Inducing a Semantic Frame Lexicon from WordNet Data

Rebecca GREEN[*†‡] and **Bonnie DORR**[*†]
*Institute for Advanced Computer Studies
†Department of Computer Science
‡College of Information Studies
University of Maryland
College Park, MD 20742 USA
{rgreen, bonnie}@umiacs.umd.edu

Abstract

This paper presents SemFrame, a system that automatically induces the names and internal structures of semantic frames. After SemFrame identifies sets of frame-evoking verb synsets, the conceptual density of nodes in the WordNet network for corresponding nouns and noun synsets is computed and analyzed. Conceptually dense nodes are candidates for frame names and frame slots. Ca. 76% of the frame names and 87% of the frame slots generated by SemFrame are rated adequate by human judges.

1 Introduction

The essence of the paraphrase problem is that semantic content may be expressed in a variety of ways. Lexical synonymy, syntactic variation, overlapping meanings, and various other phenomena interact to produce a broad range of choices for most language generation tasks (Hirst, 2003; Rinaldi et al., 2003; Kozlowski et al., 2003). At the same time, natural language understanding must recognize what remains constant across paraphrases.

Semantic frames (Fillmore 1982) address the paraphrase problem by providing slot-and-filler templates to represent frequently occurring, structured experiences. Since frames are situation-based, frame-based representations of strict paraphrases should be (very nearly) identical; the relationship between frame-based representations of looser paraphrases should be readily observable. The gestalt nature of such frame-based representations provides a psychologically plausible basis for representing text meaning.

To acquire a comprehensive set of semantic frames, we need the capacity to generate frames (semi-)automatically, since generating them by hand is labor-intensive and unsystematic. To address this need, we have developed SemFrame, a system that automatically identifies (1) sets of verb senses that evoke a common semantic frame and (2) the frame's participant structure. This paper explores the second task, which involves identifying an appropriate name for the frame and a set of frame slots.

Section 2 presents related research efforts on the generation of semantic frames and templates for information extraction. Section 3 summarizes the features of WordNet that support the automatic induction of semantic frame structures, while Section 4 sets forth the approach taken by SemFrame to accomplish this task. Section 5 presents an evaluation of SemFrame's ability to identify frame and frame slot names, while Section 6 discusses how a major weakness uncovered by the evaluation can be addressed. Section 7 summarizes our contribution.

2 Related Work

Until now, semantic frames have been generated by hand (as in Fillmore and Atkins 1992), based on native speaker intuition;[1] the FrameNet project (http://www.icsi.berkeley.edu/~framenet; Johnson et al., 2002) now couples this generation with empirical validation.

Since information extraction (IE) templates and semantic frame structures are essentially the same thing (i.e., a set of semantic types for a kind of situation), the work of Riloff and Jones (1999) on semantic-lexicon induction for information

---

[1]Generation of semantic frames here refers to the identification of situation *types*, with their participant structures. Automatic generation of semantic frame *instantiations* has been pursued to a limited degree, as in Gildea and Jurafsky (2002) and Erk et al. (2003). The Senseval-3 Automatic Labeling of Semantic Roles task (Litkowski 2004) further promotes this effort.

extraction would appear to contradict the previous claim. However, in their work both domains (which correspond to frames) and semantic categories (which correspond to frame slots) are predefined. Similarly, the work of Riloff et al. (2002) to induce information extraction systems by cross-language projection relies critically on the prior (manual) generation of IE templates. The development of the German equivalent to FrameNet (Erk et al., 2003) is likewise based largely on the re-use of semantic frames developed manually for English.

While computationally-oriented research on semantic frames has expanded dramatically in recent years, SemFrame is alone in addressing the need to generate semantic frame structures automatically. This need stems from the labor-intensive nature of generating frames by hand; for example, the FrameNet effort to generate a few hundred frames is measured in person-years.

3 Resources Used in SemFrame

The current iteration of SemFrame relies heavily on analyzing data in a pre-existing resource, WordNet, a machine-readable lexico-semantic database whose primary organizational structure is the synset—a set of synonymous word senses. Associated with each synset is a gloss; many synsets also include an example sentence. Synsets are interconnected by such relations as antonymy, hyponymy/troponymy, meronymy, cause, entailment, and verb group; WordNet 2.0 also includes links between verb synsets and noun synsets based on morphological derivation or shared subject domain (referred to as "category" domains).

Such a resource has several advantages over corpus data in identifying semantic frames. First, definitions often mention their participants using semantic-type-like nouns, thus mapping easily to the corresponding frame element. Corpus data, however, are more likely to include instantiated participants, which may not generalize readily to the frame element. Second, lexical resources provide a consistent amount of data for word senses, while the amount of data in a corpus for word senses is likely to vary widely. Third, lexical resources provide their data in a more systematic fashion than do corpora.

4 SemFrame Approach

The work reported here relies on prior analysis of

semantic relationships in the WordNet verb network and of the vocabulary used in WordNet verb synset glosses. This analysis produces numerous (ca. 2000) groups of verb synsets hypothesized to evoke the same semantic frame. A group of verb senses that evoke the same frame is referred to here as a *frameset*. For example, one frameset identified in this prior analysis grouped together the four synsets displayed in Figure 1. Further details on the process of inducing framesets are given in (Green et al., 2004).

kidnap, nobble, abduct, snatch: take away to an undisclosed location against their will and usually in order to extract a ransom
ransom, redeem: exchange or buy back for money; under threat
shanghai, impress: take (someone) against his will for compulsory service, especially on board a ship
seize: take or capture by force

Figure 1. Example Frameset

Given a verb sense frameset, SemFrame uses WordNet-based associations between these verb synsets and both nouns and noun synsets to propose a name for the overall frame and a set of frame slots (i.e., the frame's participant structure).

The overall approach is shown in Figure 2. Step 1 extracts the nouns and noun synsets related in WordNet to the verbs in an input frameset. Step 2 computes the conceptual density of WordNet synsets/nodes based on the frequency of the noun synsets from step 1 and the hierarchical relationships among the nodes. Step 3 analyzes the output of these calculations and selects a subset of the nodes with the highest conceptual density scores. These nodes serve as the participant structure of the frame, and one of the conceptually dense nodes provides a name for the frame.

4.1 Extracting Related Nouns

The major observation that underlies this approach is that the semantic arguments of a verb sense/synset (and thus the semantic types of slots in the frame evoked by the verb sense) are closely correlated with three sets of nouns found in WordNet: (1) nouns used in the glosses and example sentences for the verb synset; (2) nouns that are morphologically related to the verb synset; and (3) nouns belonging to the same category domain as the verb synset.

```
Input.  WordNet verb synset framesets (steps
        executed for each frameset separately)

Step 1. Gather nouns associated with frameset
        verb synsets through:
        (a) Extraction of nouns from Minipar-
            parsed verb synset glosses and
            example sentences
        (b) Extraction of WordNet noun synsets
            with morphological derivation links to
            verb synsets
        (c) Extraction of WordNet noun synsets
            with category domain links to verb
            synsets

Step 2. Compute conceptual density for all
        WordNet noun synsets identified in
        step 1.

Step 3. Analyze step 2 output.

Output. (a) Frame name  (b) Frame slots
```
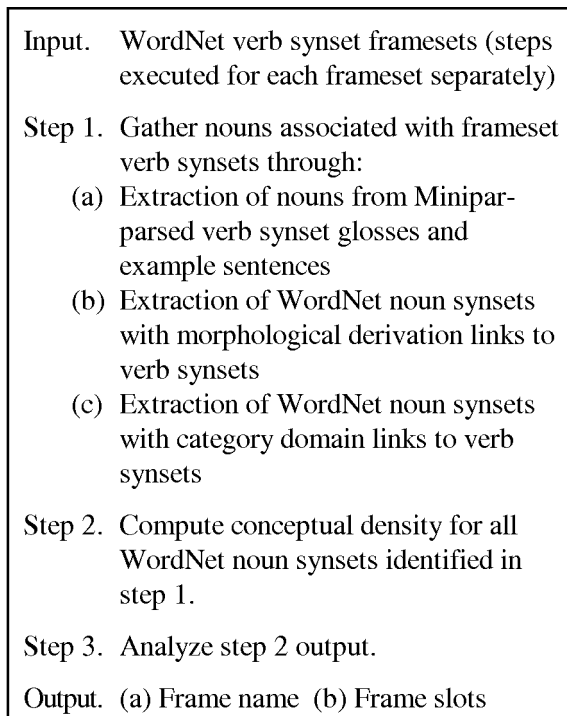
Figure 2. Semantic Frame Generation Algorithm

The first set of nouns are those found in glosses and example sentences for the verbs of a frameset. For example, a sense of *sell* that evokes the COMMERCIAL TRANSACTION frame is defined as "exchange or deliver for *money* or its equivalent," while a sense of *cost* that evokes the same frame is defined as "the total spent for *goods* or services including *money* and time and labor." Italicized words correspond closely to Money or Merchandise, two of the prominent participants in the COMMERCIAL TRANSACTION frame.

In order to isolate nouns in WordNet verb glosses potentially corresponding to a semantic argument of the verb, glosses were systematically manipulated to cast them as full sentences.[2] These sentential glosses, as well as any example sentences present, were processed by the Minipar (Lin 2001) parser. If Minipar identified a noun as a particular semantic type (e.g., time, money, number, person, location), this information was also retained.[3]

---

[2]For the most part this manipulation consisted of prepending the phrase "To $verb_1$ / $verb_2$ / $verb_3$ ... is to" to each gloss.

[3] For example, dollar amounts (e.g., "$5") are designated as MONEY, cardinal numbers (e.g., "two") as NUM, periods of the day (e.g., "night") as TIME.

The second set of nouns that shed light on the participant structure of a semantic frame is the set of nouns that are morphologically related to the verb senses in the frameset. For example, Buyer and Seller are two of the prominent participants in the COMMERCIAL TRANSACTION frame, which is evoked by *buy* and *sell*. Along with Merchandise and Money, Buyer and Seller comprise the full set of major participants of this frame. Figure 3 summarizes the morphological derivation links for the COMMERCIAL TRANSACTION senses of *buy* and *sell* provided in WordNet 2.0. Because WordNet 2.0 records such links directly between verb and noun synsets, the extraction of such data is altogether straightforward. Of 79,689 noun synsets in WordNet 2.0, 11,709 contain one or more links to morphologically-related verbs. Of 13,508 verb synsets, 8,906 contain one or more links to morphologically-related nouns.

```
v.  buy | obtain by purchase
v.  buy | be worth or be capable of buying
n.  buy | an advantageous purchase
n.  buyer | a person who buys
n.  buying | the act of buying

v.  sell | exchange or deliver for money or its
         equivalent
v.  sell | do business
n.  sell | the activity of persuading someone to
         buy
n.  seller | someone who promotes or exchanges
         goods or services for money
n.  selling | the exchange of goods for an agreed
         sum of money
```

Figure 3. Glosses for Nouns and Verbs that are
           Morphologically Related in WordNet

The third set of nouns related to a verb frameset is available only under limited circumstances. These are the noun synsets that have been identified as belonging to the same category domain as the verb synset. WordNet 2.0 includes 422 category domains. Again, because WordNet 2.0 records such links directly between verb and noun synsets, the extraction of such data is altogether straightforward.

Through the processes just described, every WordNet verb synset has associated with it a set of nouns, each of which appears in its gloss, appears in a sentence exemplifying its use, is morphologically related to it, or is in the same

category domain. For the frameset in Figure 1, the associated nouns are given in Figure 4. Each such noun is initially assigned a default weight of 1.0.

abduction, board, buy, capture, crime, exchange, force, hostage, impress, industrialist, kidnapper, kidnapping, location, men, money, order, politician, ransom, rebel, redeemer, redemption, service, shanghai, shanghaier, ship, someone, terrorist

Figure 4. Nouns Associated with
Example Frameset

Some, but not all, of the nouns are WordNet-sense-disambiguated as a result of these processes (for example, 10 of the 27 nouns in Figure 4). The remainder of the nouns are mapped to WordNet synsets through one of two methods.

The first of these methods associates nouns that received semantic type designations in the Minipar parsing with specific (i.e., corresponding) WordNet nodes. These correspondences are based on human judgment.

The other method completes the word sense disambiguation process by making default assignments for all nouns not associated with WordNet noun synsets through previous analysis. The strategy takes into account that WordNet senses are ordered by the frequency of their use in SEMCOR. Thus the first sense given for a word in WordNet has the highest *a priori* possibility of being the correct sense of the word. Specifically, the strategy is to weight the noun senses, assigning half of the noun's weight to the first sense, a fourth of the weight to a second sense, an eighth of the weight to a third sense, and so on. In this manner, the original weight assigned to the noun is distributed across its various senses in WordNet, roughly proportional to the *a priori* likelihood of the sense being a correct assignment.

As a result of these two processes, all nouns associated with a verb frameset are mapped to specific WordNet noun synsets. All these mappings have an associated weight. The various nodes within WordNet's noun network that correspond to a verb sense group constitute *evidence synsets* for the participant structure of the corresponding semantic frame.

4.2 Computing Conceptual Density

The overall idea behind transforming the list of evidence synsets into a list of participants involves

using the relationship structure of WordNet to identify an appropriately small set of concepts (i.e., synsets) within WordNet that account for (i.e., are superordinate to) as many of the nouns as possible; such synsets will be referred to as *covering synsets*. If the only constraint were to account for as many of the frame-associated nouns as possible, the clear solution would be to pick out WordNet nodes at the top of their respective hierarchies, since the number of nouns covered by a synset is greater for higher covering synsets in the WordNet hierarchy. However, nodes at the highest levels of the WordNet tree are more general and abstract than frame slots typically are. We want to identify synsets that characterize the participant structure of the frame as closely and accurately as possible. At the same time, the number of participants in a frame is generally small, perhaps only two, rarely more than four or five. This observation motivates the desire to constrain the number of covering synsets identified.

The task of identifying the participant structure of the frame evoked by sets of verb senses relies on the hypothesis that the nouns associated with them will not be randomly distributed across WordNet, but will be clustered in various subtrees within the hierarchy. In essence, the task is to identify those clusters/subtrees and then to designate the nodes at the roots of the subtrees as covering synsets (subject to the aforementioned constraints).

After mapping all the nouns associated with a semantic frame to their corresponding WordNet noun synsets, the next step is to analyze the accumulated data from the evidence synsets. It is hypothesized that the WordNet subtrees with the highest density—where density takes into account the number of evidence synsets present in the subtree, their weight, and their relative location in the subtree—are the most likely to correspond to frame slots. Intuitively, when evidence synsets cluster together, the subtrees in which they occur will be more dense than those subtrees where few or no evidence synsets occur.

Agirre and Rigau (1995) define conceptual density as the ratio between the expected area of a subtree containing a word sense and some number of "marks" and the actual area, where area is a function of the height of a node and the number of its descendants. This conceptual density measure has inspired the measure used here, which is computed using the following definition:

$$CD(n) = \frac{\sum\limits_{i \, \in \, descendants_n} (wgt_i * treesize_i)}{treesize_n}$$

The density of a node n in the WordNet noun network has two components in SemFrame: first, the occurrence of a weighted evidence synset at that node and, second, the occurrence of weighted evidence synsets at descendant nodes. For each node, record is initially made of (1) the cumulative weight associated with all nouns mapped to that node, (2) the number of nodes subordinate to that node within WordNet (its *treesize*), and (3) the node's area, which is the product of its cumulative weight and its treesize. Subsequently, a node's area is added to the area of all its ancestors. The conceptual density of a node n is computed in SemFrame as the ratio between its (cumulative) area and its (invariant) treesize:[4]

Let us look at some of the ramifications of this measure. First, if a node is not an evidence synset and none of its descendant nodes are evidence synsets, the node's density is 0 (since all weights are 0). Second, if an evidence synset has no descendant nodes that are evidence synsets, its density will equal its own cumulative weight.[5] Third, a node with descendant nodes that are evidence synsets will always have a higher density than the same node without descendant nodes that are evidence synsets.[6] Fourth, if a node is not itself an evidence synset, but all its children have density x, the node will also have density x.[7] It follows that an ancestor node that is an evidence synset, even if its weight is not very high, may have a higher

---

[4]That is to say, a node's area is affected by evidence synsets below it, but its treesize does not change on account of any of its descendants being evidence synsets.

[5]In this case, the treesize that is a multiplicative factor in computing the node's area in the numerator of the density ratio will cancel out the treesize in the denominator.

[6]Under these circumstances, the addition of the descendant's area to a node's cumulative area will increase the numerator of the density ratio, while the denominator's value will hold constant.

[7]Let S = the sum of the treesize of all the node's children. The node's cumulative area is, by the distributive law, the product Sx. Its treesize is S, and therefore its density is x.

density value than a more heavily weighted descendant node, especially if most or all of its descendants are heavily weighted.

4.3 Interpreting Conceptual Density

After the conceptual density of each evidence synset and all its ancestor synsets is computed, a final process undertakes the interpretation of the density measures. Two thresholds are used to eliminate WordNet nodes with either (1) low density values or (2) high density values but little overall support.

The number of WordNet nodes with non-zero densities for a specific semantic frame depends in large part on the number of verb senses associated with the frame. But the complexity of the frame—that is, the number of slots in its internal participant structure—tends not to vary. It is desirable then to establish a relative threshold for density values rather than an absolute threshold. We have used a relative threshold that is a multiplicative factor of the mean density of nodes within the processing of a particular semantic frame. This threshold helps retain nodes of interest while eliminating many spurious nodes.

Exceeding the density value threshold is not always enough. It is also desirable to identify nodes that receive their support from multiple sources. A node receives support in the same proportion that it is awarded a noun's weight: If a noun is associated with a single WordNet synset, it contributes 1.0 to that node's support, but if a noun's weight is distributed across multiple synsets in the disambiguation process described in Section 4.1, then it contributes .5, .25, .125, etc., according as it is the first sense, the second sense, the third sense, etc. Requiring that a node's support exceed a threshold helps minimize the effect of nouns erroneously associated with a semantic frame or mismapped to a WordNet synset.

The optimal values for density and support thresholds are still under investigation.

Of the nodes that meet these threshold criteria, many are related hierarchically. That is, there is often considerable conceptual redundancy in the set of nodes that remain. As a further filter, of nodes with a direct hierarchical relationship, only the one with the highest conceptual density is retained.

4.3.1 Identifying Semantic Frame Names

The WordNet semantic network for nouns is

divided into approximately a dozen major subnetworks. Nodes at the top of these subnetworks establish the semantic type of all their descendant nodes.

Semantic frames sought in connection with the paraphrase phenomenon are of certain semantic types. Specifically, the set of semantic types appropriate to frames includes abstractions, actions, events, phenomena, psychological features, and states. Accordingly, the node with the highest density value from among these subnetworks is designated as corresponding to the name of the semantic frame. For example, Figure 5 shows the noun synsets with the highest conceptual density values for the example frameset. The first two nodes are Person types, while the last five are Action types; of these the node labeled "Capture" has the highest value and is designated as the frame name.

shanghaier: a kidnapper who drugs men and
    takes them for compulsory service aboard
    a ship
kidnapper: someone who unlawfully seizes
    and detains a victim (usually for ransom)
impress: the act of coercing someone into
    government service
kidnapping: (law) the unlawful act of
    capturing and carrying away a person
    against their will and holding them in
    false imprisonment
capture: the act of taking of a person by force
crime: (criminal law) an act punishable by
    law; usually considered an evil act

Figure 5. Conceptually Dense Noun Nodes

The actual frame name is chosen from among the nouns in the synset. If the synset occurs within SEMCOR's Brown Corpus tagging, the noun that corresponds to that synset most frequently is designated as the frame name. If the synset does not occur within SEMCOR's tagging, the noun that is listed first in the synset is arbitrarily chosen as the frame name.

4.3.2 Identifying Semantic Frame Slots

All the remaining nodes are candidates for correspondence with slots in the internal structure of the frame. The organization of the WordNet noun network by semantic types helps further restrict the set of nodes. Comparative analysis of

sets of nodes has resulted in the insight that, with two exceptions, no more than one slot of a particular semantic type generally occurs in a frame; the exceptions are the entity and abstraction types. Thus, one final filter is implemented such that only the highest density node in any given subnetwork is output, except for the entity and abstraction subnetworks. In the example, "Capture" (but not "Impress," "Kidnapping," or "Crime") is retained as the node with the highest conceptual density value within the Action subnetwork. Although both are in the entity subnetwork, "Kidnapper" is retained but not "Shanghaier," due to its higher conceptual density score and their direct hierarchical relationship.

The resulting set of densely weighted noun synsets are not always at what intuitively seems the appropriate level for names of frame slots. To correct this situation, a set of approximately 40 nodes across the WordNet noun network that correspond most closely to standard frame elements have been picked out by hand. Nodes with high density scores that are within any of the subtrees dominated by one of these 40-odd nodes are replaced by the "standard" frame element. In the example case, "Capture" is generalized by this method to "Action." Thus, two frame slots are identified for the example frameset: Kidnapper and Action.

5   Evaluation

The semantic frames generated by SemFrame have been validated using human judgments. For each of four combinations of threshold values from the prior hypothesized verb frameset identification step (reflecting two values of verb synset connectivity, .75 and 1.0, and two clustering threshold values, .5 and 1.0^8), a set of 14 randomly selected semantic frames was presented to two or more judges.[9] Each frame was described by a set of verb synsets (a verb frameset) with glosses and example sentences, on the one hand, and a set of noun synsets with glosses, one of which is designated as the frame name and the others of

---

[8]For the clustering algorithm used, the clustering threshold range is open-ended. The values investigated in the evaluation are fairly low.

[9]The judges were graduate students recruited from a course taught by one of the authors. All had at least passing familiarity with semantic frames.

which are designated as frame slots, on the other hand.

The human judges were first asked whether all/most/a significant subset of the verb synsets could be used to convey similar information about a situation. If at least a significant subset of the verb synsets could be so used, the judges were then asked how well the frame name characterized the corresponding situation, with regard to both the single lexical unit used to name the synset and the synset's gloss. Judges were next asked how well the frame slots characterized the participants in the situation. Lastly, they were asked whether any participants were missing.

Analysis of the human judgments revealed that the lower (.75) verb synset connectivity threshold yielded significantly better results.[10] Accordingly, the performance measures cited here are drawn only from judgments for the corresponding two data sets (i.e., for either of the two clustering thresholds).

Judges found in the case of 88% of the hypothesized framesets that at least a significant subset of the verb synsets could be used to convey information about the same, a similar, or a closely related situation (this characterization was intended to capture the notion of evoking the same semantic frame). Of these, all or most of the verb synsets evoke the same frame in 70% of the cases; only a significant subset evoke the same frame in 30% of the cases. The remaining questions were all answered with respect to the largest significant subset of same-frame-evoking verb synsets.

The frame name was found to characterize the verb synsets "very well" or "OK" for 83% of the SemFrame frames, with the corresponding gloss being judged similarly adequate 76% of the time. (That is, for a small number of cases, an appropriate frame name was designated, but it was drawn from an incorrect sense of the word.)

Frame slot designations (the synset name and gloss were not separated for slots) were deemed to capture a participant type "well" or "OK" 87% of the time. There are, however, two significant drawbacks in SemFrame's current ability to identify frame slots. The first drawback is that many slots that should be identified are not being identified. Overall, SemFrame's frames average only .75 slots; those that have at least one slot

---

[10]According to χ-square tests, both frame names (df =1, $\alpha \approx$ .05) and frame slots (df = 2, α = .01) are better identified using a verb synset connectivity of .75 than using a connectivity of 1.0.

average 1.85 slots. The one judge who took the request for missing participants seriously, however, identified on average another 2.15 slots. The second drawback is that 38% of the slots that were identified are considered by the human judges to be too general.

6 Future Directions

Clearly the inadequacy of the current iteration of SemFrame in identifying frame slots is its biggest weakness. It is surmised that this deficiency is rooted in size limitations on the number of nouns associated with each verb frameset.

Reference to the initial iteration of SemFrame suggests a way in which SemFrame's capabilities relative to identifying frame slots can be improved. At first, SemFrame used data from both *Longman's Dictionary of Contemporary English* (LDOCE; Procter, 1978) and WordNet (version 1.7.1), using LDOCE as its base. It was anticipated that the restricted vocabulary used in LDOCE's definitions and example sentences would play a major role in collocating verb senses that evoke the same frame. This assumption, however, turned out to be overly optimistic. The availability of an updated version of WordNet held promise of overcoming many of the problems associated with basing SemFrame on LDOCE and of mapping WordNet verb senses to LDOCE verb senses, at the same time that its expanded array of links duplicated many of the semantic relationships found in LDOCE.

WordNet's organization into synsets is a two-edged sword for SemFrame. On the one hand, it groups synonymous senses of verbs together. Since synonymy operates within a semantic frame, WordNet's synset organization provides a solid initial base for the identification of groups of same-frame-evoking verb synsets. On the other hand, WordNet gives a single gloss for a synset, thus minimizing the data available for identifying frame names and especially frame slots.

The re-incorporation of LDOCE as a data source is likely to improve SemFrame's ability to deal with frame slots by expanding the number of nouns associated with a verb frameset. This will require creating a mapping between LDOCE verb senses and WordNet verb framesets. The relative richness of the current verb framesets should support high-quality mapping. This in turn should permit the incorporation of an expanded set of associated nouns from LDOCE's definitions, which are often better reflectors of the semantic

types of a verb's prototypical arguments than the nouns in WordNet glosses.

7 Conclusions

This paper has described how sets of verb senses that evoke a semantic frame are expanded to target the corresponding nodes in the WordNet noun network. The conceptual density of these nodes is computed, which is used to identify noun nodes most likely to correspond to slots within the corresponding frame (a strategy for improving SemFrame's capabilities in this area has been introduced), as well as to identify the WordNet node that best corresponds to the overall frame. To our knowledge this is the first time that such processes have been undertaken automatically.

Acknowledgments

References

Agirre, Eneko and German Rigau. 1995. A proposal for word sense disambiguation using conceptual distance. In *1st International Conference on Recent Advances in NLP*. Available: <http://xxx.lanl.gov/abs/cmp-lg/9510003> [2004 February 5].

Erk, Katrin, Andrea Kowalski, Sebastian Padó, and Manfred Pinkal. 2003. Towards a resource for lexical semantics: A large German corpus with extensive semantic annotation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, 537-544.

Fillmore, Charles J. 1982. Frame semantics. In *Linguistics in the Morning Calm*, 111-137. Seoul: Hanshin.

Fillmore, Charles J. and B. T. S. Atkins. 1992. Towards a frame-based lexicon: The semantics of RISK and its neighbors. In A. Lehrer and E. F. Kittay (Eds.), *Frames, Fields, and Contrasts*, 75-102. Hillsdale, NJ: Erlbaum.

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

Green, Rebecca, Bonnie J. Dorr, and Philip Resnik. 2004. Inducing Frame Semantic Verb Classes from WordNet and LDOCE. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*.

Hirst, Graeme. 2003. Paraphrasing paraphrased. Keynote address for *The Second International Workshop on Paraphrasing: Paraphrase Acquisition and Applications*, ACL 2003, <http://nlp.nagaokaut.ac.jp/IWP2003/pdf/Hirst-slides.pdf>.

Johnson, Christopher R., Charles J. Fillmore, Miriam R. L. Petruck, Collin F. Baker, Michael Ellsworth, Josef Ruppenhofer, and Esther J. Wood. 2002. *FrameNet: Theory and Practice*, version 1.0, <http://www.icsi.berkeley.edu/~framenet/book/book.html>.

Kozlowski, Raymond, Kathleen F. McCoy, and K. Vijay-Shanker. 2003. Generation of single-sentence paraphrases from predicate/argument structure using lexico-grammatical resources. In *The Second International Workshop on Paraphrasing: Paraphrase Acquisition and Applications (IWP2003)*, ACL 2003, 1-8.

Lin, Dekang. 2001. LaTaT: Language and text analysis tools. In *Proceedings of the First International Conference on Human Language Technology Research*, 222-227.

Litkowski, Ken. 2004. Senseval-3 task: Word-sense disambiguation of WordNet glosses, <http://www.clres.com/SensWNDisamb.html>.

Procter, Paul (Ed.). 1978. *Longman Dictionary of Contemporary English*. Longman Group Ltd., Essex, UK.

Riloff, Ellen and Jones, Rosie. 1999. Learning dictionaries for information extraction by multi-level bootstrapping. *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI-99)*, 474-479.

Riloff, Ellen, Charles Schafer, and David Yarowsky. 2002. Inducing information extraction systems for new languages via cross-language projection. *Proceedings of the 19th International Conference on Computational Linguistics*.

Rinaldi, Fabio, James Dowdall, Kaarel Kaljurand, Michael Hess, and Diego Mollá. 2003. Exploiting paraphrases in a question answering system. In *The Second International Workshop on Paraphrasing: Paraphrase Acquisition and Applications (IWP2003)*, ACL 2003, 25-32.