marcinkiewicz, Robert MacIntyre, Ann Bies, Mark Ferguson, Karen Katz, and Britta Schasberger. 1994. The Penn Treebank: Annotating predicate argument structure. In *Proc. ARPA Human Language Technology Workshop*.

Tom O'Hara and Janyce Wiebe. 2003. Classifying functional relations in Factotum via WordNet hypernym associations. In *Proc. Fourth International Conference on Intelligent Text Processing and Computational Linguistics (CICLing-2003)*.

Tom O'Hara, Janyce Wiebe, and Rebecca F. Bruce. 2000. Selecting decomposable models for word-sense disambiguation: The GRLING-SDM system. *Computers and the Humanities*, 34 (1-2):159–164.

J. Ross Quinlan. 1993. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Mateo, California.

Rohini Srihari, Cheng Niu, and Wei Li. 2001. A hybrid approach for named entity and sub-type tagging. In *Proc. 6th Applied Natural Language Processing Conference*.

John R. Taylor. 1993. Prepositions: patterns of polysemization and strategies of disambiguation. In Zelinsky-Wibbelt (Zelinsky-Wibbelt, 1993).

Arturo Trujillo. 1992. Locations in the machine translation of prepositional phrases. In *Proc. TMI-92*, pages 13–20.

Janyce Wiebe, Kenneth McKeever, and Rebecca Bruce. 1998. Mapping collocational properties into machine learning features. In *Proc. 6th Workshop on Very Large Corpora (WVLC-98)*, pages 225–233, Montreal, Quebec, Canada. Association for Computational Linguistics SIGDAT.

Ian H. Witten and Eibe Frank. 1999. *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*. Morgan Kaufmann.

Cornelia Zelinsky-Wibbelt, editor. 1993. *The Semantics of Prepositions: From Mental Processing to Natural Language Processing*. Mouton de Gruyter, Berlin.