

# Support Vector Learning for Semantic Argument Classification \*

Sameer Pradhan (spradhan@cslr.colorado.edu), Kadri Hacioglu (hacioglu@cslr.colorado.edu), Valerie Krugler<sup>†</sup>, (krugler@cslr.colorado.edu) Wayne Ward (whw@cslr.colorado.edu), James H. Martin (martin@cslr.colorado.edu) and Daniel Jurafsky<sup>‡</sup> (jurafsky@cslr.colorado.edu)

*The Center for Spoken Language Research, University of Colorado, Boulder, CO 80303*

**Abstract.** The natural language processing community has recently experienced a growth of interest in domain independent shallow semantic parsing – the process of assigning a WHO did WHAT to WHOM, WHEN, WHERE, WHY, HOW etc. structure to plain text. This process entails identifying groups of words in a sentence that represent these semantic arguments and assigning specific labels to them. It could play a key role in NLP tasks like Information Extraction, Question Answering and Summarization. We propose a new machine learning algorithm for semantic role parsing, extending the work of Gildea and Jurafsky (2002), Surdeanu et al. (2003) and others. Our algorithm is based on Support Vector Machines which we show give large improvement in performance over earlier classifiers. We show performance improvements through a number of new features designed to improve generalization to unseen data, such as automatic clustering of verbs. We also report on various analytic studies examining which features are most important, comparing our classifier to other machine learning algorithms in the literature, and testing its generalization to new test set from different genre. On the task of assigning semantic labels to the PropBank (Kingsbury et al., 2002) corpus, our final system has a precision of 84% and a recall of 75%, which are the best results currently reported for this task. Finally, we explore a completely different architecture which does not require a deep syntactic parse. We reformulate the task as a combined chunking and classification problem, thus allowing our algorithm to be applied to new languages or genres of text for which statistical syntactic parsers may not be available.

**Keywords:** Shallow Semantic Parsing, Support Vector Machines

## 1. Introduction

Automatic, accurate and wide-coverage techniques that can annotate naturally occurring text with semantic argument structure can facilitate the discovery of patterns of information in large text collections (Wallis and Nelson, 2001; Hearst, 1999). Shallow semantic parsing – the process of assigning a simple WHO did WHAT to WHOM, WHEN, WHERE, WHY, HOW, etc. structure to sentences in text, is the process of producing such a markup.

---

\* This research was partially supported by the ARDA AQUAINT program via contract OCG4423B and by the NSF via grant IIS-9978025

<sup>†</sup> Currently at Stanford University, Stanford, CA, 94305 vkrugler@stanford.edu

<sup>‡</sup> Currently at Stanford University, Stanford, CA, 94305 jurafsky@stanford.edu



More specifically, when presented with a sentence, a parser should, for each predicate in the sentence, identify and label the predicate's semantic arguments. This process entails identifying groups of words in a sentence that represent these semantic arguments and assigning specific labels to them. This notion of shallow semantic parsing, or case role analysis, has a long history in the computational linguistics literature. For a short review of this history, the reader is referred to Jurafsky and Martin (2000)

In recent work, a number of researchers have cast this problem as a tagging, problem and have applied various supervised machine learning techniques to it (Gildea and Jurafsky, 2000; Blaheta and Charniak, 2000; Gildea and Jurafsky, 2002; Gildea and Palmer, 2002; Gildea and Hockenmaier, 2003; Surdeanu et al., 2003; Chen and Rambow, 2003; Fleischman and Hovy, 2003; Hacioglu and Ward, 2003; Thompson et al., 2003; Pradhan et al., 2003). In this paper, we report on a series of experiments exploring this approach.

For the initial experiments, we adopted the approach described by Gildea and Jurafsky (2002) and evaluated a series of modifications to improve its performance. In the experiments reported here we first replaced their statistical classification algorithm with one that uses Support Vector Machines and then added some more features to the existing feature set. All of the experiments we report were conducted on PropBank, an approximately 300k-word corpus annotated with semantic arguments of verbs. We evaluated the system on three main tasks: identifying words/phrases that represent semantic arguments, independently classifying words/phrases known to be semantic arguments into the specific categories, and the combined task of identifying the arguments and then assigning respective labels to them. We evaluated results using both hand-corrected syntactic parses, and actual parses from the Charniak parser.

A fundamental assumption in this architecture is the presence of a full syntactic parse, which provides the bulk of the features used by the machine learning algorithm. This limits its applicability to the language and genre of text for which statistical syntactic parsers are available, or can be trained. Therefore, we reformulated the task as a combined chunking and classification problem retaining the same classifier and replacing some features that depend on the hierarchical structure of the syntactic parse, with ones derived from a flat chunked representation.

Finally, to test the sensitivity of both sets of features to change in corpus, we manually tagged about 400 sentences from the AQUAINT corpus (LDC, 2002), which is a collection of text from the New York Times Inc., Associated Press Inc., and Xinhua News Service, compared to the Wall Street Journal corpus, on which the system is trained. Perhaps not surprisingly, there is a significant drop in performance on this new source of text.

Section 2 introduces the semantic annotation schema and corpora used for the task. In Section 3 we will discuss the semantic parsing system that makes

use of features extracted from a full syntactic parse, and which assumes that the semantic argument boundaries align with those of the syntactic parser. In Section 4 we will discuss in detail the system that does not depend on a full syntactic parse, by formulating the semantic parsing problem as a chunking and classification task.

## 2. Semantic Annotation and Corpora

Two corpora are available for developing and testing semantic argument annotation – FrameNet<sup>1</sup> (Baker et al., 1998) and PropBank<sup>2</sup> (Kingsbury et al., 2002). FrameNet uses predicate specific labels such as JUDGE and JUDGEE. PropBank uses predicate independent labels – ARG0, ARG1, etc. In this paper, we will be reporting on results using PropBank, a 300k-word corpus in which predicate argument relations are marked for almost all occurrences of non-copula verbs in the Wall Street Journal (WSJ) part of the Penn Tree-Bank (Marcus et al., 1994). The arguments of a verb are labeled sequentially from ARG0 to ARG5, where ARG0 is the PROTO-AGENT (usually the subject of a transitive verb) ARG1 is the PROTO-PATIENT (usually its direct object), etc. PropBank does attempt to treat semantically related verbs consistently. In addition to these “core arguments,” additional “adjunctive arguments,” referred to as ARGMs are also marked. Some examples are ARGM-LOC, for locatives, and ARGM-TMP, for temporals. We will refer to these as ARGMs. Table I shows the argument labels associated with the predicate *operate* in PropBank.

Table I. Argument labels associated with the predicate *operate* (sense: work) in the PropBank corpus.

Tag	Description
ARG0	Agent, operator
ARG1	Thing operated
ARG2	Explicit patient (thing operated on)
ARG3	Explicit argument
ARG4	Explicit instrument

Following is an example structure extracted from the PropBank corpus. The syntax tree representation along with the argument labels is shown in Figure 1.

<sup>1</sup> <http://www.icsi.berkeley.edu/~framenet/>

<sup>2</sup> <http://www.cis.upenn.edu/~ace/>

[ARG0 It] [predicate operates] [ARG1 stores] [ARGM-LOC mostly in Iowa and Nebraska].

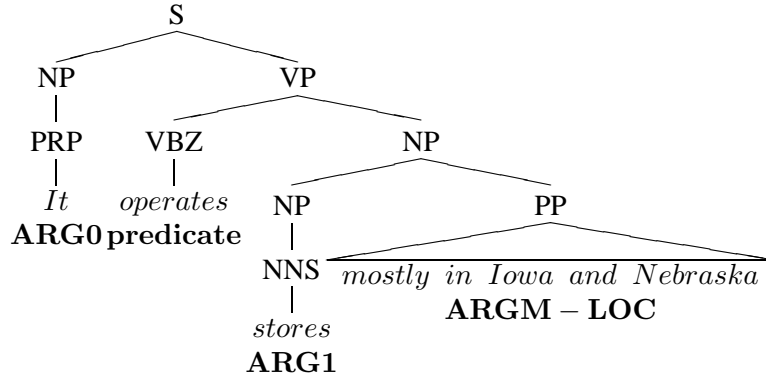


Figure 1. Syntax tree for a sentence illustrating the PropBank tags

Table II. List of adjunctive arguments in PropBank – ARGMS

Tag	Description	Examples
ARGM-LOC	Locative	<i>the museum, in Westborough, Mass.</i>
ARGM-TMP	Temporal	<i>now, by next summer</i>
ARGM-MNR	Manner	<i>heavily, clearly, at a rapid rate</i>
ARGM-DIR	Direction	<i>to market, to Bangkok</i>
ARGM-CAU	Cause	<i>In response to the ruling</i>
ARGM-DIS	Discourse	<i>for example, in part, Similarly</i>
ARGM-EXT	Extent	<i>at \$38.375, 50 points</i>
ARGM-PRP	Purpose	<i>to pay for the plant</i>
ARGM-NEG	Negation	<i>not, n't</i>
ARGM-MOD	Modal	<i>can, might, should, will</i>
ARGM-REC	Reciprocals	<i>each other</i>
ARGM-PRD	Secondary Predication	<i>to become a teacher</i>
ARGM	Bare ArgM	<i>with a police escort</i>
ARGM-ADV	Adverbials	<i>(none of the above)</i>

The adjunctive arguments are shown in Table II. Thus for example, ARG0 usually corresponds to a semantic agent, while ARG1 typically corresponds to the argument most affected by the action, usually referred to as the theme.

The PropBank annotation scheme assumes that a semantic argument of a predicate aligns with one or more nodes in the “hand-corrected” (hence TreeBank hand-corrected) parses. Although, most frequently the arguments are identified by one node in the tree, there can be cases where the arguments are discontinuous and more than one nodes are required to identify parts of the arguments. Sometimes these nodes can be \*trace\* nodes which refer to another node in the tree, but do not have any words associated with them. The following examples can help clarify these two notions.

#### EXAMPLE

PropBank also contains arguments that are coreferential. Since traces are not reproduced by a syntactic parser, we decided not to consider them in our experiments – whether or not they represent arguments of a predicate.

Most of the experiments in this paper are performed on the July 2002 release of PropBank. For these set of experiments, we treated discontinuous arguments as two instances of the same argument. We also do not differentiate between coreferential and non-coreferential arguments. Recently, a newer, larger, completely adjudicated Feb 2004 version of PropBank was released. We also report our best system performance on this dataset. In this dataset, we treat discontinuous and coreferential arguments in accordance to the CoNLL 2004 shared task on semantic role labeling. The first part of a discontinuous argument is labeled as it is, while the second part of the argument is labeled with a prefix “C-” appended to it. All coreferential arguments are labeled with a prefix “R-” appended to them.

We follow the standard convention of using Section 02 to Section 21 as the training set, Section-00 as the dev-test set and Section-23 as the test set. For the July 2002 release of PropBank, the training set comprises about 51,000 sentences, instantiating about 132,000 arguments, and the test set comprises 2,700 sentences instantiating about 7,000 arguments. The Feb 2004 release training set comprises about 85,000 sentences instantiating about 250,000 arguments and the test set comprises 5,000 sentences instantiating about 12,000 arguments.

### 3. Constituent-by-Constituent Semantic Parsing

#### 3.1. PROBLEM DESCRIPTION

The problem of shallow semantic parsing can be viewed as a two step process. First, is the process of identifying the predicate whose arguments we plan to identify. In our case, this is any non-copula verb in the sentence. The second step is to identify and classify words or phrases that represent the semantic arguments of that predicate. This process can be evaluated on three distinct tasks:

- *Argument Identification* – This is the process of identifying parse constituents in the sentence that represent valid semantic arguments of a given predicate.
- *Argument Classification* – Given constituents known to represent valid arguments of a predicate, assign the appropriate argument labels to them.
- *Argument Identification and Classification* – A combination of the above two tasks, where the system first identifies probable constituents that represent arguments of a predicate, and then assigns them the most likely argument labels.

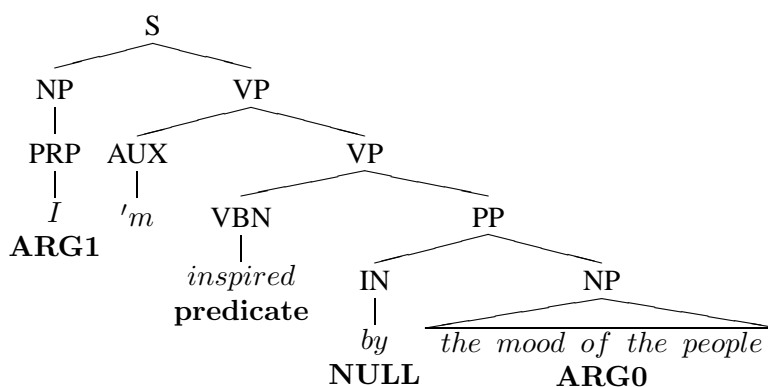


Figure 2. A sample sentence from the PropBank corpus

Each node in the parse tree can be classified as either one that represents a semantic argument (i.e., a NON-NULL node) or one that does not represent any semantic argument (i.e., a NULL node). The NON-NULL nodes can then be further classified into the set of argument labels. For example, in the tree of Figure 2, the PP that encompasses “*by the mood of the people*” is a NULL node because it does not correspond to a semantic argument. The node NP that encompasses “*the mood of the people*” is a NON-NULL node, since it does correspond to a semantic argument – ARG0.

### 3.2. BASELINE FEATURES

Our baseline system uses the same set of features introduced by Gildea and Jurafsky (2002): predicate, path, phrase type, position, voice, head word, and verb sub-categorization. All these features (except one – the predicate) are extracted from the syntactic parse tree of the sentence. Some of the features, viz., predicate, voice and verb sub-categorization are shared by all the nodes in the tree. All the others change with the constituent under consideration. For

evaluation purposes only predicates for which the arguments were annotated in PropBank were considered. In the actual system, all non-copula verbs are considered as predicates.

- **Predicate** – The predicate lemma is used as a feature.
- **Path** – The syntactic path through the parse tree from the parse constituent to the predicate being classified.

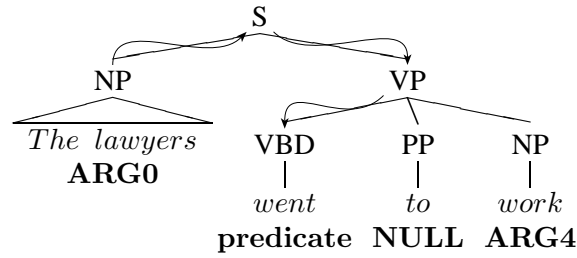


Figure 3. Illustration of path  $\text{NP}\uparrow\text{S}\downarrow\text{VP}\downarrow\text{VBD}$

For example, in Figure 3, the path from ARG0 – “The lawyers” to the predicate “went”, is represented with the string  $\text{NP}\uparrow\text{S}\downarrow\text{VP}\downarrow\text{VBD}$ .  $\uparrow$  and  $\downarrow$  represent upward and downward movement in the tree respectively.

- **Phrase Type** – This is the syntactic category (NP, PP, S, etc.) of the phrase corresponding to the semantic argument.
- **Position** – This is a binary feature identifying whether the phrase is before or after the predicate.
- **Voice** – Whether the predicate is realized as an active or passive construction. We run a set of hand-written `grep` expressions on the syntax tree to identify passive voiced predicates. In approximately 11% of the sentences in PropBank, the predicate had a passive instantiation.
- **Head Word** – The syntactic head of the phrase. This is calculated using a head word table described by Magerman (1994) and modified by Collins (1999) [Appendix. A]. This feature is not case sensitive.
- **Sub-categorization** – This is the phrase structure rule expanding the predicate’s parent node in the parse tree. For example, in Figure 3, the sub-categorization for the predicate “went” is  $\text{VP}\rightarrow\text{VBD-PP-NP}$ .

### 3.3. CLASSIFIER

Support Vector Machines (SVMs) (Vapnik, 1998; Burges, 1998) have been shown to perform well on text classification tasks, where data is represented in a high dimensional space using sparse feature vectors (Joachims, 1998; Kudo and Matsumoto, 2000; Lodhi et al., 2002). We formulate the parsing problem as a multi-class classification problem and use a Support Vector Machine (SVM) classifier (Hacioglu et al., 2003; Pradhan et al., 2003). However, SVMs are binary classifiers. There are two common approaches for extending SVMs to multi-class classification problems (Allwein et al., 2000). The first is known as the PAIRWISE approach, where a separate binary classifier is trained for each of the class pairs and their outputs are combined to predict the classes. This approach requires the training of  $\frac{N(N-1)}{2}$  binary classifiers. The second, known as the ONE vs ALL (OVA) approach, involves training  $n$  classifiers for a  $n$ -class problem. The classifiers are trained to discriminate between examples of each class, and those belonging to all other classes combined. During testing, the classifier scores on an example are combined to predict its class label.

While some experiments have been reported that the pairwise approach outperforms the OVA approach (Krebel, 1999), our initial experiments show better performance for OVA. Therefore, in this paper, all results are using the OVA classification approach.

### 3.4. SYSTEM IMPLEMENTATION

The system can be viewed as comprising two stages – the training stage and the testing stage. We will first discuss how the SVM is trained for this task. Since the training time taken by SVMs scales exponentially with the number of examples, and about 90% of the nodes in a syntactic tree have NULL argument labels, we found it efficient to divide the training process into two stages:

1. Filter out the nodes that have a very high probability of being NULL. A binary NULL vs NON-NULL classifier is trained on the entire dataset. A sigmoid function is fitted to the raw scores to convert the scores to probabilities as described by Platt (2000). All the training examples are run through this classifier and the respective scores for NULL and NON-NULL assignments are converted to probabilities using the sigmoid function. Nodes that are most likely NULL (probability > 0.90) are pruned from the training set. This reduces the number of NULL nodes by about 90%. This is accompanied by a very negligible (about 1%) pruning of nodes that are NON-NULL.
2. The remaining training data is used to train a OVA classifier that contains the all the classes along with the NULL class.



With this strategy only one classifier (NULL vs NON-NULL) has to be trained on all of the data. The remaining OVA classifiers are trained on the nodes passed by the filter (approximately 20% of the total), resulting in a considerable savings in training time.

In the testing stage, we do not perform any filtering of NULL nodes in a first pass. Instead, all the nodes are classified directly as NULL or one of the arguments using the classifier trained in step 2 above. We observe a slight decrease in recall if we filter the test examples using a NULL vs NON-NULL classifier in a first pass, as we do in the training process.

For our experiments, we used tinySVM<sup>3</sup> along with YamCha<sup>4</sup> (Kudo and Matsumoto, 2000; Kudo and Matsumoto, 2001) as the SVM training and test software. The system uses a polynomial kernel with degree 2; the cost per unit violation of the margin,  $C=1$ ; and, tolerance of the termination criterion,  $e=0.001$ .

### 3.5. BASELINE SYSTEM PERFORMANCE

Table III shows the baseline performance numbers on all the three tasks mentioned earlier in Section 3.1 using the PropBank corpus with “hand-corrected” parses, i.e., the parses from the hand-corrected Penn TreeBank<sup>5</sup>. The set features listed in Section 3.2 were used.

Table III. Baseline performance on all the three tasks using “hand-corrected” parses

PROPBANK ARGUMENTS	P	R	F <sub>1</sub>	A
Identification	91.2	89.6	90.4	
Classification	-	-	-	87.1
Identification + Classification	83.3	78.5	80.8	

For the argument identification and the combined identification and classification tasks, we report the precision (P), recall (R) and the  $F_{\beta}$ <sup>6</sup> scores, and for the argument classification task we report the classification accuracy (A). This test set and all test sets, unless noted otherwise are Section-23 of PropBank.

<sup>3</sup> <http://cl.aist-nara.ac.jp/~talus-Au/software/TinySVM/>

<sup>4</sup> <http://cl.aist-nara.ac.jp/~taku-Au/software/yamcha/>

<sup>5</sup> We report on hand-corrected parses because it is easier to compare with results from the literature; we will also report on performance using an actual errorful parser.

<sup>6</sup>  $F_1 = \frac{2 \times P \times R}{P + R}$

### 3.6. SYSTEM IMPROVEMENTS

#### 3.6.1. *Disallowing Overlaps*

The system as described above might label two constituents even if they overlap in words. This is a problem since we would like to have each word to belong to at most one semantic argument as constrained by the PropBank labeling process. We propose to choose among the overlapping constituents by retaining the one for which the SVM has the highest confidence, and labeling the others NULL. Since we are dealing with strictly hierarchical trees, nodes overlapping in words always have an ancestor-descendent relationship, and therefore the overlaps are restricted to subsumptions only. The probabilities obtained by applying the sigmoid function to the raw SVM scores are used as the measure of confidence. Table IV shows the performance improvement on the task of identifying and labeling semantic arguments using the “hand-corrected” parses. In this system, the overlap-removal decisions are taken independently of each other. For this experiment and all other experiments where we report on performance improvement, we conduct a  $\chi^2$  test of significance to determine whether the difference in number of responses over all the confusion categories (correct, wrong, false positive and false negative) are statistically significant at  $p = 0.05$ . All the statistically significant improvements are marked with an \*.

Table IV. Improvements on the task of argument identification and classification upon disallowing overlapping constituents using hand

	P	R	F <sub>1</sub>
Baseline	83.3	78.5	80.8
No Overlaps	85.4	78.1	*81.6

#### 3.6.2. *New Features*

We tested several new features, two of which were obtained from the literature – named entities in constituents and head word part of speech. The others are novel features.

##### *Named Entities in Constituents*

Surdeanu et al. (2003) reported a performance improvement on classifying the semantic role of the constituent by relying on the presence of a named entity in the constituent. We expect that some of these entities such as location and time are particularly important for the adjunctive arguments ARGM-LOC and ARGM-TMP. We evaluated a similar use of named entities in our system. We tagged 7 named entities (PERSON, ORGANIZATION, LOCATION,

PERCENT, MONEY, TIME, DATE) using Identifinder (Bikel et al., 1999) and added them as 7 binary features.

#### *Head Word Part of Speech*

Surdeanu et al. (2003) showed that adding the part of speech of the head word of a constituent as a feature in the task of argument identification gave a significant performance boost to their decision tree based system. We experimented with this feature to our system. improvement in performance on the task

#### *Verb Clustering*

The predicate is one of the most salient features in predicting the argument class. Since our training data is relatively limited, any real world test set will contain predicates that have not been seen in training. In these cases, we can benefit from some information about the predicate by using predicate cluster as a feature. The distance function used for clustering is based on the intuition that verbs with similar semantics will tend to have similar direct objects. For example, verbs such as “eat”, “devour”, “savor”, will tend to all occur with direct objects describing food. The clustering algorithm uses a database of verb-direct-object relations extracted by Lin (1998). The verbs were clustered into 64 classes using the probabilistic co-occurrence model of Hofmann Hofmann and Puzicha (1998). We then use the verb class of the current predicate as a feature.

#### *Generalizing the Path Feature*

As we will see in Section 3.11, for the argument identification task, path is one of the most salient features. However, it is also the most data sparse feature. To overcome this problem, we tried generalizing the path in three different ways:

1. Compressing sequences of identical labels into one following the intuition that successive embedding of the of the same phrase in the tree might not add additional information.
2. Removing the direction in the path, thus making insignificant the point at which it changes direction in the tree, and
3. Using only that part of the path from the constituent to the lowest common ancestor of the predicate and the constituent – “Partial Path”. For example, the partial path for the path illustrated in Figure 3 is NP↑S.

*Verb Sense Information*

The arguments that a predicate can take depend on the word sense of the predicate. Each predicate tagged in the PropBank corpus is assigned a separate set of arguments depending on the sense in which it is used. Table V illustrates the argument sets for the predicate *talk*. Depending on the sense of the predicate *talk*, either ARG1 or ARG2 can identify the *hearer*. Absence of this information can be potentially confusing to the learning mechanism.

Table V. Argument labels associated with the two senses of predicate *talk* in PropBank corpus.

Talk		sense 1: speak		sense 2: persuade/dissuade	
	Tag	Description	Tag	Description	
	ARG0	Talker	ARG0	Talker	
	ARG1	Subject	ARG1	Talked to	
	ARG2	Hearer	ARG2	Secondary action	

We added the oracle sense information extracted from PropBank, to our features by treating each sense of a predicate as a distinct predicate.

*Head Word of Prepositional Phrases*

Many adjunctive arguments, such as temporals and locatives, occur as prepositional phrases in a sentence, and it is often the case that the head words of those phrases, which are always prepositions, are not very discriminative, eg., “in the city”, “in a few minutes”, both share the same head word “in” and neither contain a named entity, but the former is ARG<sub>M</sub>-LOC, whereas the latter is ARG<sub>M</sub>-TMP. Therefore, we tried replacing the head word of a prepositional phrase, with that of the first noun phrase inside the prepositional phrase. We retained the preposition information by appending it to the phrase type, eg., “PP-in” instead of “PP”.

*First and Last Word/POS in Constituent*

Some arguments tend to contain discriminative first and last words so we tried using them along with their part of speech as four new features.

*Ordinal constituent position*

In order to avoid false positives of the type where constituents far away from the predicate are spuriously identified as arguments, we added this feature which is a concatenation of the constituent type and its ordinal position from the predicate.

*Constituent tree distance*

This is a finer way of specifying the already present position feature.

*Constituent relative features*

These are nine features representing the phrase type, head word and head word part of speech of the parent, and left and right siblings of the constituent in focus. These were added on the intuition that encoding the tree context this way might add robustness and improve generalization.

*Temporal cue words*

There are several temporal cue words that are not captured by the named entity tagger and were considered for addition as a binary feature indicating their presence.

*Dynamic class context*

In the task of argument classification, these are dynamic features that represent the hypotheses of at most previous two nodes belonging to the same tree as the node being classified.

*3.6.3. Alternative Pruning Strategies*

In the baseline system architecture, we opted for the two step training strategy mentioned in Section 3.4. To evaluate possible performance degradation accompanying the savings in training time, we tested two other pruning strategies. The three strategies in order of increased training time are:

- **Two-pass hard-prune** strategy, which uses a NULL *vs* NON-NULL classifier trained on the entire data in a first pass. All nodes labeled NULL are filtered. And then, ONE *vs* ALL classifiers are trained on the data containing *only* NON-NULL examples. There is no NULL *vs* ALL classifier in the second stage.
- **Two-pass soft-prune** strategy (baseline), uses a NULL *vs* NON-NULL classifier trained on the entire data, filters out nodes with high confidence (probability > 0.90) of being NULL in a first pass and then trains ONE *vs* ALL classifiers on the remaining data *including* the NULL class.
- **One-pass no-prune** strategy that uses all the data to train ONE *vs* ALL classifiers – one for each argument including the NULL argument. This has considerably higher training time as compared to the other two.

Table VI shows performance on the task of identification and classification of PropBank arguments. These performance numbers are computed over the reference examples, and not just the ones that pass the filter. There is no statistically significant difference between the *two-pass soft-prune* strategy,

and the *one-pass no-prune* strategy ( $\chi^2; p = 0.08$ ). However, both are better than the *two-pass hard-prune* strategy. Our initial choice of training strategy was dictated by the following efficiency considerations: i) SVM training is a convex optimization problem that scales exponentially with the size of the training set, ii) On an average about 90% of the nodes in a tree are NULL arguments, and iii) We have to optimize only one classifier on the entire data. We continued to use the *two-pass soft-prune* strategy

Table VI. Comparing pruning strategies on the task of argument identification and classification using hand-corrected treebank parses.

	No Overlaps		
	P	R	F <sub>1</sub>
Two-pass hard-prune	84	80	*81.9
Two-pass soft-prune	86	81	83.4
One-pass no-prune	87	80	83.3

### 3.7. ARGUMENT SEQUENCE INFORMATION

In order to improve the performance of their statistical argument tagger, G&J used the fact that a predicate is likely to instantiate a certain set of arguments. We use a similar strategy, with some additional constraints: i) argument ordering information is retained, and ii) the predicate is considered as an argument and is part of the sequence. We achieve this by training a trigram language model on the argument sequences, so unlike G&J, we can also estimate the probability of argument sets not seen in the training data. We first convert the raw SVM scores to probabilities using a sigmoid function. Then, for each sentence being parsed, we generate an argument lattice using the  $n$ -best hypotheses for each node in the syntax tree. We then perform a Viterbi search through the lattice using the probabilities assigned by the sigmoid as the observation probabilities, along with the language model probabilities, to find the maximum likelihood path through the lattice, such that each node is either assigned a value belonging to the PROPBANK ARGUMENTS, or NULL.

The search is constrained in such a way that no two NON-NULL nodes overlap with each other. To simplify the search, we allowed only NULL assignments to nodes having a NULL likelihood above a threshold. While training the language model, we can either use the actual predicate to estimate the transition probabilities in and out of the predicate, or we can perform a joint estimation over all the predicates. We implemented both cases considering two best hypotheses, which always includes a NULL (we add NULL to the list if it is not among the top two). On performing the search, we found that the

Table VII. Improvements on the task of argument identification and classification using hand-corrected treebank parses, after performing

CORE ARGS/ Hand-corrected parses	P	R	F <sub>1</sub>
Baseline w/o overlaps	90.5	87.4	88.9
Common predicate	91.2	86.9	89.0
Specific predicate lemma	91.7	87.2	*89.4

overall performance improvement was not much different than that obtained by resolving overlaps as mentioned earlier. However, we found that there was an improvement in the CORE ARGUMENT accuracy on the combined task of identifying and assigning semantic arguments, given hand-corrected parses, whereas the accuracy of the ADJUNCTIVE ARGUMENTS slightly deteriorated. This seems to be logical considering the fact that the ADJUNCTIVE ARGUMENTS are not linguistically constrained in any way as to their position in the sequence of arguments, or even the quantity. We therefore decided to use this strategy only for the CORE ARGUMENTS. Although, there was an increase in F<sub>1</sub> score when the language model probabilities were jointly estimated over all the predicates, this improvement is not statistically significant. However, estimating the same using specific predicate lemmas, showed a significant improvement in accuracy. The performance improvement is shown in Table VII.

### 3.8. BEST SYSTEM PERFORMANCE

It was found that a subset of the features helped argument identification and another subset helped argument classification. We will look at the individual feature contribution in Section 3.11. The best system is trained by first filtering the most likely nulls using the best NULL vs NON-NULL classifier trained using all the features whose argument identification F<sub>1</sub> score is marked in bold in Table XI, and then training a ONE vs ALL classifier using the data remaining after performing the filtering and using the features that contribute positively to the classification task – ones whose accuracies are marked in bold in Table XI. Table VIII shows the performance of this system.

### 3.9. USING AUTOMATIC PARSES

Thus far, we have reported results using hand-corrected “hand-corrected” parses. In real-world applications, the system will have to extract features from an automatically generated parse. To evaluate this scenario, we used

Table VIII. Best system performance on all tasks using hand-corrected parses.

Classes	Task	Hand-corrected parses			
		P	R	F <sub>1</sub>	A
ALL	Id.	95.2	92.5	93.8	
ARGS	Classification	-	-	-	91.0
	Id. + Classification	88.9	84.6	86.7	
CORE	Id.	96.2	93.0	94.6	
ARGS	Classification	-	-	-	93.9
	Id. + Classification	90.5	87.4	88.9	

the Charniak parser (Charniak, 2001) to generate parses for PropBank training and test data. We chose Charniak parser over Collins' parser for two reasons: i) At the time the source code only for Charniak parser was available and it could be modified to accept data from standard input required for the interactive parsing application that we were developing, and ii) Preliminary experiments indicated that the Charniak parser was faster than the Collins' parser. The predicate lemma was extracted using the XTAG morphology database<sup>7</sup> (Daniel et al., 1992).

Table IX. Performance degradation when using automatic parses instead of hand-corrected ones.

Classes	Task	Automatic parses			
		P	R	F <sub>1</sub>	A
ALL	Id.	89.3	82.9	86.0	
ARGS	Classification	-	-	-	90.0
	Id. + Classification	84.0	75.3	79.4	
CORE	Id.	92.0	83.3	87.4	
ARGS	Classification	-	-	-	90.5
	Id. + Classification	86.4	78.4	82.2	

In these generated parses, there are about 6% arguments whose boundaries do not align exactly with any of the hand-generated phrase boundaries. We tried to recover training information in such cases by automatically correcting the slight misalignment caused by a class of parser errors – as done before

<sup>7</sup> <ftp://ftp.cis.upenn.edu/pub/xtag/morph-1.5/morph-1.5.tar.gz>



by Gildea and Jurafsky (2002). In order to align the automatically generated constituents with the hand generated ones, words were dropped one by one from the right end of a hand-generated constituent until it matched one of the automatically generated constituents, with which it was then aligned. This procedure was followed for both training and test data. Table IX shows the performance degradation when automatically generated parses are used.

### 3.10. USING LATEST PROPBANK DATA

Owing to the Feb 2004 release of much more and completely adjudicated PropBank data, we have a chance to report our performance numbers on this data set. Table X shows the same information as in previous Tables VIII and IX, but generated using the new data. Owing to time limitations, we could not get the results on the argument identification task and the combined argument identification and classification task using automatic parses.

Table X. Best system performance on all tasks using hand-corrected parses using the latest PropBank data.

ALL ARGS	Task	P	R	F <sub>1</sub>	A
HAND	Id.	96.2	95.8	96.0	
	Classification	-	-	-	93.0
	Id. + Classification	89.9	89.0	89.4	
AUTOMATIC	Classification	-	-	-	90.1

### 3.11. SYSTEM ANALYSIS

#### *Feature Performance*

Table XI shows the effect each feature has on the argument classification and argument identification tasks, when added individually to the baseline. Addition of named entities to the NULL vs NON-NULL classifier degraded its performance. We attribute this to a combination of two things: i) there are a significant number of constituents that contain named entities, but are not arguments of a particular predicate (the parent of an argument node will also contain the same named entity) therefore this provides a noisy feature for NULL vs NON-NULL classification. We then tried using this as a feature solely in the task of classifying constituents known to represent arguments, using features extracted from “hand-corrected” parses. On this task, overall classification accuracy increased from 87.1% to \*88.1%. As expected, the most significant improvement was for adjunct arguments like temporals (ARGM-TMP) and locatives (ARGM-LOC) as shown in Table XII. Since the number of locative and temporal arguments in the test set is quite low (<

Table XI. Effect of each feature on the argument classification task and argument identification task when added to the baseline system

Features	Classification	ARGUMENT ID		
	Accuracy	P	R	F <sub>1</sub>
Baseline	87.9	93.7	88.9	91.3
+ Named entities	<b>88.1</b>	-	-	-
+ Head POS	<b>*88.6</b>	94.4	90.1	<b>*92.2</b>
+ Verb cluster	<b>88.1</b>	94.1	89.0	<b>91.5</b>
+ Partial path	<b>88.2</b>	93.3	88.9	91.1
+ Verb sense	<b>88.1</b>	93.7	89.5	<b>91.5</b>
+ Noun head PP (only POS)	<b>*88.6</b>	94.4	90.0	<b>*92.2</b>
+ Noun head PP (only head)	<b>*89.8</b>	94.0	89.4	<b>91.7</b>
+ Noun head PP (both)	<b>*89.9</b>	94.7	90.5	<b>*92.6</b>
+ First word in constituent	<b>*89.0</b>	94.4	91.1	<b>*92.7</b>
+ Last word in constituent	<b>*89.4</b>	93.8	89.4	<b>91.6</b>
+ First POS in constituent	<b>88.4</b>	94.4	90.6	<b>*92.5</b>
+ Last POS in constituent	<b>88.3</b>	93.6	89.1	91.3
+ Ordinal const. pos. concat.	87.7	93.7	89.2	<b>91.4</b>
+ Const. tree distance	<b>88.0</b>	93.7	89.5	<b>91.5</b>
+ Parent constituent	<b>87.9</b>	94.2	90.2	<b>*92.2</b>
+ Parent head	85.8	94.2	90.5	<b>*92.3</b>
+ Parent head POS	<b>*88.5</b>	94.3	90.3	<b>*92.3</b>
+ Right sibling constituent	<b>87.9</b>	94.0	89.9	<b>91.9</b>
+ Right sibling head	<b>87.9</b>	94.4	89.9	<b>*92.1</b>
+ Right sibling head POS	<b>88.1</b>	94.1	89.9	<b>92.0</b>
+ Left sibling constituent	<b>*88.6</b>	93.6	89.6	<b>91.6</b>
+ Left sibling head	86.9	93.9	86.1	89.9
+ Left sibling head POS	<b>*88.8</b>	93.5	89.3	<b>91.4</b>
+ Temporal cue words	<b>*88.6</b>	-	-	-
+ Dynamic class context	<b>88.4</b>	-	-	-

10%) as compared to the core arguments, the increase in performance on these does not boost the overall performance significantly.

Adding head word POS as a feature significantly improves both the argument classification and the argument identification tasks. We tried two other ways of generalizing the head word: i) adding the head word cluster as a feature, and ii) replacing the head word with a named entity if it belonged to any of the seven named entities mentioned earlier. Neither method showed significant improvement. The improvement on adding the verb cluster infor-

Table XII. Improvement in classification accuracies after adding named entity information.

	ARGM-LOC			ARGM-TMP		
	P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
Baseline	61.8	56.4	59.0	76.4	81.4	78.8
With Named Entities	70.7	67.0	*68.8	81.1	83.7	*82.4

mation was obtained despite of the fact that more than 99% of the predicates in the PropBank test set are present in the training set. Table XI also shows the contribution of replacing the head word and the head word POS separately in the feature where the head of a prepositional phrase is replaced by the head word of the noun phrase inside it. Apparently, a combination of relative features seem to have a significant improvement on either or both the classification and identification tasks, and so do the first and last words in the constituent.

#### *Feature Saliency*

In analyzing the performance of the system, it is useful to estimate the relative contribution of the various feature sets used. Table XIII shows the argument classification accuracies for combinations of features on the “hand-corrected” training and test set for all PropBank arguments.

Table XIII. Performance of various feature combinations on the task of argument classification.

Features	Accuracy
<i>All</i>	91.0
<i>All except Path</i>	90.8
<i>All except Phrase Type</i>	90.8
<i>All except HW and HW-POS</i>	90.7
<i>All except All Phrases</i>	*83.6
<i>All except Predicate</i>	*82.4
<i>All except HW and FW and LW-POS</i>	*75.1
<i>Only Path and Predicate</i>	74.4
<i>Only Path Phrase Type</i>	47.2
<i>Only Head Word</i>	37.7
<i>Only Path</i>	28.0

Table XIV. Performance of various feature combinations on the task of argument identification

Features	P	R	F <sub>1</sub>
<i>All</i>	95.2	92.5	93.8
<i>All except HW</i>	95.1	92.3	93.7
<i>All except Predicate</i>	94.5	91.9	93.2
<i>All except HW and FW and LW-POS</i>	91.8	88.5	90.1
<i>All except Path and Partial Path</i>	88.4	88.9	88.6
<i>All except Path and HW and Predicate Info.</i>			
<i>Only Path and HW</i>	88.5	84.3	86.3
<i>Only Path and Predicate</i>	89.3	81.2	85.1
<i>Only HW</i>	33.0	00.1	00.2

In the upper part of Table XIII we see the degradation in performance by leaving out one feature at a time. The features are arranged in the order of increasing salience. Removing all predicate related information has the most detrimental effect on the performance. The lower part of the table shows the performance of some feature combinations by themselves.

Table XIV shows the feature salience on the task of argument identification. As opposed to the argument classification task, where removing the path has the least effect on performance, on the task of argument identification, removing the path causes the convergence in SVM training to be prohibitively slow.

#### *Size of Training Data*

One important concern in any supervised learning method is the amount of training examples required for the near optimum performance of a classifier. To address this issue, we trained the classifier on varying amounts of training data. The resulting plots are shown in Figure 4. We approximate the curves by plotting the accuracies at six data sizes. The first curve from the top indicates the change in F<sub>1</sub> score on the task of argument identification alone. The third curve indicates the F<sub>1</sub> score on the task of argument identification and classification. Interestingly both the curves run almost parallel to each other indicating that there is constant loss due to classification errors throughout the data range. We also plotted the recall values for the combined task of identification and classification to show the fact that most of the trend in the F<sub>1</sub> score curve is contributed by the recall, whereas there is very little increase in precision with increasing data.

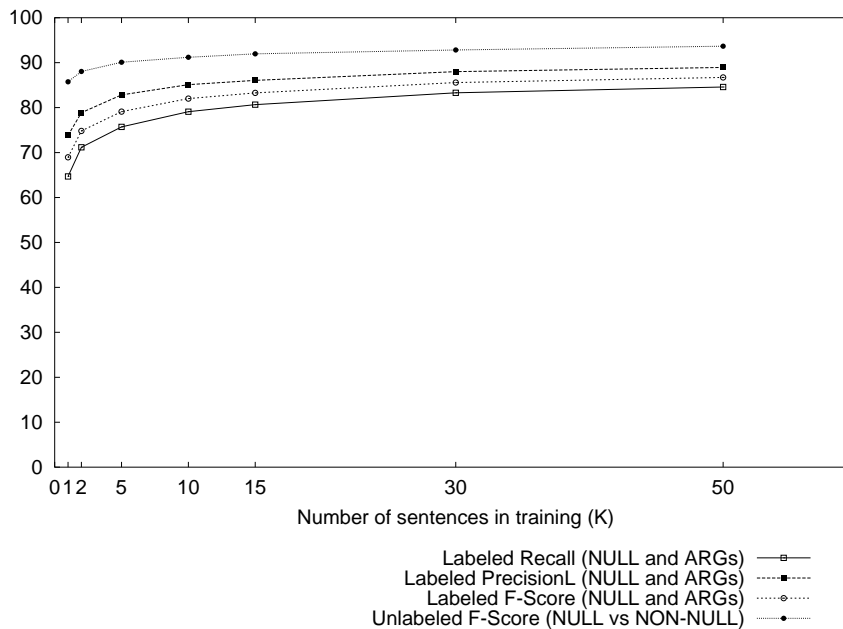


Figure 4. Learning Curve for the task of identifying and classifying arguments using hand-corrected parses.

### Performance Tuning

The expectations of a semantic parser by a real-world application would vary considerably. Some applications might favor precision over recall, and others vice-versa. Since we have a mechanism of assigning confidence to the hypothesis generated by the parser, we can tune its performance accordingly. Table XV shows the achievable increase in precision with some drop in recall for the task of identifying and classifying semantic arguments using the “hand-corrected” parses. The arguments that get assigned probability below the confidence threshold are assigned a null label.

### 3.12. COMPARISON WITH SIMILAR SYSTEMS

We evaluated our system in a number of ways. First, we compare it against 4 other shallow parsers in the literature. In comparing the systems, results are reported for the three types of tasks mentioned earlier. Not all the systems report performance on all these tasks, so we could only compare against tasks for which performance numbers are reported by the different systems. This results in some of the cells in the comparison tables being empty.

Table XV. Precision/Recall table for the combined task of argument identification and labeling using hand-corrected parses.

Confidence Threshold	P	R
0.10	88.9	84.6
0.25	89.0	84.6
0.50	91.7	81.4
0.75	96.7	65.2
0.90	98.7	32.3
0.95	98.5	12.3

#### *The Gildea and Palmer (G&P) System*

Gildea and Palmer (2002) use the same features used by Gildea and Jurafsky (2002), which are also our initial set of features. Two of the features – Head Word and Path – are found to be the most discriminative for argument identification. Maximum likelihood probabilities that the constituent is an argument, based on these two features –  $P(is\_argument|Path, Predicate)$  and  $P(is\_argument|Head, Predicate)$ , are interpolated to calculate the probability that the constituent under consideration represents an argument.

In the second step, the constituents that are retained are assigned a normalized probability that is calculated by interpolating distributions conditioned on various sets of features using a backoff lattice. The most probable argument is selected. They report results on the December 2001 release of PropBank.

#### *The Surdeanu et al. System*

Surdeanu et al. (2003) report results for a system that uses the same features as Gildea and Jurafsky (2002) (Surdeanu System I). They then report performance gains by adding additional features (Surdeanu System II) All the feature are combined with a decision tree classifier – C5 (Quinlan, 1986; Quinlan, ). Using the built-in boosting capabilities of this classifier show a slight improvement on their baseline performance. They use the July 2002 release of PropBank.

The additional features used by Surdeanu System II are:

- **Content Word** – This is a lexicalized feature that selects the most informative word from a phrase using some heuristics.
- **Part of Speech of the Content Word**
- **Named Entity class of the Content Word**

- **Boolean Named Entity Flags** – The value of these features are set to true or false depending on whether the corresponding named entity is present in the phrase.
- **Phrasal verb collocations** – This feature comprises frequency statistics related to the verb and the immediately following preposition.

*The Gildea and Hockenmaier (G&H) System*

This system is similar to the Gildea and Palmer (2002) system, but uses features extracted from a CCG grammar, which is a dependency grammar, instead of a phrase structure grammar. The features used are:

- **Path** – This was replaced with the concatenation of the category that the head word belongs to, the slot that it fills and an indicator of whether the word is a categorical functor. When the category information is unavailable, the path through the binary tree from the constituent to the predicate is used.
- **Phrase type** – This is the maximal projection of the PropBank argument’s head word in the CCG parse tree.
- **Voice**
- **Head Word**

Gildea and Hockenmaier (2003) report on both the core arguments and the adjunctive arguments on the November 2002 release of the PropBank, which contains approximately 20% more data, so we cannot directly compare to their system. However, since 80% of the data is shared by the two system, and we use the same test set, a comparison is valuable. We will refer to this as “G&H System I”.

*The Chen and Rambow (C&R) System*

Chen and Rambow (2003) also report results using decision tree classifier C4.5 (Quinlan, 1986). They report results using two different sets of features: i) Surface syntactic features much like the Gildea and Palmer (2002) system, ii) Additional features that result from the extraction of a Tree Adjoining Grammar (TAG) from the Penn TreeBank. They chose a Tree Adjoining Grammar because of its ability to address long distance dependencies in text. The additional features they introduced are:

- **Supertag Path** – This is the same as the feature “Path” that we saw earlier, but derived from a TAG rather than from a CFG.

- **Supertag** – This can be the tree-frame corresponding to the predicate or the argument.
- **Surface syntactic role** – This is the surface syntactic role of the argument.
- **Surface sub-categorization** – This is the surface sub-categorization frame.
- **Deep syntactic role** – This is the deep syntactic role of an argument.
- **Deep sub-categorization** – This is the deep syntactic sub-categorization frame.
- **Semantic sub-categorization** – This is the semantic sub-categorization frame.

We compare results to the system based on surface syntactic features (C&R System I), and the one that uses deep syntactic features (C&R System II). Unlike all the other systems we discuss, they also report results on core arguments (ARG0-5). Furthermore, owing to the absence of some semantic annotations in the Penn TreeBank, they could extract features for about 87% of the test set. Their numbers are not therefore directly comparable to ours, but since 87% of the test sets are overlapping, a rough comparison is still valuable.

### 3.12.1. *Comparing Classifiers*

Since two systems, in addition to ours, report results using the same set of features on the same data, we can directly assess the influence of the classifiers. Gildea and Palmer (2002) system estimates the posterior probabilities using several different feature sets and interpolate the estimates, while Surdeanu et al. (2003) use a decision tree classifier. Table XVI shows a comparison between the three systems for the task of argument classification. Keeping the same features and changing only the classifier produces a 8% absolute increase in performance on the same test set.

Table XVI. Argument classification accuracies using same features but different classifiers

System	Accuracy
SVM	87
Decision Tree (Surdeanu et al., 2003)	79
Lattice Backoff (Gildea and Palmer, 2002)	77



### 3.12.2. *Argument Identification*

Table XVII compares results on the argument identification task. As in the argument classification task, we see that the SVM system performs better than “Surdeanu System I” using the same set of features. It also outperforms the “Surdeanu System II”, which includes some more informative features.

Table XVII. Argument identification

Classes	System	Hand			Automatic		
		P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
ALL	SVM	95	92	94	89	83	86
ARGS	Surdeanu System II	-	-	89	-	-	-
	Surdeanu System I	85	84	85	-	-	-

### 3.12.3. *Argument Classification*

Table XVIII compares the argument classification accuracies of the systems.

Table XVIII. Argument classification

Classes	System	Hand	Automatic
		Accuracy	Accuracy
ALL	SVM	91	90
ARGS	G&P	77	74
	Surdeanu System II	84	-
	Surdeanu System I	79	-
CORE	SVM	93.9	90.5
ARGS	C&R System II	93.5	-
	C&R System I	92.4	-

### 3.12.4. *Argument Identification and Classification*

Table XIX shows the results for combined argument identification and classification task.

Table XIX. Identification and classification

Classes	System	Hand			Automatic		
		P	R	F <sub>1</sub>	P	R	F <sub>1</sub>
ALL	SVM	89	85	87	84	75	79
ARGs	G&H System I	76	68	72	71	63	67
	G&P	71	64	67	58	50	54
CORE	SVM System	90	87	89	86	78	82
ARGs	G&H System I	82	79	80	76	73	75
	C&R System II	-	-	-	65	75	70

## 4. Word-by-Word Semantic Parsing

### 4.1. PROBLEM DESCRIPTION

In the Constituent-by-Constituent (C-by-C) classification approach the candidate chunks are provided by the full syntactic parse of a sentence. So the segmentation process is performed separately from the classification task. In this method the classification is done on constituents. It is also possible to classify at the word-level. In (Ramshaw and Marcus, 1995), chunking was proposed as a tagging task. Here, each word in a sentence is labeled with a tag; I means that the word is inside a chunk, O means that the word is outside a chunk, and B means that the word is the beginning of a chunk. Using a variant of IOB representation, known as IOB2 (Sang and Veenstra, 1999), the words in the semantic chunks for the following example can be tagged as shown:

[ARG1I] 'm [predicateinspired] by [ARG0the mood of the people].

[B-ARG1 I] [O 'm] [O inspired] [O by] [B-ARG0 the][I-ARG0 mood] [I-ARG0 of] [I-ARG0 the] [I-ARG0 people]

In this scheme,  $N$  arguments would result in  $2 * N + 1$  classes. Experiments using this method have been reported in (Hacioglu and Ward, 2003; Hacioglu et al., 2003)

### 4.2. FEATURES

The features are very similar to the features used in the C-by-C system, except the features are derived for each word, not for each constituent. The features are

- **Predicate:** The verb in focus.
- **Part of Speech:** The part-of-speech category of the word.
- **Phrase Position:** The position of the word in a phrase based on IOB representation (e.g. B-NP, I-NP, O etc.)
- **Word Position:** The position of the word with respect to the predicate. It has two values as “before” and “after”.
- **Voice:** Indicates whether the sentence is in active or passive form with respect to the predicate.
- **Path:** This is same as the path feature used in the C-by-C system, except the path is relative to the word and not the constituent.

For each word to be tagged, a set of features is created from a fixed size context that surrounds the word. A 5-word sliding window centered at the current word defines the set of features. In addition to the above features, we also use semantic IOB tags that have already been assigned to the words that precede the word in focus and appear in the context. Figure 5 clarifies the notion of context and illustrates the features used. It should be noted that some of the features attached to the words are sentence-level features.

#### 4.3. SYSTEM IMPLEMENTATION

As in the C-by-C classification framework we have used SVMs in the ONE vs ALL setup. All SVM classifiers are realized using the TinySVM with the polynomial kernel of degree 2 and the general purpose SVM based chunker YamCha as distributed, without any changes. Figures 7 and 8 illustrate two implementations for feature extraction. As in the baseline system, the first system extracts features from a full syntactic parser. Even though classification is done on a word-by-word basis, features are extracted from a full parse. This system is referred to as W-by-W Chunker-1. In the second system, the full-syntactic parser is replaced by a part of speech tagger followed by a syntactic phrase chunker. The POS tagger and syntactic chunker were both implemented using Yamcha. They were trained on Sections 15 through 18 of the Penn Treebank. The POS tagger used the words, word prefixes/suffixes, word types (e.g numeric, capitalized etc.) and previous tag decisions as features. The syntactic chunker used the words and POS tags as features. In both systems, the features were created using a sliding window of size 5 centered at the word to be tagged. The second system is referred to as W-by-W Chunker-2. As mentioned earlier, the motivations for using a chunking parser instead of a full syntactic parser are: i) full parsing is computationally expensive than chunking, ii) chunkers can be developed more easily for new languages and

Word	POS	Phrase	Path	Predicate	Position	Voice	Class
I	PRP	B-NP	PRP<-NP<-S->VP->VBN	inspire	before	ACT	B-ARG0
'm	AUX	O	AUX<-VP->VP->VBN	inspire	before	ACT	O
inspired	VBN	B-VP	O	inspire	-	ACT	O
by	IN	B-PP	VBN<-VP->PP->IN	inspire	after	ACT	?
the	DT	B-NP	VBN<-VP->PP->NP->NP->DT	inspire	after	ACT	current prediction
mood	NN	I-NP	VBN<-VP->PP->NP->NP->NN	inspire	after	ACT	features
of	IN	B-PP	VBN<-VP->PP->NP->PP->IN	inspire	after	ACT	
the	DT	B-NP	VBN<-VP->PP->NP->PP->NP->DT	inspire	after	ACT	
people	NNS	I-NP	VBN<-VP->PP->NP->PP->NP->NNS	inspire	after	ACT	

Figure 5. Illustration of the context and the features used to classify a word – Chunker-1

Word	POS	Phrase	Path	Predicate	Position	Voice	Class
I	PRP	B-NP	PRP->NP->VP->VBN	inspire	before	ACT	B-ARG0
'm	VBP	B-VP	VPB->VP->VBN	inspire	before	ACT	O
inspired	VBN	I-VP	O	inspire	-	ACT	O
by	IN	B-PP	IN->PP->VP->VBN	inspire	after	ACT	?
the	DT	B-NP	DT->NP->PP->VP->VBN	inspire	after	ACT	current prediction
mood	NN	I-NP	NN->NP->PP->VP->VBN	inspire	after	ACT	features
of	IN	B-PP	IN->PP->NP->PP->VP->VBN	inspire	after	ACT	
the	DT	B-NP	DT->NP->PP->NP->PP->VP->VBN	inspire	after	ACT	
people	NNS	I-NP	NNS->NP->PP->NP->PP->VP->VBN	inspire	after	ACT	

Figure 6. Illustration of the context and the features used to classify a word – Chunker-2

genre of text. The features created for these two systems are shown in Figure 6 for the example exhibited in Figure 5. There are slight differences in POS and phrase position features and major differences in the path features. The path feature for W-by-W Chunker 2 is uni-directional (from the word in focus to the predicate) and created using the flat structure:

[NP (PRP I) ] [VP (VBP 'm) (VBN inspired) ] [PP (IN by) ] [NP (DT the) (NN mood) ] [PP (IN of) ] [NP (DT the) (NNS people) ]

The flat path is defined as a chain of phrase chunk labels terminated with the POS tags of the word in focus and predicate. Consecutive chunks with identical labels are collapsed into one. For example, the path from *people* to the predicate *inspired* is NNS→NP→PP→NP→PP→VP→VBN.

We are experimenting with both implementations to determine trade-offs between coverage, efficiency and accuracy. We present some preliminary results in the next section that compare the accuracy of Chunker-2 to Chunker-1, and also we compare their performance to the C-by-C parser described in Section 3.

#### 4.4. EXPERIMENTAL RESULTS

Since training SVM on the entire PropBank corpus takes a long time, these experiments were carried out by training both C-by-C and W-by-W systems on 10,000 randomly selected sentences from the training data. As always,

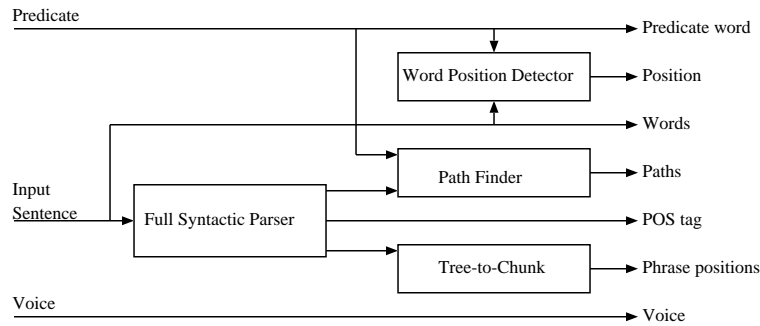


Figure 7. Chunker-1 architecture

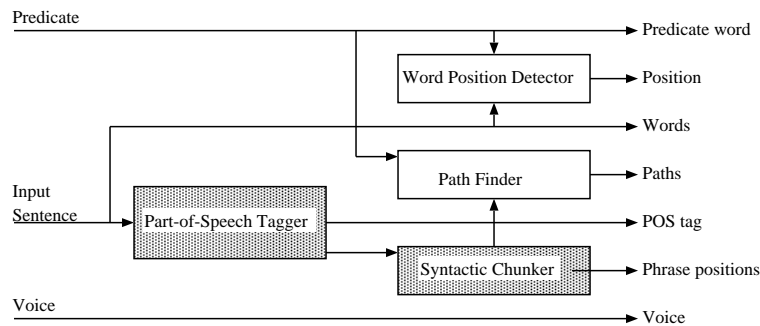


Figure 8. Chunker-2 architecture

Section-23 was used as the test set. We used the Charniak parser to obtain syntactic parse trees. The results are shown in Table XX.

The entire training set has approximately 1.5 million words. Training SVMs using such data is still prohibitive. Unfortunately, this kept us from carrying out extensive experiments with the W-by-W approach using the entire training set, additional features and other system settings.

Table XX. Comparison of C-by-C and W-by-W classifiers

System	P	R	F <sub>1</sub>
C-by-C	80.6	67.1	73.2
Chunker-1, W-by-W	70.7	60.5	65.2
Chunker-2, W-by-W	66.2	54.9	60.0

When features are derived from a flat chunked parse, instead of a syntactic parse tree, we observed a significant drop in performance. The major difference between these systems is the derivation of the path feature. To

Table XXI. Comparison W-by-W classifiers without the path feature

System	P	R	F <sub>1</sub>
Chunker-1, W-by-W	60.4	51.4	55.6
Chunker-2, W-by-W	59.8	50.8	54.9

further illustrate the effect, we ran both systems without using the path feature. These results are shown in Table XXI. As can be seen the performance of these systems is very similar, indicating the salience of the path feature. These experiments have also shown that the performance obtained by the W-by-W paradigm fell short of that obtained by the C-by-C paradigm. We think that an important step towards bridging this gap would be to adopt a two pass approach in the W-by-W paradigm, analogous in the C-by-C paradigm. That is, first chunk the sentence into arguments (NON-NULLS) and non-arguments(NULLS), and then label the NON-NULLS with the respective arguments. We are currently working along those directions to improve the performance.

## 5. Generalization to a New Domain

Thus far, all experiments have been tested on the same genre of data that they have been trained on. In order to see how well the features transfer to a new domain, we used both the W-by-W and the C-by-C classifiers trained on PropBank to test data drawn from the AQUAINT corpus (LDC, 2002). We annotated 400 sentences from the AQUAINT corpus with PropBank arguments. This is a collection of text from the New York Times Inc., Associated Press Inc., and Xinhua News Service (PropBank by comparison is drawn from Wall Street Journal). The results of the C-by-C classifier are shown in Table XXII.

There is a significant drop in the precision and recall numbers for the AQUAINT test set (compared to the precision and recall numbers for the PropBank test set which were 84% and 75% respectively). One possible reason for the drop in performance is relative coverage of the features on the two test sets. The head word, path and predicate features all have a large number of possible values and could contribute to lower coverage when moving from one domain to another. Also, being more specific, they might not transfer well across domains. All the other features are less likely to have been a source of the problem since they can take small enough values that the amount of training data is sufficient to estimate their statistics.

Table XXII. Performance on the AQUAINT test set.

Task		P	R	F <sub>1</sub>	A
ALL	Id.	75.8	71.4	73.5	-
ARGs	Classification	-	-	-	83.8
	Id. + Classification	65.2	61.5	63.3	-
CORE	Id.	88.4	74.4	80.8	-
ARGs	Classification	-	-	-	84.0
	Id. + Classification	75.2	63.3	68.7	-

Table XXIII. Feature Coverage on PropBank test set using parser trained on PropBank training set.

Features	Arguments	non-Arguments
<i>Predicate, Path</i>	87.6	2.9
<i>Predicate, Head Word</i>	48.9	26.5
<i>Cluster, Path</i>	96.3	5.0
<i>Cluster, Head Word</i>	83.8	60.1
<i>Path</i>	99.1	15.1
<i>Head Word</i>	93.0	90.6

Table XXIV. Coverage of features on AQUAINT test set using parser trained on PropBank training set.

Features	Arguments	non-Arguments
<i>Predicate, Path</i>	62.1	4.7
<i>Predicate, Head Word</i>	30.3	17.4
<i>Cluster, Path</i>	87.2	10.7
<i>Cluster, Head Word</i>	65.8	45.4
<i>Path</i>	96.5	29.3
<i>Head Word</i>	84.6	83.5

Table XXIII shows the coverage for features on the “hand-corrected” PropBank test set. The tables show feature coverage for constituents that were Arguments and constituents that were NULL. About 99% of the predicates in the AQUAINT test set were seen in the PropBank training set. Table XXIV shows coverage for the same features on the AQUAINT test set. We believe that the drop in coverage of the more predictive feature combinations explains

part of the drop in performance. We also ran the Word-by-Word chunking algorithms on the AQUAINT test set. The results are shown in Table XXV. This would compare to the Table XX. We also observed a degradation in the  $F_1$  score for this condition, from 65.2% to 52.2% for Chunker-1 and 60.0% to 45.7% for Chunker-2.

Table XXV. Identification and Classification performance of the chunking systems on the AQUAINT test set

Classes	System	P	R	$F_1$
ALL	Chunker-1, W-by-W	54.2	50.4	52.2
ARGUMENTS	Chunker-2, W-by-W	49.5	42.4	45.7

## 6. Conclusions

We have described an algorithm which significantly improves the state-of-the-art in shallow semantic parsing. Like previous work, our parser is based on a supervised machine learning approach. It achieves the best published performance on the task of identifying and labeling semantic arguments in PropBank. Key aspects of our results include significant (8% absolute) improvement via an SVM classifier, improvement from new features, and a series of analytic experiments on the contributions of the features.

For example, we found that path is the most salient feature for argument identification, but is difficult to generalize. Only the partial-path generalization seemed to have any benefit. For the task of classification, head word and predicate were the most salient features, but likewise, may be difficult to estimate because of data sparsity. Adding features that are generalizations of the more specific features also seemed to help. These features were named entities, head word part of speech and verb clusters.

We also report on a series of experiments using a completely different architecture which does not require a syntactic parser. This architecture performed significantly worse than our parser-based architecture, but may nonetheless be useful in new languages or genres in which statistical syntactic parsers are not available.

## Acknowledgements

We would like to thank Ralph Weischedel and Scott Miller of BBN Inc. for letting us use their named entity tagger – *IdentiFinder*; Martha Palmer for



providing us with the PropBank data and useful advice; Daniel Gildea for helpful discussions, and all the anonymous reviewers for their comments and suggestions.

## References

- Allwein, E. L., R. E. Schapire, and Y. Singer: 2000, ‘Reducing Multiclass to Binary: A Unifying Approach for Margin Classifiers’. In: *Proc. 17th International Conf. on Machine Learning*. pp. 9–16, Morgan Kaufmann, San Francisco, CA.
- Baker, C. F., C. J. Fillmore, and J. B. Lowe: 1998, ‘The Berkeley FrameNet project’. In: *COLING/ACL-98*. Montreal, pp. 86–90.
- Bikel, D. M., R. Schwartz, and R. M. Weischedel: 1999, ‘An Algorithm that Learns What’s in a Name’. *Machine Learning* **34**, 211–231.
- Blaheta, D. and E. Charniak: 2000, ‘Assigning Function Tags to Parsed Text’. In: *Proceedings of the 1st Annual Meeting of the North American Chapter of the ACL (NAACL)*. Seattle, Washington, pp. 234–240.
- Burges, C. J. C.: 1998, ‘A Tutorial on Support Vector Machines for Pattern Recognition’. *Data Mining and Knowledge Discovery* **2**(2), 121–167.
- Charniak, E.: 2001, ‘Immediate-Head Parsing for Language Models’. In: *Proceedings of the 39th Annual Conference of the Association for Computational Linguistics (ACL-01)*. Toulouse, France.
- Chen, J. and O. Rambow: 2003, ‘Use of Deep Linguistics Features for the Recognition and Labeling of Semantic Arguments’. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Sapporo, Japan.
- Collins, M. J.: 1999, ‘Head-driven Statistical Models for Natural Language Parsing’. Ph.D. thesis, University of Pennsylvania, Philadelphia.
- Daniel, K., Y. Schabes, M. Zaidel, and D. Egedi: 1992, ‘A freely available wide coverage morphological analyzer for English’. In: *Proceedings of the 14th International Conference on Computational Linguistics (COLING-92)*. Nantes, France.
- Fleischman, M. and E. Hovy: 2003, ‘A Maximum Entropy Approach to FrameNet Tagging’. In: *Proceedings of the Human Language Technology Conference*. Edmonton, Canada.
- Gildea, D. and J. Hockenmaier: 2003, ‘Identifying Semantic Roles Using Combinatory Categorical Grammar’. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Sapporo, Japan.
- Gildea, D. and D. Jurafsky: 2000, ‘Automatic Labeling of Semantic Roles’. In: *Proceedings of the 38th Annual Conference of the Association for Computational Linguistics (ACL-00)*. Hong Kong, pp. 512–520.
- Gildea, D. and D. Jurafsky: 2002, ‘Automatic Labeling of Semantic Roles’. *Computational Linguistics* **28**(3), 245–288.
- Gildea, D. and M. Palmer: 2002, ‘The Necessity of Syntactic Parsing for Predicate Argument Recognition’. In: *Proceedings of the 40th Annual Conference of the Association for Computational Linguistics (ACL-02)*. Philadelphia, PA.
- Hacioglu, K., S. Pradhan, W. Ward, J. Martin, and D. Jurafsky: 2003, ‘Shallow Semantic Parsing using Support Vector Machines’. Technical Report TR-CSLR-2003-1, Center for Spoken Language Research, Boulder, Colorado.
- Hacioglu, K. and W. Ward: 2003, ‘Target Word Detection and Semantic Role Chunking using Support Vector Machines’. In: *Proceedings of the Human Language Technology Conference*. Edmonton, Canada.

- Hearst, M.: 1999, 'Untangling Text Data Mining'. In: *Proceedings of the 37th Annual Meeting of the ACL*. College Park, Maryland, pp. 3–10.
- Hofmann, T. and J. Puzicha: 1998, 'Statistical Models For Co-occurrence Data'. Memo, Massachusetts Institute of Technology Artificial Intelligence Laboratory.
- Joachims, T.: 1998, 'Text Categorization with Support Vector Machines: Learning with Many Relevant Features'. In: *Proceedings of the European Conference on Machine Learning (ECML)*.
- Kingsbury, P., M. Palmer, and M. Marcus: 2002, 'Adding Semantic annotation to the Penn Treebank'. In: *Proceedings of the Human Language Technology Conference*. San Diego, CA.
- Krebel, U. H. G.: 1999, 'Pairwise Classification and Support Vector Machines'. In: *Advances in Kernel Methods*.
- Kudo, T. and Y. Matsumoto: 2000, 'Use of Support Vector Learning for Chunk Identification'. In: *Proceedings of the 4th Conference on CoNLL-2000 and LLL-2000*. pp. 142–144.
- Kudo, T. and Y. Matsumoto: 2001, 'Chunking with Support Vector Machines'. In: *Proceedings of the 2nd Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL-2001)*.
- LDC: 2002, 'The AQUAINT Corpus of English News Text, Catalog no. LDC2002T31'.
- Lin, D.: 1998, 'Automatic Retrieval and Clustering of Similar Words'. In: *Proceedings of the COLING-ACL*. Montreal, Canada.
- Lodhi, H., C. Saunders, J. Shawe-Taylor, N. Cristianini, and C. Watkins: 2002, 'Text Classification using String Kernels'. *Journal of Machine Learning Research* **2**(Feb), 419–444.
- Magerman, D.: 1994, 'Natural Language Parsing as Statistical Pattern Recognition'. Ph.D. thesis, Stanford University, CA.
- Marcus, M., G. Kim, M. A. Marcinkiewicz, R. MacIntyre, A. Bies, M. Ferguson, K. Katz, and B. Schasberger: 1994, 'The Penn treebank: Annotating predicate argument structure'.
- Platt, J.: 2000, 'Probabilities for Support Vector Machines'. In: A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans (eds.): *Advances in Large Margin Classifiers*. Cambridge, MA: MIT press.
- Pradhan, S., K. Hacioglu, W. Ward, J. Martin, and D. Jurafsky: 2003, 'Semantic Role Parsing: Adding Semantic Structure to Unstructured Text'. In: *Proceedings of the International Conference on Data Mining (ICDM 2003)*. Melbourne, Florida.
- Quinlan, J. R.: 1986, 'Induction of Decision Trees'. *Machine Learning* **1**(1), 81–106.
- Quinlan, R., 'Data Mining Tools See5 and C5.0, <http://www.rulequest.com/see-info.html>, 2003.
- Ramshaw, L. A. and M. P. Marcus: 1995, 'Text Chunking using transformation-based learning'. In: *Proceedings of the Third Annual Workshop on Very Large Corpora*. pp. 82–94.
- Sang, E. F. T. K. and J. Veenstra: 1999, 'Representing Text Chunks'. In: *Proc. of EACL*. pp. 173–179.
- Surdeanu, M., S. Harabagiu, J. Williams, and P. Aarseth: 2003, 'Using Predicate-Argument Structures for Information Extraction'. In: *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*. Sapporo, Japan.
- Thompson, C. A., R. Levy, and C. D. Manning: 2003, 'A Generative Model for Semantic Role Labeling'. In: *Proceedings of the European Conference on Machine Learning (ECML)*.
- Vapnik, V.: 1998, *Statistical Learning Theory*. New York: John Wiley and Sons Inc.
- Wallis, S. and G. Nelson: 2001, 'Knowledge Discovery in Grammatically Analysed Corpora'. *Data Mining and Knowledge Discovery* **5**(4), 305–335.