

VerbNet Class Assignment as a WSD Task

Susan Windisch Brown, Dmitriy Dligach, and Martha Palmer

Abstract The VerbNet lexical resource classifies English verbs based on semantic and syntactic regularities and has been used for numerous NLP tasks, most notably, semantic role labeling. Since, in addition to thematic roles, it also provides semantic predicates, it can serve as a foundation for further inferencing. Many verbs belong to multiple VerbNet classes, with each class membership corresponding roughly to a different sense of the verb. A VerbNet token classifier is essential for current applications using the resource and could provide the basis for a deep semantic parsing system, one that made full use of VerbNet's extensive syntactic and semantic information. We describe our VerbNet classifier, which uses rich syntactic and semantic features to label verb instances with their appropriate VerbNet class. It achieves an accuracy of 88.67 % with multiclass verbs, which is a 49 % error reduction over the most frequent class baseline.

1 Introduction

Rich verb representations are central to deep semantic parsing, requiring the identification of not only a verb's meaning but also how it connects the participants in the sentence. Disambiguating verbs using a lexicon that has already been enriched with syntactic and semantic information, rather than a more traditional lexicon, can bring end systems a step closer to accurate knowledge representation and reasoning. One such lexical resource, VerbNet, groups verbs into classes based on commonalities

S.W. Brown

Universita degli Studi di Firenze, Florence, Italy

e-mail: susanwbrown@att.net

D. Dligach

Boston Children's Hospital, Harvard Medical School, 300 Longwood Avenue, Boston, MA 02115, USA

e-mail: dmitriy.dligach@childrens.harvard.edu

M. Palmer (✉)

Department of Linguistics, University of Colorado at Boulder, Boulder, CO 80309-0295, USA

e-mail: martha.palmer@colorado.edu

in their semantic and syntactic behavior. It is widely used for a number of semantic processing tasks, including semantic role labeling (Swier and Stevenson 2004), the creation of conceptual graphs (Hensman and Dunnion 2004), and the creation of semantic parse trees (Shi and Mihalcea 2005). In addition, the detailed semantic predicates associated with each VerbNet class have the potential to contribute to text-specific semantic representations and, thereby, to inferencing tasks. However, application of VerbNet’s semantic and syntactic information to specific text requires first identifying the appropriate VerbNet class of each verb token. This is equivalent to a word sense disambiguation task.

Studies that have made use of VerbNet have dealt with the issue of multiclass verbs in different ways. When deciding on the class for a particular token of a verb in text, Zapirain et al. (2008) simply assigned the most frequent class for the verb rather than attempt to disambiguate. Their data consisted of any sentences in the Semlink corpus (Loper et al. 2007) in which the thematic roles mapped completely between PropBank and VerbNet, which resulted in a corpus that contained about 56 % of the original. For the data in their study, the most frequent class label was accurate 97 % of the time. Multiclass verbs throughout the entire Semlink corpus, however, have a most frequent class baseline of 73.8 %.

Other systems seem to have set aside the problem of multiclass verbs. For example, Bobrow et al. (2007) describe using VerbNet’s semantic predicates in the PARC’s question-answering system to derive pre- and post-conditions of events, such as the change of location of entities. For a verb like *leave*, the system attempts to use the semantic predicates provided by the VerbNet Leave-51.2 class:

```
MOTION(DURING(E), THEME)
LOCATION(START(E), THEME, SOURCE)
NOT(LOCATION(END(E), THEME, SOURCE))
DIRECTION(DURING(E), FROM, THEME, SOURCE)
```

to show that an entity was located in one place before the event and was in another location after the event. However, *leave* has multiple usages, not all of them involving physical change of location.

Table 1 shows its VerbNet classes and their semantic predicates. The PARC system would need to identify only those instances in their data where *leave* has the change of location meaning.

Zaenen et al. (2008) explain that the problem of automatically selecting only those instances that fit the desired class remains to be solved, especially in terms of dividing metaphorical from literal tokens of a verb: “We ignore the problem of metaphorical extensions for the relevant verbs. Resources other than VerbNet will need to be exploited to insure that these non-physical interpretations are excluded.” Although they do not state which ones are the relevant verbs, for many verbs this problem could be alleviated by disambiguating the class assignment for a specific verb instance. To continue our example, *leave* has six VerbNet classes: Escape, Fulfilling, Future_having, Keep, Leave and Resign. Only the Leave class and the Resign class have the START location and END location information they are looking for, and, for the Resign class, the CHANGE OF LOCATION is metaphorical. Therefore,

Table 1 VerbNet classes and semantic predicates for the verb *leave*

VerbNet class	Example	VerbNet semantics
Escape-51.1	The students left.	MOTION(DURING(E), THEME) DIRECTION(DURING(E), PREP_DIR, THEME)
Leave-51.2	Elvis has left the building.	MOTION(DURING(E), THEME) LOCATION(START(E), THEME, SOURCE) NOT(LOCATION(END(E), THEME, SOURCE)) DIRECTION(DURING(E), FROM, THEME, SOURCE)
Resign-10.11	He left Microsoft in 2008.	CAUSE(AGENT, E) LOCATION(START(E), SOURCE) NOT(LOCATION(END(E), SOURCE))
Fulfilling-13.1.4	He left the tenant with his business card.	HAS_POSSESSION(START(E), AGENT, THEME) HAS_POSSESSION(END(E), RECIPIENT, THEME) TRANSFER(DURING(E), THEME) CAUSE(AGENT, E)
Future_having-13.3	He left Sam his stamp collection.	HAS_POSSESSION(START(E), AGENT, THEME) FUTURE_POSSESSION(END(E), RECIPIENT, THEME) CAUSE(AGENT, E)
Keep-15.2	She left the papers in her desk.	PREP(DURING(E), THEME, LOCATION) CAUSE(AGENT, E)

the Leave class is the only class for this verb that suits their purposes. Classifying instances with the appropriate VerbNet class would enable them to apply the Location predicate to only those instances that are relevant. For the Semlink corpus, applying a most frequent class heuristic for *leave* would result in only 59 % accuracy. This is only one example of how an accurate, automatic VerbNet classifier would be useful.

2 Related Work

We know of only two previous efforts to create a VerbNet class disambiguator for verb tokens, those of Girju et al. (2005) and Abend et al. (2008). Girju et al. used a supervised machine learning methodology, with features from the words within three positions of the verb. These features included lemma, part of speech tag, phrase type from a syntactic chunker and named entity information. First, however, they faced the problem of creating a training set tagged with VerbNet class labels. They automatically constructed one by mapping from PropBank roleset labels to VerbNet classes, choosing to label only those verb instances in which the PropBank roleset mapped to only one VerbNet class. This methodology resulted in a set of target verbs in which 96 % belonged to only one VerbNet class. The high most-frequent-class

baseline of 96.5 % reflects the predominance of monosemous verbs and explains the low level of improvement over it: only 2 %. Because our classifier uses only multiclass verbs and a gold standard corpus with VerbNet class labels, it is not comparable to the Girju classifier.

The disambiguator developed by Abend et al. (2008) supports a much closer comparison. They also approach the task as a supervised machine learning problem, training and testing on the Semlink corpus. Polysemous verbs account for 58 % of their data, and they report results for all verbs and for just polysemous verbs. The Semlink corpus has annotated the verbs in the Wall Street Journal corpus with VerbNet classes. They selected instances that had been labeled with a VerbNet class, disregarding those verb instances that had been labeled as having no appropriate VerbNet class. Their system achieved 96.4 % accuracy, which was a 2.9 % increase over the 93.7 % baseline. The high baseline can also be attributed to the large number of monosemous verbs in their data. Considering only the polysemous verbs and the model using an automatic parser, the scenario most closely resembling our experimental setup, the most frequent class baseline was 88.6 % and the system accuracy was 91.9 %, which represents an error reduction of 28.95 %.

The results of the Abend et al. study suggest that automatic disambiguation of VerbNet classes is a reasonable line of research, and a possible method for verb sense disambiguation. The classifier relies on lexical and syntactic features, such as part of speech and heads of phrases. The classifier we describe is similar in several ways, although it adds several unique syntactic and semantic features and trains and tests only on multiclass verbs. The following sections will include comparisons of features and results where appropriate.

3 Method

To achieve verb token classification with VerbNet classes, we use a supervised machine learning approach. Using a corpus annotated with VerbNet class labels, we create a feature vector for each verb instance. A learning algorithm is then applied to generate a classifier. The following sections describe the data, the features and the experimental setup.

3.1 The Data

The training and test data are drawn from the Semlink corpus (Loper et al. 2007), which consists of the Penn Treebank portions of the Wall Street Journal corpus. A combination of automatic and manual techniques was used to label each verb instance with the appropriate VerbNet class. The resulting corpus is the largest repository of VerbNet token classification available. The corpus contains 113 K verb instances, 97 K of which are verbs represented in at least one VerbNet class (i.e., 86 %). Semlink includes 495 verbs that have instances labeled with more than one class (including verbs labeled with a single VerbNet class and None). We have

Table 2 Classifier features

Lexical	All open class words from target sentence and the surrounding sentences The two words preceding the target and their POS tags The two words following the target and their POS tags
Syntactic	The path through the parse tree from the target verb to its arguments Whether the target has a subj or obj and their head wds and POS Whether the target has a subordinate clause Whether the target has a PP adjunct The subcategorization frame The verb's voice (active or passive)
Semantic	Named-entity tags of the target's arguments WN hypernyms of the target's arguments WN synonyms of the target's arguments Dynamic Dependency Neighbors (DDNs)

trained and tested with all of these verbs that have 10 or more instances, resulting in a set of 344 verbs. The average number of classes for these verbs is 2.7, and the average number of instances was 133. All instances in the corpus for each verb were used, which created a dataset of 45,584 instances.

3.2 Features

We use a wide variety of features, including lexical, syntactic and semantic features, all derived automatically. Previous work has focused on lexical and syntactic features possibly because of the strong association of a VerbNet class to its syntactic alternations. However, a verb's membership in different classes also depends on its meaning, making the inclusion of semantic features a possible benefit. As mentioned earlier, multiple class memberships usually correlate with different senses of the verb, making VerbNet class disambiguation much like verb sense disambiguation. For this reason, we thought it was appropriate to treat the task as a verb sense disambiguation task. Some of the features are fairly standard ones used for general word sense disambiguation, but we have added some rich syntactic and semantic features that have proven useful for sense disambiguation of verbs. All features, which were previously also shown to be useful for WSD (Dligach and Palmer 2008) are summarized in Table 2 and explained more fully in the sections that follow.

3.2.1 Lexical Features

The lexical features include all open class words drawn from the target sentence and the sentence directly before and the sentence directly after it. In addition, we use a

feature that pairs each of the two words before and the two words after the target verb with their respective part-of-speech tag.

3.2.2 Syntactic Features

The syntactic features are drawn from syntactic parses automatically created with the Bikel Parser (Bikel et al. 1999). These features focus on the type of patterns that often distinguish one verb sense from another and that help delineate VerbNet classes. These include whether the target verb is in an active or passive form, whether it has a subject, an object, a subordinate clause, or a prepositional phrase adjunct. For each of these dependent items, the head word and its part of speech are included as features.

We also implement several unusual syntactic features that seem particularly well suited for VerbNet class disambiguation. The first is the path through the parse tree from the target verb to the verb's arguments, and the second is the sentence's subcategorization frame, as used in semantic role labeling. Because syntactic alternations, or patterns of subcategorization frames, play a large role in the organization of VerbNet classes, we expect these final two features to be particularly useful.

3.2.3 Semantic Features

Our use of semantic features is motivated by the work of Patrick Hanks (1996), who proposed that sense distinctions in verbs often rely on the membership of the verb's arguments in narrowly defined verb-specific semantic classes that he called lexical sets. A lexical set could consist, for example, of such nouns as *fist*, *finger*, *hand*, etc. (but not all body-parts); its members, when used as objects of *shake*, form instances of the communicative act sense of *shake*. This view corroborates our motivation that states the necessity of capturing the semantics of the verb's arguments and semantic similarities among them.

To illustrate with an example from our data, the verb *fix* falls into two VerbNet classes: (1) Preparing-26.3, (e.g., *He fixed lunch for the team; My mom fixed me a peanut butter and bacon sandwich*) and (2) Price-54.4, with the sense of "establish" (e.g., *They fixed the interest rate at 3 %; The lawyers fixed the terms of the agreement at their last meeting*). These two senses can be distinguished largely on the basis of the objects *lunch*, *sandwich*, *rate* and *terms*, the first two indicating the Preparing-26.3 class and the latter two indicating the Price-54.4 class. Not surprisingly, semantic features drawn from a target verb's arguments have been shown to improve verb sense disambiguation above and beyond lexical and syntactic features (Dligach and Palmer 2008).

Another study that reinforces a similar idea was reported by Federici et al. (1999). They describe their SENSE system that relies on inter-contextual analogies between tagged and untagged instances of a word to infer that word's sense. For example, if a verb's sense is preserved when used with two different objects, it is often possible

Table 3 Semantic features for one sense of the verb *fix*

Object	NE tag	WN synset	WN hypernyms	Sample DDNs
price	n/a	price, terms, damage	cost	raise, bring, increase, put, reduce, cut, have, offer, set
terms	n/a	price, terms, damage	cost	reduce, cut, have, offer, set
rate	n/a	charge per unit	cost	raise, bring, increase, put, reduce, cut, have, offer, set
number	n/a	figure	amount	raise, bring, increase, put, reduce, cut, have, offer, set

to conclude by analogy that the sense of another verb is also preserved when it is used with the same two objects.

In word sense disambiguation, the existing approaches to extracting semantic features are often based on obtaining lexical knowledge about the target verb's arguments from electronic dictionaries such as WordNet (Fellbaum 1998). WordNet synonyms and hypernyms are often used as semantic features (Dang and Palmer 2002; Dligach and Palmer 2008). Named entity tags, another source of lexical knowledge, can be obtained from the output of a named-entity tagger such as IdentiFinder (Bikel et al. 1999).

Four types of semantic features are used, all derived from the arguments of the target verb: (1) named entity tags for all of the arguments of the target verb, extracted using IdentiFinder; (2) synonyms of the arguments as listed in their synonym sets in WordNet; (3) hypernyms of the arguments, also taken from WordNet; and (4) dynamic dependency neighbors (DDNs) (Dligach and Palmer 2008), which connect objects of the verb based on the type of verbs they frequently occur with in object position. In this paper we utilized object-based DDNs to capture the semantics of the target verb's object. Elsewhere (citation below) we also experimented with subject-based DDNs in the context of verb sense disambiguation. We discovered that subject-based DDNs do not improve the performance over and above object-based DDNs. For these experiments the DDNs were calculated from the verbs' and objects' occurrence in the English Gigaword corpus, parsed with the dependency MaltParser (Nivre et al. 2007).

This last feature finds similarities between objects that can be missed by the other three, as can be seen in Table 3. The similarity in the first two objects, *price* and *terms*, is captured by the WordNet synset. The third object, *rate*, can be grouped with these via the WordNet hypernym. The fourth object, however, has none of these features in common with the others. Even moving up the WN hypernym hierarchy, *number* does not connect to the others until the very general category of Abstract Entity. However, objects with very different hypernyms or named entity tags may still be common objects of the same verbs. Objects grouped in this way can often help identify the particular sense of a verb (Dligach and Palmer 2008). Comparing lists of the top 50 verbs that each object occurs with shows a great deal of overlap and notably draws the noun number into a group with the other three.

3.3 *Experimental Setup*

Like all supervised word sense disambiguation, each verb required the training and testing of its own classifier. We classified using support vector machines (Chang and Lin 2001). Accuracy and error rates were computed with 5-fold cross validation. Baselines were established for each target verb type by calculating the accuracy that would be achieved if all instances of a verb were labeled with its most frequent VerbNet class. The average baseline for our verb set was 77.78 %.

4 Results

The average accuracy of the system with the target verbs was 88.67 %, which represents an error reduction of 49 % over the baseline of 77.78 %. The closest comparison to the Abend et al. classifier is to their results based on only polysemous verbs and using features drawn from an automatic parser. In this scenario, their classifier had an accuracy of 91.9 %, with an error reduction of 28.95 % over their baseline of 88.6 %.

In order to assess the contribution of the features we use to the performance of the classifier, we developed several different models composed of various combinations of our features. In addition we created a dedicated test set using 30 % of the Semlink corpus so that each model would be evaluated on identical training and test sets, assuring consistent comparisons. Using this test set, the overall performance of our classifier (the model with all features) was 84.64 %. This result is somewhat lower than the classifier accuracy using 5-fold cross-validation described above, possibly because of the smaller amount of training data used for this method. Compared to the most frequent class baseline, this figure still represents an error reduction of 31 %.

Lexical features are generally the most standard in supervised WSD systems and seem to contribute the most to the accuracy. Therefore, we used a model containing only the lexical features as our most stripped-down model. This model had an accuracy of 83.07 %. The second model added syntactic features to that, and achieved an accuracy of 84.44 %. Adding semantic features brought the accuracy to 84.65 %. We were particularly interested in assessing the contribution of the DDN feature, given that it can be generated automatically and requires no manually built lexical resource. For that reason, we also created a model with all the features but the DDN and a model with all the features but the non-DDN semantic features, which resulted in accuracies of 84.12 % and 84.89 % respectively, validating the efficacy of the DDN feature. See Table 4 for a summary of these results, along with error reduction figures.

5 Discussion

The accuracy of our VerbNet classifier approaches 90 %, the level that several researchers have indicated is needed for useful WSD (Sanderson 2000; Ide and Wilks

Table 4 Accuracy and error reduction of models using various features

Model	Baseline (%)	Accuracy (%)	Error reduction (%)
Lexical features only	77.78	83.07	23.81
Lexical + syntactic	77.78	84.44	29.97
Lexical + semantic	77.78	83.75	26.87
All but DDN	77.78	84.12	28.53
Lexical + syntactic + DDN	77.78	84.89	32.00
All features	77.78	84.65	30.92

2006). Using VerbNet classes as sense distinctions makes available sets of semantic predicates that can be used for deeper analysis. WSD is not an end in itself; it is only useful in so far as it improves more complex applications. By substituting VerbNet classification for verb sense disambiguation, we would gain both a coarse-grained sense of the verb and direct mappings to VerbNet's class-specific syntactic and semantic information. With the goal of improving future VerbNet classifiers, we discuss several pertinent issues in the following sections.

5.1 Contributions of the Features

The difference between the model with only lexical features and that with both lexical and syntactic features was statistically significant ($p = .0005$), suggesting that our syntactic features were a notable improvement to the model. Given the strong basis of VerbNet classes on syntactic alternations, we expected that syntactic features focused on argument structure would improve the system, and this comparison supports that hypothesis.

The semantic features showed a more complex pattern. A model with lexical and semantic features achieved an accuracy of 83.75%. Compared to the accuracy of the lexical-only model, this was a significant improvement ($p = .0182$), although less strongly so than the syntactic features. Interestingly, when the lexical+syntactic model (no semantic features) was compared to one with lexical, syntactic and semantic features, the difference in accuracy was not significant ($p = .6982$), suggesting that the small improvement we saw with the semantic features was only replicating some of the information the system was gaining from the syntactic features.

When the semantic features were tested separately, however, we found that the DDN feature substantially improved the system, while the other semantic features did not help the system. A model with all the features but the DDN feature showed no significant improvement over the lexical+syntactic model. This suggests that the named entity, WordNet synset, and WordNet hypernym features added nothing to the model. In a head-to-head comparison between the model with all features but the DDN and one with lexical, syntactic, and only DDNs, we found that the DDN

feature significantly improved the system ($p < .05$). With an error reduction of 32 %, the lexical + syntactic + DDN model performed the best of all those we tested.

These results suggest that the system could be streamlined by removing the named entity tag, WordNet synset, and WordNet hypernym features and leaving the DDNs as the only semantic features. This would reduce the system's dependence on other resources with no loss of accuracy. In addition, the DDN feature is created dynamically, and can be done with any corpus, increasing the portability of this system to new domains.

5.2 Semlink Annotation

A couple of matters came to light during a close examination of some of the Semlink annotation in our dataset. First, for some of the verbs, the mapping from PropBank to VerbNet that was the basis of the semiautomatic labeling inappropriately mapped some VerbNet classes. For example, the verb *fix* belongs to the Preparing class, which primarily describes events of food preparation. The thematic roles and semantic predicates for this class indicate the creation of some entity, such as *He fixed me a sandwich*. This class was used in the Semlink data to label such instances, but also to label instances of *fix* as a repair event, such as *We had to fix his car*, a usage that is currently not covered by any VerbNet class. Accuracy for this verb was still high at 89 %, possibly because the feature patterns were still consistent when these instances were labeled with the Preparing class.

The consequences of inappropriate labeling in this case are mixed. If thematic roles were assigned based on this label, they would likely still be correct. Both senses of *fix* call for an Agent and a Patient. The subject in *We had to fix his car* would be correctly labeled as an Agent and the object would be correctly labeled as a Patient. For semantic role labeling, this sort of error should have little negative effect. Any inferences based on the semantic predicates, however, would be misleading. In a Repair event, such as *We had to fix his car* no new entity is created, but the Preparing class label would incorrectly imply that the car is a newly created entity. It is not clear whether such inappropriate mapping is an isolated problem or not. In Sect. 7 we discuss some methods for assessing the existing annotation and for efficiently augmenting it.

5.3 Metaphorical Interpretations

A more common issue concerns the extension of VerbNet classes to metaphorical or figurative usages of a verb. Although some classes include metaphorical usages of the member verbs, such as the Amalgamate-22.2 class, others restrict the uses to literal events. For example, the Bump-18.4 class describes events of contact between a Theme and a Location, such as *The grocery cart hit the wall*. The class restricts

both the Theme and Location to [+concrete] arguments. A natural extension of this sense of *hit* would apply to abstract arguments and metaphorical events of contact, such as *The Bank of England was hit hard by the financial slump*. This usage of *hit* would not strictly fit the Bump-18.4 class because the financial slump (the Theme) is not a concrete entity and the Bank of England would not qualify as a concrete location, at least as it is used in this sentence. There is currently no VerbNet class, however, that would accommodate this usage of *hit*.

For several verbs in our set, including *hit* and *pay*, class labels were applied to metaphorical sense extensions. It is unclear whether this affected the accuracy of the classifier; for these two examples, the accuracy for *hit* was 75 %, whereas for *pay*, it was 97 %. More importantly, in terms of applying the labeled data to further semantic processing, metaphorical extensions should have little detrimental effect. Any thematic roles assigned based on the class label would be correct, although the semantic restrictions on the roles (e.g., +concrete) would not. The semantic predicates would also be correct, as long as they were interpreted metaphorically as well.

6 Conclusion

The VerbNet class disambiguator we present in this paper achieves 89 % accuracy with polysemous verbs, which is a 49 % error reduction over the most frequent class baseline. Given that most applications that currently use verb mappings to VerbNet classes rely on a most-frequent-class heuristic (or hand-selected data), this classifier should improve the functioning of these applications.

In addition, we have demonstrated that VerbNet class disambiguation often corresponds to coarse-grained verb sense disambiguation. However, unlike sense disambiguation with more traditional lexicons, VerbNet class disambiguation would not only help disambiguate the senses of verbs in context, it would automatically connect that context to detailed information about likely thematic roles, semantic representations, and related verbs. In combination with a syntactic parse of the sentence, knowing the appropriate VerbNet class could help select a semantic representation of the events in the sentence. By choosing VerbNet as a sense inventory, the next steps in complex knowledge representation and reasoning tasks could be facilitated.

7 Future Work

Some additional steps can be taken to improve the usefulness of VerbNet class labeling. The coverage of verbs and verb senses could be improved, both in the Semlink corpus and in VerbNet itself: 25 % of the verb tokens in the Semlink corpus have no VerbNet class label. However, Semlink is based on version 2.1 of VerbNet. The current version, 3.1, incorporates over 700 new verb senses, many of which introduce very common verbs, such as *seem*, *involve*, and *own*. Updating the corpus with

annotations for these new verbs and verb senses would improve coverage. A more long-term goal is to annotate data from other types of corpora than the WSJ, which would likely improve any VerbNet classifier's portability to new domains.

We plan to increase VerbNet annotation in the Semlink corpus using methods that take advantage of existing mappings between PropBank and VerbNet and efficient manual annotation (Dligach 2011). Semlink expansion can be accomplished in two ways. First, more data can be labeled using some form of active learning (Settles 2010) (e.g., batch mode uncertainty sampling). Once more annotated data has been acquired, it may be a good idea to double annotate all or parts of the data, leading to a more error-free labeled corpus. Various error detection techniques can be used to reduce the amount of the second round of annotation (Dligach 2011). These methods can also be used to judge the reliability of the semiautomatic annotation that has already been done, which should indicate how widespread mislabeling is (such as with the verb *fix*, see Sect. 5.2).

The question of metaphorical extensions in the VerbNet annotation is currently being addressed by the VerbNet team. Plans are underway to enhance VerbNet classes with metaphorical information, where appropriate. These enhancements will indicate any changes in thematic role restrictions with a metaphoric usage, and any changes necessary for a semantic predicate to be interpreted correctly.

Given the success of the DDN feature, we would like to see if expanding its contribution would further enhance our classifier. Currently, the DDN feature is only calculated for objects of the verb, but the feature could be encoded for the subject of the verb as well.

We see this classifier as an important step toward using VerbNet for deep semantic analysis. We have shown that verbs in multiple VerbNet classes can be disambiguated with close to 90 % accuracy. Another related task, semantic role labeling, has made great strides lately (Palmer et al. 2010). Using the output from both these tasks should enable us to identify the specific VerbNet frame and semantic predicate for the sentence. For example, VerbNet class disambiguation and semantic role labeling would identify the sentence *He left Sam his stamp collection* as

Agent V(class:Future-having-13.3) Recipient Theme.

Only one frame in the Futurehaving13.3 class has that pattern: the NP V NPdative NP frame. Its semantic predicates are

HAS_POSSESSION(START(E), AGENT, THEME)
 FUTURE_POSSESSION(END(E), RECIPIENT, THEME)
 CAUSE(AGENT, E).

Given the argument labels from the semantic role labeling, it is straightforward to map from the original sentence to the semantic representation:

HAS_POSSESSION(START(E), HE, THE STAMP COLLECTION)
 FUTURE_POSSESSION(END(E), SAM, THE STAMP COLLECTION)
 CAUSE(HE, E).

Recent work in coreference resolution (Haghighi and Klein 2009) and implicit argument resolution (Gerber and Chai 2010) suggest how this representation could

be enriched by identifying the referent of *he* from the surrounding text. All of these pieces of the semantic puzzle have the potential to fit together into a richer and deeper semantic representation of text. To further this goal, we intend to develop our classifier for all of the verbs in VerbNet and release the system to the public, along with an expanded version of the Semlink corpus.¹

References

- Abend, O., Reichart, R., & Rappoport, A. (2008). A supervised algorithm for verb disambiguation into verbnet classes. In *Proceedings of the 22nd international conference on computational linguistics – Volume 1* (pp. 9–16). Stroudsburg: Association for Computational Linguistics.
- Bikel, D. M., Schwartz, R., & Weischedel, R. M. (1999). An algorithm that learns what's in a name. *Machine Learning*, 34(1–3), 211–231.
- Bobrow, D., Cheslow, B., Condoravdi, C., Karttunen, L., King, T., Nairn, R., de Paiva, V., Price, C., & Zaenen, A. (2007). PARC's bridge and question answering system. In *Grammar engineering across frameworks* (pp. 46–66).
- Chang, C.-C., & Lin, C.-J. (2001). LIBSVM: A library for support vector machines.
- Dang, H. T., & Palmer, M. (2002). Combining contextual features for word sense disambiguation. In *Proceedings of the ACL-02 workshop on word sense disambiguation*, Morristown, NJ, USA (pp. 88–94). Stroudsburg: Association for Computational Linguistics.
- Dligach, D. (2011). *High-performance word sense disambiguation with less manual effort*. Ph.D. thesis, University of Colorado at Boulder.
- Dligach, D., & Palmer, M. (2008). Novel semantic features for verb sense disambiguation. In *HLT '08: Proceedings of the 46th annual meeting of the Association for Computational Linguistics on human language technologies*, Morristown, NJ, USA (pp. 29–32). Stroudsburg: Association for Computational Linguistics.
- Federici, S., Montemagni, S., & Pirrelli, V. (1999). Sense: An analogy-based word sense disambiguation system. *Natural Language Engineering*, 5(2), 207–218.
- Fellbaum, C. (1998). *WordNet: An electronic lexical database*. Cambridge: MIT Press.
- Gerber, M., & Chai, J. (2010). Beyond nombank: A study of implicit arguments for nominal predicates. In *Proceedings of the 48th annual meeting of the Association for Computational Linguistics* (pp. 1583–1592). Stroudsburg: Association for Computational Linguistics.
- Girju, R., Roth, D., & Sammons, M. (2005). Token-level disambiguation of verbnet classes. *Urbana*, 51, 61801.
- Haghighi, A., & Klein, D. (2009). Simple coreference resolution with rich syntactic and semantic features. In *Proceedings of the 2009 conference on empirical methods in natural language processing: Volume 3* (pp. 1152–1161). Stroudsburg: Association for Computational Linguistics.
- Hanks, P. (1996). Contextual dependency and lexical sets. *Dépendance contextuelle et ensembles lexicaux. International Journal of Corpus Linguistics*, 1(1), 75–98.
- Hensman, S., & Dunnion, J. (2004). Automatically building conceptual graphs using verbnet and wordnet. In *Proceedings of the 2004 international symposium on information and communication technologies* (pp. 115–120). Dublin: Trinity College.
- Ide, N., & Wilks, Y. (2006). Making sense about sense. In *Word sense disambiguation* (pp. 47–73).
- Loper, E., Yi, S., & Palmer, M. (2007). Combining lexical resources: Mapping between propbank and verbnet. In *Proceedings of the 7th international workshop on computational linguistics*, Tilburg, the Netherlands.

¹This expanded corpus is now available at https://verbs.colorado.edu/wiki/index.php/Main_Page.

- Nivre, J., Hall, J., Nilsson, J., Chanev, A., Eryigit, G., Kübler, S., Marinov, S., & Marsi, E. (2007). MaltParser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(02), 95–135.
- Palmer, M., Gildea, D., & Xue, N. (2010). Semantic role labeling. *Synthesis Lectures on Human Language Technologies*, 3(1), 1–103.
- Sanderson, M. (2000). Retrieving with good sense. *Information Retrieval*, 2(1), 49–69.
- Settles, B. (2010). *Active learning literature survey*. Madison: University of Wisconsin.
- Shi, L., & Mihalcea, R. (2005). Putting pieces together: Combining framenet, verbnets and wordnet for robust semantic parsing. In *Computational linguistics and intelligent text processing* (pp. 100–111).
- Swier, R., & Stevenson, S. (2004). Unsupervised semantic role labelling. In *Proceedings of the 2004 conference on empirical methods in natural language processing* (pp. 95–102).
- Zaenen, A., Condoravdi, C., & Bobrow, D. (2008). The encoding of lexical implications in verbnets: predicates of change of locations. 716/922 EF 8.
- Zapirain, B., Agirre, E., & Marquez, L. (2008). Robustness and generalization of role sets: Propbank vs. verbnets. In *Proc. of ACL*.