A Semantic Feature for Verbal Predicate and Semantic Role Labeling using SVMs

Hansen A. Schwartz and **Fernando Gomez** and **Christopher Millward**
School of Electrical Engineering and Computer Science
University of Central Florida, Harris Center
{hschwartz, gomez, cmillward}@cs.ucf.edu

Abstract

This paper shows that semantic role labeling is a consequence of accurate verbal predicate labeling. In doing so, the paper presents a novel type of semantic feature for verbal predicate labeling using a new corpus. The corpus contains verbal predicates, serving as verb senses, that have semantic roles associated with each argument. Although much work has been done using feature vectors with machine learning algorithms for various types of semantic classification tasks, past work has primarily shown effective use of syntactic or lexical information. Our new type of semantic feature, ontological regions, proves highly effective when used in addition to or in place of syntactic and lexical features for support vector classification, increasing accuracy of verbal predicate labeling from 65.4% to 78.8%.

Introduction

Verbal predicates and semantic roles are useful for a wide range of tasks including knowledge extraction, machine translation, and question answering. Most machine learning approaches to this type of semantic interpretation base features heavily on syntax or collocations of lexical elements. Information about a sentence is passed to a machine learning algorithm via features, which allow the algorithm to classify a word or phrase. We investigate a method of exploiting semantic information effectively in a feature based classifier. Although we theorize our semantic features can benefit many types of semantic annotation, we take on the tasks of verb predicate labeling using support vector machines (SVMs), and subsequently, of annotating semantic roles of verb predicates.

We introduce the notion of *ontological regions* as an effective type of semantic feature using the WordNet noun ontology (Miller *et al.* 1993). Ontological regions are essentially candidate selectional restrictions, which the SVM can learn to associate with a given predicate. The improvement of these features over grammatical, head word, and other features is very apparent. Furthermore, with these new features, we are able to reduce the need for some types of features and perform classification with a much smaller feature-space.

We use a corpus annotated with the verbal predicates of Gomez (Gomez 2003). The verb predicates in the corpus are based on the notion of verb predicate in the sense of (Grimshaw 1990; Pinker 1989). A predicate corresponds to a verb sense with its semantic roles (Grimshaw's a-structure). This is an important feature of our corpus, because it allows us to essentially do two labeling tasks at once, labeling both verbs and arguments. The corpus also provides the advantage of disambiguated noun senses of an argument's head word which can be useful information in the training portion of the corpus to help pick semantic features for each predicate. There are few instances of each predicate for each verb, thus providing a challenging classification task. First, we evaluate the annotation of verb predicates with a SVM, and then the annotation of semantic roles with another SVM.

Related Research

Our work is related to many other topics from semantic role labeling to verb sense disambiguation. Throughout the other work mentioned in this section, several corpora and classes of verbs are very apparent. Levin provides verb classes based on alternation patterns (Levin 1993). In PropBank, verb senses are annotated with arguments (Kingsbury, Palmer, & Marcus 2002). Similarly, FrameNet provides sentences annotated with lexical units as part of semantic frames (Baker, Fillmore, & Lowe 1998). A unique characteristic of our corpus is that the meaning of the verb, or verb predicate, is resolved, and the semantic roles associated with each predicate are linked to the WordNet ontology (Gomez 2004). Other notable differences include that our corpus provides noun senses for head nouns of grammatical relations, and there is a distinction between arguments of the verb predicates and adjuncts. We believe our semantic features may be useful with another corpus or type of verb class, but for this first work with *ontological regions*, we do not stray too far from the selectional restrictions from which the idea is based. Thus, we selected Gomez's corpus to validate the notion of *ontological region* with the benefit of also being able to show the consequence of semantic role labeling when performing predicate labeling.

The WordNet ontology has been used as a type of semantic feature in portions of past work. (Gildea & Jurafsky 2002) only used the direct hypernyms of nouns as features and found they did not improve semantic role classification. When (Dang & Palmer 2005) used the entire trace of hypernyms as features they were able to improve verb sense disambiguation results at the expense of a huge feature space. *Ontological regions* address both of these issues. They are very effective in improving results, and they do not take up a vast amount of feature space.

Grammatical relations, the basis for all syntactic features of our system, play a role similar to that of syntactic or lexical features in other work. Lexical features have been used to classify verbs into argument based classes (Merlo & Stevenson 2001). Lapata and Brew take this further and show that verb classes along with syntactic frames or selectional preferences can be useful in verb sense disambiguation (Lapata & Brew 2004) . In a general sense, this notion is captured in our work through using the grammatical relation based features. In particular, the *indicator* feature captures the grammatical context in which a verb appears. The *ontological regions* feature improves results over grammatical relations alone. (Dang & Palmer 2005) use a simple method for extracting grammatical relations (subject, direct object, indirect object, clausal complement) from a tree parse produced by (Collins 1997). In order to show that our features were still effective with an automatically parsed approach, we used the parser of Charniak (Charniak 2000) and the same method to determine grammatical relations automatically. As expected, the overall results were better when using the accurate parse of the corpus rather than this automatic method. However, Table 4 shows that we still saw a clear benefit in using *ontological regions*.

Corpus and Features

Corpus

Our corpus consisted of a set of annotated sentences from encyclopedia text (World Book Encyclopedia, World Book Inc). Half of these sentences can be downloaded from Gomez's homepage[1], while the other half are available upon request. For our experiments we used a selection of 697 clauses chosen based on verb predicates described below. The sentences are annotated with grammatical relations, semantic roles, verb predicates, and WordNet 1.6 noun senses. After each annotated sentence, a hierarchy of predicates is provided which we do not use in this work. The selectional restrictions for each predicate are not included. An example extracted from the corpus is below (information which we did not use is edited out of the corpus):

```
(Clause CL68 (SUBJ : ((NOUN SURFACE)
(NOUN TENSION)) (PHYSICAL_PHENOMENON1
SURFACE_TENSION1) ⟨(INANIMATE-CAUSE)⟩ ) (VERB
: DRAW ((AUX (WILL)) (MAIN-VERB DRAW DRAW))
⟨DRAW-FLUID:(DRAW8)⟩ ) (OBJ : ((NOUN LIQUID))
(LIQUID LIQUID1) ⟨(THEME)⟩ ) (PREP : INTO
(PREP-NP: ((UDT A) (NOUN CAPILLARY)) (TUBE1
CAPILLARY1) ⟨(GOAL)⟩ ) ) )
```

Note that the verbal predicate and the semantic roles are denoted: ⟨PREDICATE⟩ and ⟨(ROLE)⟩, where a WordNet verb sense is provided next to the predicate, and a WordNet noun sense precedes the role.

[1] http://www.eecs.ucf.edu/∽gomez/

verb	pred.	occ.	verb	pred.	occ.
abandon	2	15	blend	3	16
absorb	6	40	blow	4	23
accept	7	50	bombard	2	12
achieve	2	22	bring	7	58
act	3	16	burn	3	19
adapt	4	21	drain	2	11
adopt	2	16	draw	5	32
affect	2	15	drive	4	26
appeal	3	17	drop	3	19
approach	2	10	empty	2	16
arrange	2	22	enjoy	2	13
arrive	2	13	enter	4	26
ask	2	12	escape	3	24
assume	3	19	establish	5	21
attack	2	15	help	2	15
beat	3	9	sell	2	16
believe	3	38	**total**	**103**	**697**

Table 1: Predicates for each verb(pred.) and total occurrences (occ.)

The predicates, which map to WordNet 1.6 verb senses and have associated semantic roles, are our focus for classification. We chose the most frequent polysemous verbs and individual predicates with 3 or more instances in the corpus. The resulting list of 33 verbs and 103 predicates is included in Table 1. The average predicate has < 6.8 instances in this data. Distribution of predicates among verbs was not entirely uniform, and occurrences of particular predicates favored those found more commonly in text.

Grammatical Relations

Grammatical relations serve as the basis for nearly all of our features. When referring to grammatical relations, we are not talking about a specific feature, rather the information of a sentence on which many features are based. As mentioned previously our corpus provides grammatical relations. Below are the definitions of each relation:

subject preverbal noun phrase
object first postverbal noun phrase
object2 second postverbal noun phrase
prepositional-phrases
prepositional phrases modifying the verb

Base Features

Indicator This feature simply indicates if the grammatical relation exists.

Head Noun As defined by (Collins 1999), the head word is typically the most important word of a phrase. The base form of the head word of each noun phrase from each grammatical relation is encoded, including the noun phrases within all applicable prepositional phrases. If the content of a grammatical relation is a subclause (clausal-complement), we indicate so instead of encoding the head word as a feature.

Synsets We include WordNet 2.1 synsets as features in some of our experiments. These are groups of synonyms of nouns (Miller *et al.* 1993). The idea is that each noun can be represented more robustly through its synset.

Preposition This is the head word, as a preposition, for the entire prepositional phrase ("in", "on", etc..).

Passive This feature indicates if the voice is passive. It is important to note that we adjust passive clauses such that the *subject* becomes the *object*. Although many times a subject may be inside a prepositional phrase, it is ambiguous, and thus we do not modify prepositions. In experiments we received slightly better results on average by doing this, as would be expected since these adjustments correspond to the definition of passive voice. We use this reconstruction for all experiments reported here.

Ontological Regions

We define a type of feature which takes advantage of the WordNet noun ontology. An *ontological region* is a noun sense in WordNet which represents an area of the ontology. A noun belongs to an *ontological region* if any of it's senses can be traced via hypernym (is-a) relations to the representative of that region.

Ontological Regions as SVM Features Each verb has a set of many regions for each grammatical relation. The way in which regions are chosen is presented in the next section. Assuming one has already chosen these *ontological regions*, we must decide whether each region, as a feature, is present in the testing and training examples. In order to do this, all senses of head nouns for each grammatical relation are examined. In the case that a head word is a proper noun, we use a named entity tagger (Metzler & Sammons 2005), and map the results to WordNet 2.1 nouns in the following fashion: (LOC ↦ "location", ORG ↦ "organization", PER ↦ "person", MISC ↦ "thing"). A complete search from each noun sense to the root is performed, and all ontological regions for the particular grammatical relation are recorded in the feature vectors. Below is an example from our corpus with all semantic information removed (corresponding to a predicate testing example):

```
  (CL279 (SUBJ ((NOUN JET) (NOUN ENGINES)))
(VERB DRAW ⟨?⟩) (OBJ ((NOUN OXYGEN))) (PREP
FROM (PREP-NP ((DFART THE) (NOUN AIR)))))
```

Since "engine" is the head word for the subject, we would trace the three WordNet 2.1 senses of "engine" via the hypernym relationship to see which ontological regions the senses belong to and record those regions as features for the SVM. Similarly, the object head word "oxygen" would need to be traced to see in which regions its senses lie. If our ontological regions for the object of "draw" are *physical_entity-1, visual_communication1, psychological_feature1, substance1, object1, and content5*, then "oxygen" could be traced to *physical_entity1* and *substance1*. In this case, the object has only one sense, the correct predicate is DRAW-FLUID. Figure 1 shows the trace to the regions. In short, if any sense of a head noun is part of a region, the region is marked as a feature.

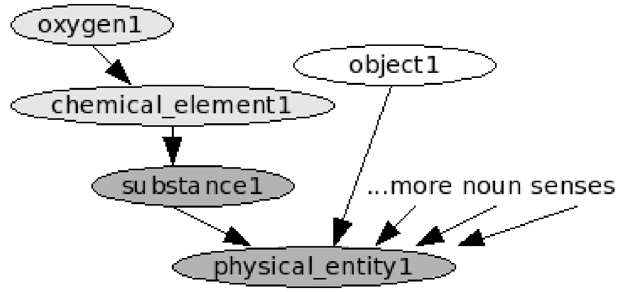

Figure 1: *physical_entity1* ontological region and trace from oxygen to this point in WordNet. The region covers all noun senses which can be traced through hypernym relationships to *physical_entity1*.

It may become clear that *ontological regions* are very similar to selectional restrictions. For a given verbal predicate, selectional restrictions state areas of the ontology a given argument (grammatical relation) must satisfy (Gomez 2004). However, for the case of feature-based classification we are trying to choose the most likely predicate rather than determine if a given verbal predicate's selectional restrictions are satisfied. Therefore, the difference between *ontological regions* and selectional restrictions is that the *ontological regions* are not specific to a predicate. They are simply an implementation of the concept behind selectional restrictions, which allow the SVM to learn which regions correspond to which predicates.

It is important to note that the *representative* for an *ontological region* is essentially a synset. However, the ontological region feature indicates if a noun sense is within the region of the ontology rooted by the *representative*, while the synset feature simply indicates if a sense of the noun belongs to a specific synset.

Choosing WordNet Representatives for Regions This section describes how WordNet's senses are chosen as *ontological regions*. There are two approaches to the selection of such representatives used in experiments in this paper. The first approach is to hand select a static set of *ontological regions* for all verbs and grammatical relations. The method requires a bit of human experience with the WordNet ontology and tends to stick with more general concepts of the ontology. The resulting list had 145 concepts[2]. Although we experimented with adding and removing regions from the general list, we found this amount to be sufficient to show the effectiveness of such a feature. Future work may focus on the optimum number of regions.

A second approach, which was used for most of our experiments, analyzes the verbal arguments in the training corpus to automatically choose representatives. An advantage

[2]The list of concepts can be found at:
http://www.eecs.ucf.edu/~hschwartz/GenORs.txt

of our corpus is that it provides WordNet 1.6 noun senses for head words of verbal arguments. The sense information is taken only from the training corpus and mapped to WordNet 2.1 senses (Daudé, Padró, & Rigau 2000). We group senses according to verb, predicate (verb sense), and grammatical relation. Each sense is traced through the ontology to the root via the hypernym relationship, such that any concept in the ontology receives a score equal to the sum of sub-concepts used as verbal arguments for the predicate. This produces a list of scores for each trio of verb, predicate, and grammatical relation (gr).

The next step is to combine the scores($Score_p$) from all predicates of a particular *(verb, gr)* pair. Each score is combined as follows:

$$total_{verb,gr}(cncpt) = \frac{\sum_{p \in Preds} score_p(cncpt)}{cnt^2}$$

The variable cnt is the number of predicates which had a score for the concept; it is intended to decrease the score of a concept that is used for multiple predicates.

Finally, we produce a set of representative concepts for every *(verb, gr)* pair. The top 75% of concepts are examined according to the *total*. Once again using the ontology, we find *significant siblings*, hyponyms of the same parent with a score at least $1/3$ that of the parent's score. If at least two *significant siblings* are found under a single parent, the siblings are chosen as *representatives*. Note that this produces choices of concepts which distinguish points of convergence in the ontology.

Experiments

SVM Setup

The success of our ability to compare features depends highly on the implementation of the SVM. Different parameters, kernel functions, and multi-class approaches can greatly affect the results of an SVM (Hsu & Lin 2001). Using a ”one-against-one” approach to multi-class classification, we implemented a separate C-SVC classifier for each verb. The program LIBSVM provided functionality needed for our tests (Chang & Lin 2001).

We use a ”10-fold” cross-validation method in order to evaluate our feature sets. The corpus is randomly broken into 10 equal-sized subsets, except for a restriction that each test predicate have at least two corresponding training examples. We use parameters based only on the training set (chosen by another cross-validation step within the training set). This insures the testing information is never in the training data, including *ontological region* choices. Figure 2 shows the steps taken. Notice step 1 chooses which ontological regions will be searched for in the corpus from only the training data. Thus, we had 10 different sets of ontological regions corresponding to each of the 10 training corpora. During steps 2 and 4 we are extracting these regions for each example (among other data) as features for input to the SVM. Therefore, the only pieces of information we assume to be in the testing corpus are the grammatical relations, but we also perform experiments with these removed using an automatic parser.

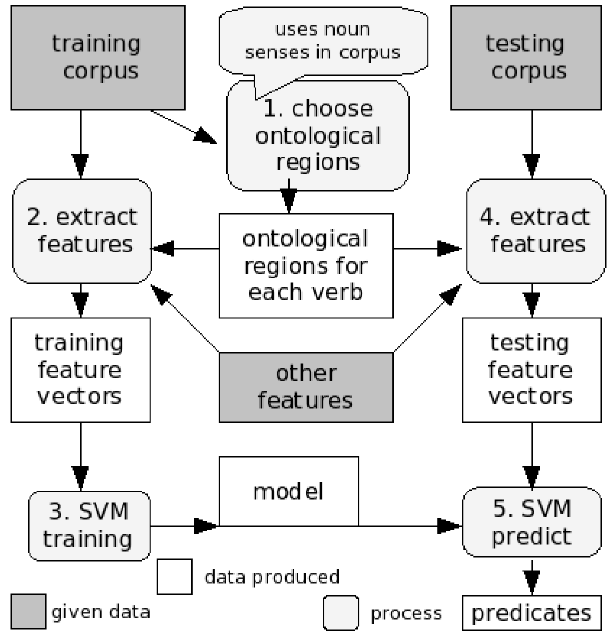

Figure 2: Steps in performing experiments: demonstrating how choosing *ontological regions* fits in.

Test:	**I**	**H**	**S**	**IH**	**IS**	**IHS**
Mean:	64.6	65	63.8	65.4	65.4	65.4

Table 2: Baseline results: results for predicate labeling with features in addition to *preposition* and *passive* features. I=*indicator*, H=*head noun*, S=*synonym*.

Verbal Predicate Labeling

We perform many experiments in order to record the effect of *ontological regions* in various situations for verbal predicate labeling. Although we alternate between many of the grammatical-based features, for all experiments we include the *preposition* and *passive* features described earlier. Unless otherwise indicated, we used a sigmoid kernel function for the SVM, finding results to be best with such a function. Results are reported based on the accuracy of labeling all instances of the predicates.

Baseline Results First, as a baseline for comparison, we provide results using features that are not new to semantic annotation systems. These include the *indicator*, *head noun*, and *synonym* features described earlier. The results are presented in Table 2. Note that the *head noun* and/or *synset* features are based on grammatical relations so they may supplement the information of the *indicator* feature when used on their own.

Ontological Region Results Next, Table 3 shows the effect of adding *ontological regions* in an ideal situation. The situation is ideal due to accurate grammatical relations being provided by the corpus, and the use of noun senses from the training corpus before the training stage. We see a drastic improvement over our baseline of *indicator*, *head*

Verb \ Test	IHS	O	IO	IHO	ISO	IHSO
abandon	80.0	100.0	93.3	86.7	93.3	93.3
absorb	55.0	77.5	72.5	72.5	77.5	72.5
accept	56.0	74.0	74.0	74.0	78.0	70.0
achieve	81.8	77.3	81.8	81.8	77.3	81.8
act	81.3	87.5	87.5	87.5	81.3	81.3
adapt	42.9	66.7	71.4	71.4	66.7	66.7
adopt	87.5	68.8	62.5	68.8	75.0	81.3
affect	66.7	73.3	73.3	66.7	73.3	73.3
appeal	58.8	94.1	94.1	94.1	94.1	88.2
approach	70.0	70.0	70.0	70.0	70.0	70.0
arrange	72.7	95.5	90.9	90.9	90.9	90.9
arrive	53.8	53.8	53.8	53.8	53.8	53.8
ask	66.7	83.3	83.3	91.7	91.7	91.7
assume	73.7	68.4	78.9	78.9	73.7	78.9
attack	40.0	66.7	73.3	66.7	73.3	66.7
beat	33.3	88.9	100.0	88.9	88.9	88.9
believe	100.0	94.7	94.7	94.7	94.7	94.7
blend	56.3	93.8	93.8	93.8	93.8	87.5
blow	39.1	52.2	56.5	56.5	56.5	52.2
bombard	100.0	100.0	100.0	100.0	100.0	100.0
bring	56.9	74.1	72.4	72.4	70.7	67.2
burn	84.2	89.5	84.2	84.2	84.2	78.9
drain	90.9	90.9	100.0	90.9	100.0	90.9
draw	50.0	75.0	78.1	78.1	71.9	75.0
drive	84.6	84.6	88.5	88.5	92.3	84.6
drop	63.2	89.5	89.5	89.5	89.5	89.5
empty	75.0	62.5	75.0	75.0	68.8	75.0
enjoy	69.2	61.5	69.2	61.5	61.5	76.9
enter	57.7	69.2	73.1	76.9	73.1	80.8
escape	50.0	75.0	75.0	75.0	83.3	75.0
establish	38.1	66.7	66.7	66.7	61.9	57.1
help	73.3	80.0	73.3	73.3	80.0	73.3
sell	93.8	81.3	81.3	81.3	81.3	87.5
Mean:	65.4	77.9	78.8	78.3	78.8	77.3

Table 3: Accuracy of predicate labeling when adding *ontological regions* (O) with different combinations of other grammar-based features (see Table 2 for the meaning of other symbols).

noun, and *synonym* features, increasing highest accuracy from 65.4% to 78.8%. Note that ontological regions by themselves are quite effective. The feature-space is significantly smaller when *head nouns* are not included as they require a significant number of words to be encoded. Additionally, we see that for some verbs, such as “appeal” and “beat”, the regions seem to make a drastic difference, while for others, such as “believe” or “sell”, they seem to just get in the way as unnecessary features.

Automatic Parser Results Table 4 shows the results when using the parser of Charniak (Charniak 2000) and an automatic method to choose *subject*, *direct object* and *indirect object* from the parse tree as was done for the syntactic features of (Dang & Palmer 2005). This experiment shows the effect of *ontological regions* in a situation without given grammatical relations. The feature maintains the same improvements although the overall accuracy is lower by roughly 5%. The idea that an accurate parsing results in

Test:	**AIH**	**AIO**	**AIHO**	**IG**	**IHG**	**AIHG**
Mean:	56.7	73.1	73.1	77.5	76.8	73

Table 4: Accuracy results for predicate labeling when using an automatic grammatical parse (A), and a general set of *ontological regions* (G) (see previous tables for the meaning of other symbols).

improvements seems to agree with other related work such as (Gildea & Palmer 2002; Pradhan *et al.* 2005).

General Ontological Regions Results We implemented a general set of *ontological regions* in order to simulate the use of a training corpus that did not have noun senses for the verbal arguments. Note that this used the hand selected approach to choosing representatives for regions. These results can also be found in Table 4. Given that the same general regions were selected in the ontology for every *(verb, gr)* pair, this approach did not require the noun senses provided by the corpus. Using a general static set of regions seems to reduce accuracy by around 1%, arguably an insignificant amount. We provide an additional test with these features and using an automatic parsing (*AIHG*) in order to show how one might use these features with another corpus. In this case, using the general set of *ontological regions* does not seem to hinder the classification accuracy compared to the use of an automatic parse.

Semantic Role Labeling

Finally, we test a semantic role labeler which uses predicates as a feature. As we described previously, the predicates of our corpus have semantic roles which follow according to the predicate’s arguments. Thus, the samples for this experiment came from clauses with correctly annotated predicates of the previous experiments. The same 10-fold cross-validation was performed, and we attempted to label roles which occurred in the training data along with a predicate and grammatical relations at least two times, resulting in a total of 471 attempted role instances. The features used for the SVM were the *passive*, *indicator*, *ontological regions*, as well as two new types of features: the *verbal predicate* of the clause is encoded as well as the *grammatical relation* of the argument itself. Alternatively, we could have created an SVM for each verbal predicate and grammatical relation pair, but this was not necessary since labeling errors seemed to come from a lack of training data rather than noise in the feature-space. Although the task may seem trivial for an SVM, the idea is that the SVM should be able to figure out this pattern with very few training instances if it is indeed true that the roles of arguments fall in place when the verbal predicates are correct.

Semantic Role Results We use the results from the *IO* experiment of predicate labeling in order to determine the *verbal predicate* feature. Without using *ontological regions* we received an accuracy of 98.5%. However, we noted that in a few instances the subject could be either an *AGENT* or an *INANIMATE-CAUSE* depending on the semantics of the subject itself. Thus, we included the *ontological region* features for just the subject and received an accuracy of 99.2%. As a final note, we should mention the same results were achieved with a linear kernel function, which has the advantage of less computational time compared with the sigmoid function.

Conclusion

We performed labeling of verbal predicates and semantic roles of a new corpus using support vector machines. Our goal was to explore the use of a semantic feature based on the WordNet ontology, namely *ontological regions*, and to verify the idea that semantic role labeling is a consequence of accurate verbal predicate labeling. From the results, it is clear that the *ontological regions* are an effective semantic feature, improving accuracy roughly 13% over the use of only syntactic and lexical types of features. Furthermore, we showed that once a predicate is chosen correctly, the semantic roles for its arguments follow around 99.2% of the time based on the grammatical relations.

Ontological regions represent a region of WordNet from which one searches for arguments of a verb to lie. A semantic classification system based on lexical features would require, ideally, the ability to represent every single noun that may occur. Our regions only require a limited portion of nouns (as a concept in WordNet) to be encoded, and thus with the removal of lexical features such as *head nouns*, the feature space is greatly reduced. Additionally, consider that there were only a small number of examples of each predicate on which the system was trained. This implies the *ontological region* features can provide learnable information with few examples. Lastly, although extra information in the corpus was helpful for the best results, the regions were still effective when some of the benefits of our corpus, such as the grammatical relations and noun senses of head nouns in training sets, were taken away.

We believe an important general implication can be drawn from our work. Under machine learning, perhaps predicate labeling and semantic role labeling should not be treated as distinct tasks. If one was to consider our annotation of verbal predicates and semantic roles as a single task, the accuracy would only reduce from 78.8% (of predicate labeling) to around 78.2% (of a combined predicate/semantic role labeling). Accordingly, future corpora should consider using the notion that a verb predicate is a verb sense with associated semantic roles.

References

Baker, C. F.; Fillmore, C. J.; and Lowe, J. B. 1998. The berkeley framenet project. In *Proceedings of the 17th International conference on Computational Linguistics*, 86–90. Morristown, NJ, USA: Association for Computational Linguistics.

Chang, C.-C., and Lin, C.-J. 2001. *LIBSVM: a library for support vector machines*. Software available at http://www.csie.ntu.edu.tw/ cjlin/libsvm.

Charniak, E. 2000. A maximum-entropy-inspired parser. In *Proceedings of the first conference on North American chapter of the Association for Computational Linguistics*, 132–139. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.

Collins, M. 1997. Three generative, lexicalized models for statistical parsing. In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, 16–23. Association for Computational Linguistics.

Collins, M. 1999. *Head-Driven Statistical Models for Natural Language Parsing*. Ph.D. Dissertation, University of Pennsylvania.

Dang, H. T., and Palmer, M. 2005. The role of semantic roles in disambiguating verb senses. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*.

Daudé, J.; Padró, L.; and Rigau, G. 2000. Mapping wordnets using structural information. In *Proceedings of the 38th Annual Meeting of the Association for Computational Linguistics*, 504–511.

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

Gildea, D., and Palmer, M. 2002. The necessity of syntactic parsing for predicate argument recognition. In *Proceedings of the 40th Annual Conference of the Association for Computational Linguistics*, 239–246.

Gomez, F. 2003. Automatic semantic annotation of texts. Technical report, University of Central Florida.

Gomez, F. 2004. Building verb predicates: A computational view. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, 359–366.

Grimshaw, J. 1990. *Argument Structure*. Cambridge, MA: MIT Press.

Hsu, C., and Lin, C. 2001. A comparison of methods for multi-class support vector machines. Technical report, National Taiwain University.

Kingsbury, P.; Palmer, M.; and Marcus, M. 2002. Adding semantic annotation to the penn treebank. In *Proceedings of the Human Language Technology Conference (HLT'02)*.

Lapata, M., and Brew, C. 2004. Verb class disambiguation using informative priors. *Computational Linguistics* 30(1).

Levin, B. 1993. *English Verb Classes and Alternations: A Priliminary Investigation*. Chicago: University of Chicago Press.

Merlo, P., and Stevenson, S. 2001. Automatic verb classification based on statistical distributions of argument structure. *Computational Linguistics* 27(3):373–408.

Metzler, J., and Sammons, M. 2005. Ner package 1.2. http://l2r.cs.uiuc.edu/ cogcomp/asoftware.php?skey=NE.

Miller, G.; Beckwith, R.; Fellbaum, C.; Gross, D.; and Miller, K. 1993. Five papers on wordnet. Technical report, Princeton University.

Pinker, S. 1989. *Learnability and Cognition: The acquisition of Argument Structure*. Cambridge, Mass.: MIT Press.

Pradhan, S.; Ward, W.; Hacioglu, K.; Martin, J.; and Jurafsky, D. 2005. Semantic role labeling using different syntactic views. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*.