Automatic Annotation for All Semantic Layers in FrameNet

Richard Johansson and Pierre Nugues

Department of Computer Science, Lund University

Box 118

SE-221 00 Lund, Sweden

{richard, pierre}@cs.lth.se

Abstract

We describe a system for automatic annotation of English text in the FrameNet standard. In addition to the conventional annotation of frame elements and their semantic roles, we annotate additional semantic information such as support verbs and prepositions, aspectual markers, copular verbs, null arguments, and slot fillers. As far as we are aware, this is the first system that finds this information automatically.

1 Introduction

Shallow semantic parsing has been an active area of research during the last few years. Semantic parsers, which are typically based on the FrameNet (Baker et al., 1998) or PropBank formalisms, have proven useful in a number of NLP projects, such as information extraction and question answering. The main reason for their popularity is that they can produce a flat layer of semantic structure with a fair degree of robustness.

Building English semantic parsers for the FrameNet standard has been studied widely (Gildea and Jurafsky, 2002; Litkowski, 2004). These systems typically address the task of identifying and classifying Frame Elements (FEs), that is semantic arguments of predicates, for a given target word (predicate).

Although the FE layer is arguably the most central, the FrameNet annotation standard defines a number of additional semantic layers, which contain information about support expressions (verbs and prepositions), copulas, null arguments, slot-fillers, and aspectual particles. This information can for example be used in a semantic parser to

refine the meaning of a predicate, to link predicates in a sentence together, or possibly to improve detection and classification of FEs. The task of automatic reconstruction of the additional semantic layers has not been addressed by any previous system. In this work, we describe a system that automatically identifies the entities in those layers.

2 Introduction to FrameNet

FrameNet (Baker et al., 1998; Johnson et al., 2003) is a comprehensive lexical database that lists descriptions of words in the frame-semantic paradigm (Fillmore, 1976). The core concept is the *frame*, which is conceptual structure that represents a type of situation, object, or event, coupled with a semantic valence description that describes what kinds of semantic arguments (*frame elements*) are allowed or required for that particular frame. The frames are arranged in an ontology using relations such as inheritance (such as the relation between COMMUNICATION and COMMUNICATION_NOISE) and causative-of (such as KILLING and DEATH).

For each frame, FrameNet lists a set of lemmas or *lexical units* (mostly nouns, verbs, and adjectives, but also a few prepositions and adverbs). When such a word occurs in a sentence, it is called a *target word* that *evokes* the frame. FrameNet comes with a large set of manually annotated example sentences, which is typically used by statistical systems for training and testing. Figure 1 shows an example of such a sentence. Here, the target word *eat* evokes the INGESTION frame. Three FEs are present: INGESTOR, INGESTIBLES, and PLACE.

Often [an informal group]_{INGESTOR} will *eat* [lunch]_{INGESTIBLES} [near a machine or other work station]_{PLACE}, even though a canteen is available.

Figure 1: A sentence from the FrameNet example corpus, with FEs bracketed and the target word in italics.

3 Semantic Entities in FrameNet

The semantic annotation in FrameNet consists of a set of layers. One of the layers defines the target, and the other layers provide additional information with respect to the target. The following layers are used:

- The FE layer, which defines the spans and semantic roles of the arguments of the predicate.
- A part-of-speech-specific layer, which contains aspectual information for verbs; and copulas, support expressions, and slot filling information for nouns and adjectives.
- The “Other” layer, containing special cases such as null arguments.

The semantic entities that we consider in this article are defined in the second and third of these layers.

3.1 Support Expressions

Some noun targets, typically denoting events, are often constructed using *support verbs*. In this case, the noun carries most of the semantics (that is, it evokes the frame), while the verb allows the slots of the frame to be filled. Thus, the dependents of a support verb are annotated as FEs, just like for a verb target. Support verbs are annotated using the SUPP label on the Noun or Adjective layer. In the following sentence, there is a support verb (**underwent**) for the noun target (*operation*).

[Frances Patterson]_{PATIENT} **underwent** an *operation* at RMH today and is expected to be hospitalized for a week or more.

The support verbs do not change the core semantics of the noun target (that is, they bear no relation to the frame). However, they may determine the relation between the FEs and the target (“point-of-view supports”, such as “undergo an operation”

or “perform an operation”) or provide aspectual information (such as “start an operation”).

The following sentence shows an example where a governing verb is *not* a support verb of the noun target. An automatic system must be able to distinguish support verbs from other verbs.

A senior nurse observed the *operation*.

Although a large majority of the support expressions are verbs, there are additionally some cases of support prepositions, such as the following example:

Secret agents **of** this *ilk* are at work all the time.

3.2 Copulas

Copular verbs, typically *be*, may be seen as a special kind of support verb. They are marked using the COP label on the Noun or Adjective layer. There are several uses of copulas:

- Class membership: John **is** a *sailor*.
- Qualities: Your literary masterpiece **was** *delicious*.
- Location: This **was** *inside* a desk drawer.
- Identity: Smithers **is** the *vice-president* of the arm-chair division.

In FrameNet annotation, these uses of the copular verb are not distinguished.

3.3 Null Arguments

There are constructions that require special arguments to be syntactically valid, but where these arguments have no relation to the semantics of the sentence. In the example below, **it** is an example of this phenomenon.

I *hate* **it** when you do that.

Other common cases include existential constructions (“there are”) and subject requirement of zero-place predicates (“it rains”). These null arguments are tagged as NULL on the Other layer.

3.4 Aspectual Particles

Verb particles that indicate aspectual information are marked using the ASPECT label. These particles must be distinguished from particles that are parts of multiword units, such as *carry out*.

They just *moan* **on and on** about Fergie this and Fergie that and I’ve simply had enough.

3.5 Slot Fillers: GOV and X

FrameNet annotation contains some information about the relation of predicates in the same sentence when one predicate is a slot filler (that is, an argument) of the other. This is most common for noun target words, typically referring to natural kinds or artifacts.

In the following example, the target word *fingertips* evokes the OBSERVABLE_BODYPARTS frame, involving two FEs: POSSESSOR (“his”) and BODY_PART (“fingertips”). This noun phrase is also a slot filler (that is, an argument) of another predicate in the sentence: *cling on*. In FrameNet, such predicates are annotated using the GOV label. The constituent that contains the slot filler in question is called (for lack of a better name) X.

Shares will boom and John Major will
[cling on]_{GOV} [by [his]_{POSSESSOR}
[fingertips]_{BODY_PART}]X.

If GOV and X are present, all FEs must be contained in the span of the X node, such as BODY_PART and POSSESSOR above. This may be of use for automatic FE identifiers.

4 Identifying Semantic Entities

To find the semantic entities in the text, we used the method that has previously been used for FE detection: classification of nodes in a parse tree. We divide the identification process into two stages:

- The first stage finds SUPP, COP, and GOV.
- The second stage finds NULL, ASP, and X.

The reason for this division is that we expect that the knowledge of the presence of SUPP, COP, and GOV, which are almost always verbs, is useful when detecting the other entities. The second stage makes use of the information found in the first stage. Above all, it is necessary to have information about GOV to be able to detect X.

To train the classifiers, we selected the 150 most common frames and divided the annotated example sentences for those frames into a training set of 100,000 sentences and a test set of 8,000 sentences.

The classifiers used the Support Vector learning method using the LIBSVM package (Chang and Lin, 2001). The features used by the classifiers are listed in Table 1. Apart from the features used by

Features for first and second stage
Target lemma
Target POS
Voice
Available semantic role labels
Position (before or after target)
Head word and POS
Phrase type
Parse tree path from target to node
Features for second stage only
Has SUPP
Has COP
Has GOV
Parse tree path from SUPP to node
Parse tree path from COP to node
Parse tree path from GOV to node

Table 1: Features used by the classifiers.

Stage 2, most of them are well-known from previous literature on FE identification and labeling (Gildea and Jurafsky, 2002; Litkowski, 2004). For all path features, we used both the traditional constituent parse tree path (as by Gildea and Jurafsky (2002)) and a dependency tree path (as by Ahn et al. (2004)). We produced the parse trees using the parser of Collins (1999).

5 Evaluation

We applied the system to a test set consisting of approximately 8,000 sentences.

Because of inconsistent annotation, we did not evaluate the performance of detection of the EXIST tag used in existential constructions. Preliminary experiments indicated that the performance was very poor.

The results, with confidence intervals at the 95% level, are shown in Table 2. They demonstrate that the classical approach for FE identification, that is classification of nodes in the parse tree, is as well a viable method for detection of other kinds of semantic information. The detection of X shows the poorest performance. This is to be expected, since it is very dependent on a GOV to have been detected in the first stage.

The results for detection of aspectual particles is not very reliable (the confidence interval was ± 0.17 for precision and ± 0.19 for recall), since test corpus contained just 25 of these particles.

	P	R	$F_{\beta=1}$
SUPP	0.85 ± 0.046	0.64 ± 0.054	0.73
COP	0.90 ± 0.027	0.87 ± 0.030	0.88
NULL	0.76 ± 0.082	0.80 ± 0.080	0.78
ASP	0.83 ± 0.17	0.6 ± 0.19	0.70
GOV	0.79 ± 0.029	0.64 ± 0.030	0.71
X	0.59 ± 0.035	0.49 ± 0.032	0.54

Table 2: Results with 95% confidence intervals on the test set.

6 Conclusion and Future Work

We have described a system that reconstructs all semantic layers in FrameNet: in addition to the traditional task of building the FE layer, it marks up support expressions, aspectual particles, copulas, null arguments, and slot filling information (GOV/X). As far as we know, no previous system has addressed these tasks.

In the future, we would like to study how the information provided by the additional layers influence the performance of the traditional task for a semantic parser. FE identification, especially for noun and adjective target words, may be made easier by knowledge of the additional layers. As mentioned above, if a support verb is present, its dependents are arguments of the predicate. The same holds for copular verbs. GOV/X nodes also restrict where FEs may occur. In addition, support verbs (such as “perform” or “undergo” an operation) may be beneficial when determining the relationship between the FE and the predicate, that is when assigning semantic roles.

References

- David Ahn, Sisay Fissaha, Valentin Jijkoun, and Maarten de Rijke. 2004. The university of Amsterdam at Senseval-3: Semantic roles and logic forms. In *Proceedings of SENSEVAL-3*.
- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. In *Proceedings of COLING-ACL’98*, pages 86–90, Montréal, Canada.
- Chih-Chung Chang and Chih-Jen Lin, 2001. *LIBSVM: a library for support vector machines*.
- Michael J. Collins. 1999. Head-driven statistical models for natural language parsing. Ph.D. thesis, University of Pennsylvania, Philadelphia.
- Charles J. Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York*

Academy of Sciences: Conference on the Origin and Development of Language, 280:20–32.

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

Christopher Johnson, Miriam Petruck, Collin Baker, Michael Ellsworth, Josef Ruppenhofer, and Charles Fillmore. 2003. *FrameNet: Theory and Practice*.

Ken Litkowski. 2004. Senseval-3 task: Automatic labeling of semantic roles. In Rada Mihalcea and Phil Edmonds, editors, *Senseval-3: Third International Workshop on the Evaluation of Systems for the Semantic Analysis of Text*, pages 9–12, Barcelona, Spain, July. Association for Computational Linguistics.