Improving Semantic Parsing Using Statistical Word Sense Disambiguation

Ritwik Bose, Siddharth Vashishtha & James Allen

University of Rochester

{rbose, svashis3, james}@cs.rochester.edu

Objective

Improve sense decisions made by logical semantic parsing using the outputs of a statistical WSD system while honoring semantic role restrictions.

We consider the case study of improving a semantic parser, the TRIPS Parser [1], which is a best-first bottom-up chart-parser with a hand-built, lexicalized context-free grammar. It consists a TRIPS ontology which is a hand-crafted single-inheritance hierarchy where nodes specialize or override hierarchical features and argument templates from their parent. We improve the sense disambiguation of this logical semantic parser by integrating advice from a statistical Word Sense Disambiguation (WSD) system, SupWSD [3].

TRIPS Ontology and Lexicon

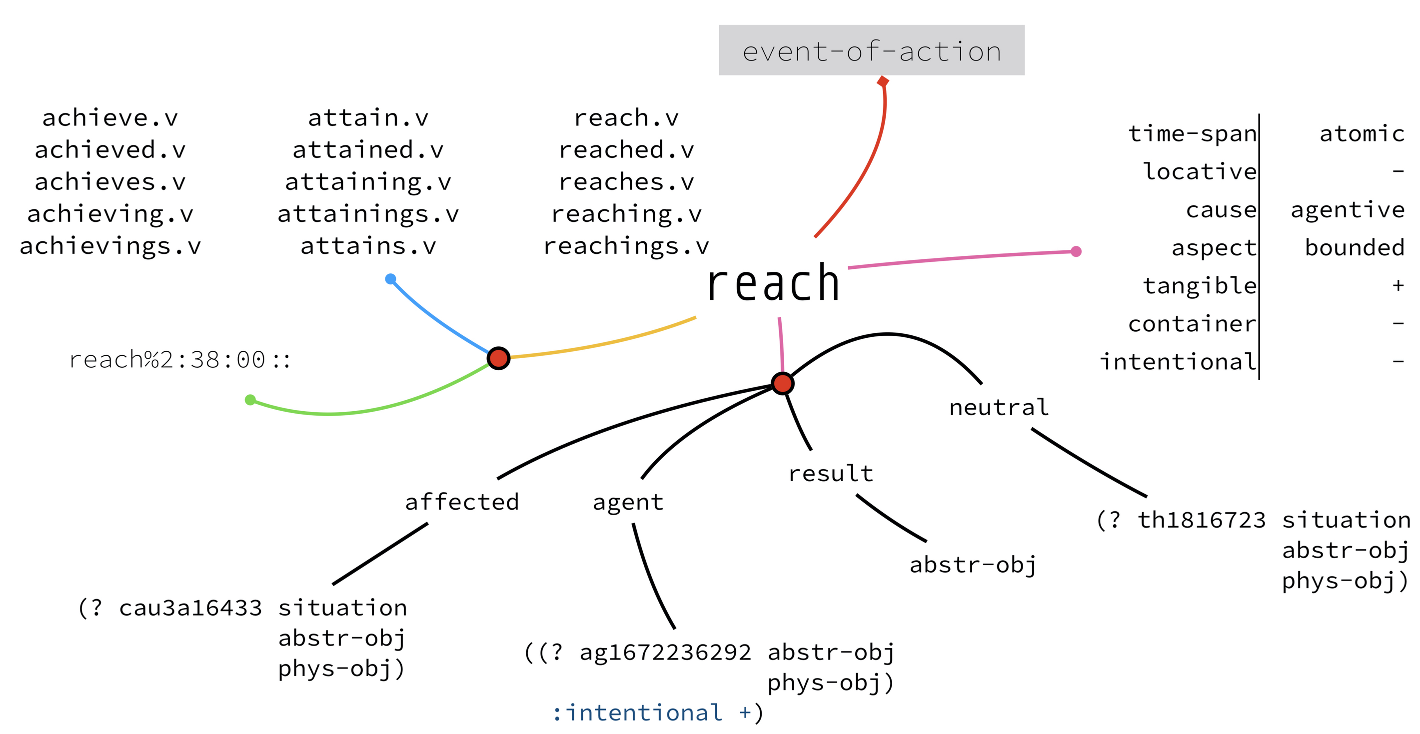

Figure 1: Entry for `ONT::REACH` in the TRIPS Ontology

SupWSD

We use a supervised-system SupWSD [3] which provides a **probability distribution** over **WordNet senses** for a given sentence. The system uses an **SVM classifier** to predict the word senses and its **features** include various linguistics properties such as POS tags, syntactic relations, local collocation, word-embeddings, and information about surrounding words. We use the version which is trained on **SEMCOR**.

Hinting

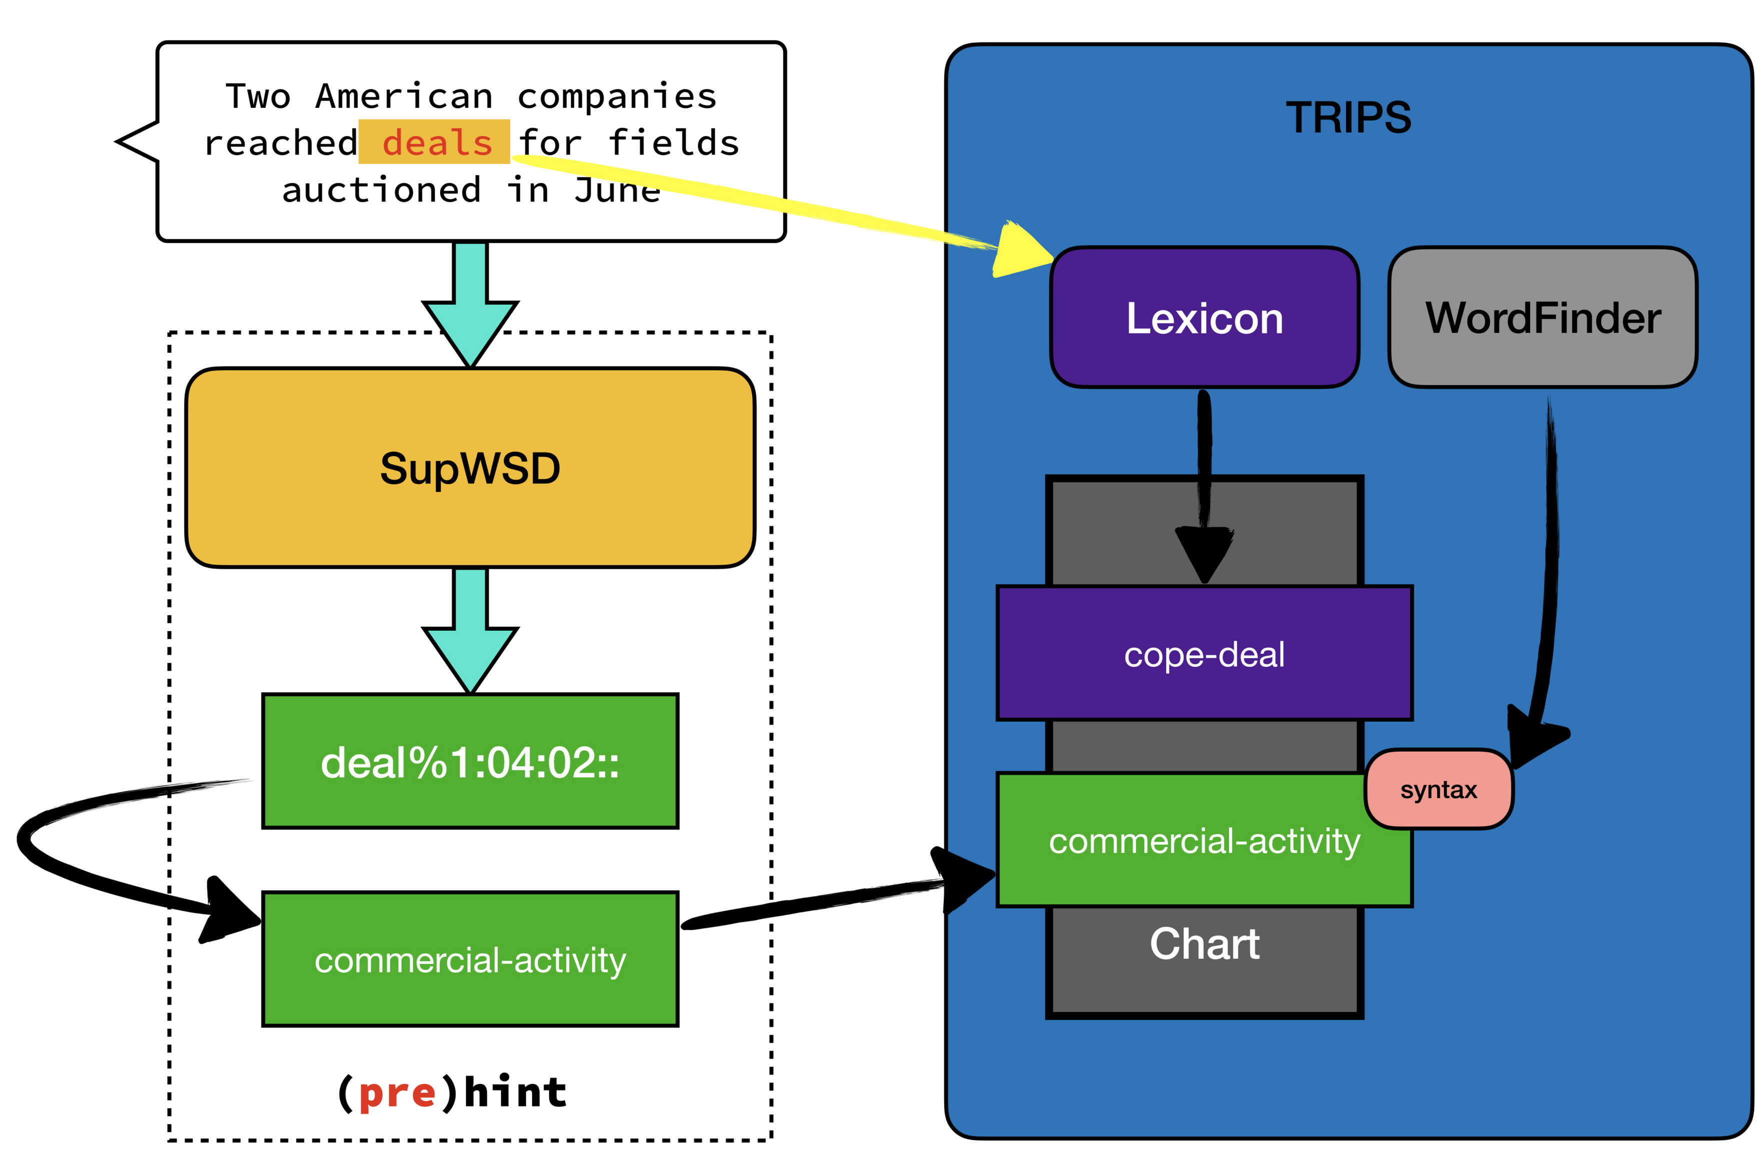

Figure 2: Hinting Pipeline: A hint for the word deals is produced by SupWSD, generating a wordnet sense. The WordNet sense is mapped to a TRIPS type and marked up with a syntax template and features and added to the chart.

Strategies

Hints can be provided to the parser by curating inputs to the chart or by actively reordering constituents as they are built.

- **Pre-Hinting:** The chart is populated with the top candidates from SupWSD
- **Progressive Hinting:** Constituents are reordered to prefer the best outputs from SupWSD
- **Combined Hinting:** All of the above

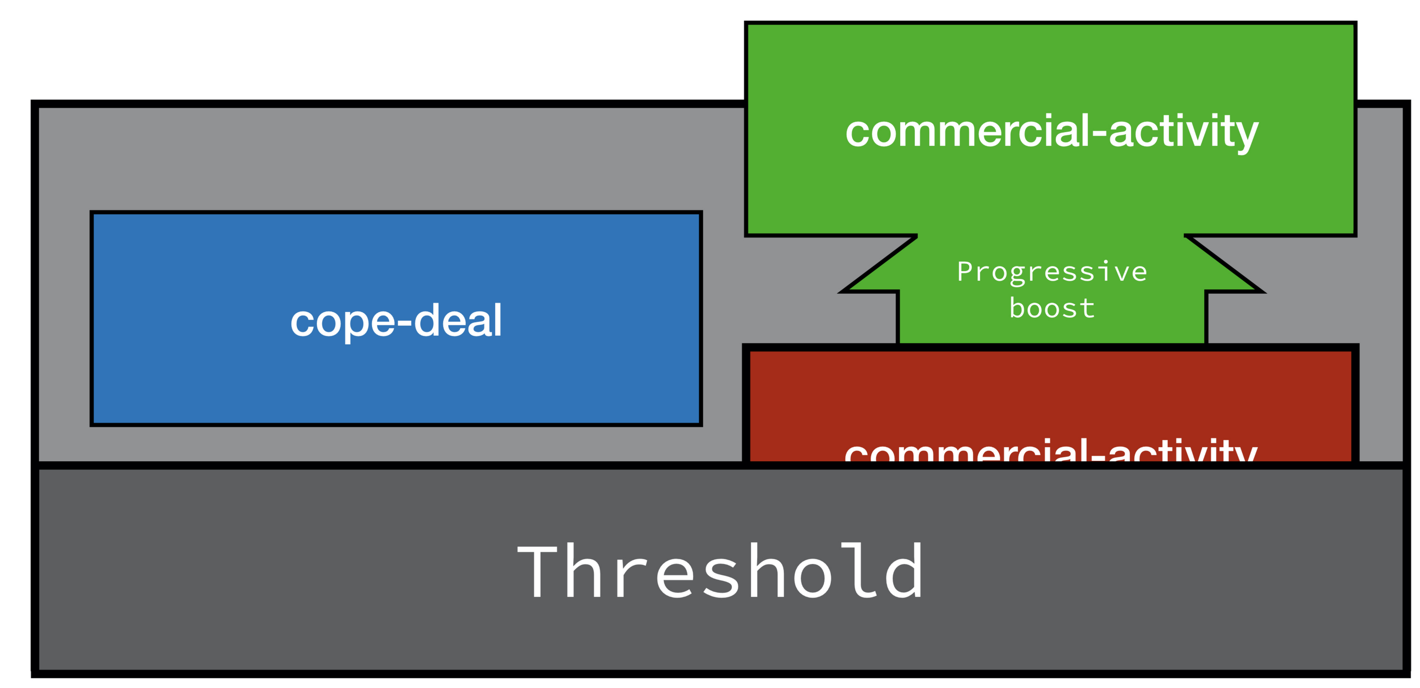

Figure 3: Progressive hinting strategy

Results

We report three metrics to evaluate the performance of the semantic parser – (i) accuracy (exact sense agreement) (ii) Mean Wu-Palmer similarity (Wu and Palmer 1994) and (iii) Mean accuracy over semantic factors (computed as Wu-Palmer over the factorized ontology).

Metric	SupWSD	Plain	Pre	Prog	Comb
Accuracy	66.19	39.42	50.33	42.45	53.19
WuP	84.22	73.97	79.12	75.75	80.53
Sem-fac	70.26	49.33	60.76	53.38	63.37

Table 1: Results from evaluation on SemEval2013 [2].

- Pre-hinting produces the greatest individual benefit
  Missing sense mappings are a problem.
- Plain ⟶ Progressive has almost the same improvement as Pre ⟶ Combined
  Missing senses and mis-ranked senses are likely separate issues

Future Work

- Universal Decompositional Semantics provides super-sense vectors
- Abstract Meaning Representation/Unscoped Logical Forms can provide sense bracketing and argument structures as further restrictions on chart entries.
- Discourse analysis can produce better global sense decisions

References

[1] James F Allen, Mary Swift, and Will De Beaumont. Deep semantic analysis of text. In *Proceedings of the 2008 Conference on Semantics in Text Processing*, pages 343–354. Association for Computational Linguistics, 2008.

[2] Roberto Navigli, David Jurgens, and Daniele Vannella. Semeval-2013 task 12: Multilingual word sense disambiguation. In *Second Joint Conference on Lexical and Computational Semantics (* SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, volume 2, pages 222–231, 2013.

[3] Simone Papandrea, Alessandro Raganato, and Claudio Delli Bovi. Supwsd: A flexible toolkit for supervised word sense disambiguation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 103–108, 2017.