Augmenting FrameNet Via PPDB

Pushpendre Rastogi[1] And Benjamin Van Durme[1,2]

[1]Center for Language and Speech Processing

[2]Human Language Technology Center of Excellence

Introduction

Frame semantic parsers analyze a sentence and annotate its frame evoking words and the corresponding frame elements/arguments.

The task of frame semantic parsing was introduced by Gildea and Jurafsky (2002) shortly after the release of FrameNet (Fillmore and Baker, 2001). CMU's SEMAFOR system (Das et al., 2013) is the current state-of-the-art.

Like other semantic resources, the coverage of FrameNet is incomplete. State-of-the-art frame semantic parsers thus employ heuristics to identify the frame evoked by out-of-vocabulary items (OOVs).

For example the word *jettison* does not appear in FrameNet. A parser may look up its synonym from WordNet, e.g., *Abandon* and its associated frame in FrameNet and assign to *jettison* that frame.

PPDB

Ganitkevitch et al. (2013) released a large collection of lexical, phrasal and syntactic paraphrases collectively called PPDB.[a]

Lexical : Two words with the same meaning.

Phrasal : Two strings of words with the same meaning.

Syntactic : Two strings consisting of words and non-terminal categories that have the same meaning, e.g.,

```
[S/PP] ||| NP/NN explosion VP/PP
||| NP/NN blast VP/PP
```

A lexical rule in PPDB looks like the following:
```
[VB] ||| help ||| assist |||
p(e|f)=2.832, p(f|e)=1.551, ...
```

This rule conveys that the log-probability that *help* would be paraphrased by the word *assist* is -2.832 but the log probability of *assist* being paraphrased as *help* is -1.551.

---
[a]paraphrase.org

References

Roni Ben Aharon, Idan Szpektor, and Ido Dagan. Generating entailment rules from framenet. In *Proceedings of the ACL 2010 Conference*, 2010.

Dipanjan Das, Desai Chen, André F. T. Martins, Nathan Schneider, and Noah A. Smith. Frame-semantic parsing. *Computational Linguistics*, 2013.

Oscar Ferrández, Michael Ellsworth, Rafael Munoz, and Collin F Baker. Aligning framenet and wordnet based on semantic neighborhoods. In *LREC*, 2010.

Charles J Fillmore and Collin F Baker. Frame semantics for text understanding. In *Proceedings of Word-Net and Other Lexical Resources Workshop, NAACL*, 2001.

Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. PPDB: The paraphrase database. In *Proceedings of NAACL-HLT*. ACL, June 2013.

Daniel Gildea and Daniel Jurafsky. Automatic labeling of semantic roles. *Computational linguistics*, 2002.

Contact

pushpendre@jhu.edu, vandurme@cs.jhu.edu

cs.jhu.edu/~prastog3, cs.jhu.edu/~vandurme

Objective of this work

We increase the coverage of the FrameNet corpus by expanding the collection of training examples via PPDB and n-gram language models.

Results

Method 1 : We generated a set of candidate paraphrases for every target word in a sentence by first finding the paraphrases of that word in PPDB.

Method 2 : We first enlarged the set of lexical units mapped to a frame by merging lexical units of frames that were related to the target word's frame through either of the following relations:
`Is_Perspectivized_In, Is_Inherited_By, Has_Subframe` (s)

PPDB	1a	1b	2a	2b
S	4,582	2,574	1,064,926	1,022,533
M	15,459	9,752	1,314,169	1,263,087
L	73,763	55,517	2,417,760	2,347,656
XL	340,406	283,126	–	–

Table 1: The total number of paraphrases generated for the 23,226 input targets versus different sizes of PPDB. The paraphrase count excludes the input. Column 1a and 2a represent unfiltered paraphrases as opposed to 1b and 2b where they have been filtered using WordNet v3.0.

We used the lexical rules along with a 5-gram Kneser-Ney language model trained on the raw English sequence of Annotated Gigaword to paraphrase the frame annotated sentences of FrameNet. We used a combination of the WordNet morphological analyzer and Morpha for lemmatization and Morphg for generation.

PPDB Size	1	2	3	%(1+2)	%(1)
S	0	0	0	–	–
M	6	1	2	77.77	66.67
L	27	15	11	86.25	50.94
L rank 3	23	12	7	83.33	54.76
XL	110	85	50	79.60	44.89
XL rank 3	47	16	9	87.5	65.27
XL rank 5	69	28	13	88.18	62.72
XL rank 10	105	52	32	83.07	55.55

Table 2: Average quality of all paraphrases for 25 random sentences. Rows marked *A rank B* convey that we used PPDB of size *A* and kept only the top *B* sentences after sorting them by their language model score. Column %(1) indicates the percentage of output which was grammatical and evoked the correct frame. Column %(1+2) represents the output that evoked the correct frame.

The General Assembly should set aside **money** for a new state health lab , millions of doses of antiviral drugs and a fund to help meet basic needs after a disaster , a legislative panel recommended Thursday .
1: The General Assembly should set aside **cash** ...
2: The General Assembly should set aside **fund** ...
1: The General Assembly should set aside **dough** ...
3: The General Assembly should set aside **silver** ...

Table 3: Examples and their judgements, with the last being debatable.

Conclusion and Error Analysis

We found that over a sample of 25 target words the top three paraphrases produced by PPDB XL evoked the correct frame and were grammatical 65% of the time.[a]

We found two major reasons for ungrammaticality of lexical paraphrases.

1. Within FrameNet some sentences will have a single token annotated as trigger, when in fact it is part of a multi-word expression. For example, it was grammatically infelicitous to replace *part* by *portion* in the expression *part of the answer*.

2. The other major source of error was inaccuracies in PPDB itself. We found that for a large number of cases when PPDB XL did not have a high number of paraphrases, the paraphrases were wrong (e.g., PPDB XL had only 2 paraphrases for the words *lab* and *millions*.)

---
[a]We have released the generated corpus as well as the manual annotations at http://cs.jhu.edu/~prastog3/res/fnppdb.html

Acknowledgments