

Demonstrating an Interactive Semantic Role Labeling System

Vasin Punyakanok Dan Roth Mark Sammons

Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA
{punyakan,danr,mssammon}@uiuc.edu

Wen-tau Yih

Microsoft Research
Redmond, WA 98052, USA
scottyih@microsoft.com

Abstract

Semantic Role Labeling (SRL) is the task of performing a shallow semantic analysis of text (i.e., *Who did What to Whom, When, Where, How*). This is a crucial step toward deeper understanding of text and has many immediate applications. Preprocessed information on text, mostly syntactic, has been shown to be important for SRL. Current research focuses on improving the performance assuming that this lower level information is given without any attention to the overall efficiency of the final system, although minimizing execution time is a necessity in order to support real world applications. The goal of our demonstration is to present an interactive SRL system that can be used both as a research and an educational tool. Its architecture is based on the state-of-the-art system (the top system in the 2005 CoNLL shared task), modified to process raw text through the addition of lower level processors, while achieving effective real time performance.

1 Introduction

Semantic parsing of sentences is believed to be an important subtask toward natural language understanding, and has immediate applications in tasks such information extraction and question answering.

We study *semantic role labeling (SRL)*, defined as follows: for each verb in a sentence, the goal is to

identify all constituents that fill a semantic role, and to determine their roles (such as Agent, Patient or Instrument) and their adjuncts (such as Locative, Temporal or Manner). The PropBank project (Kingsbury and Palmer, 2002), which provides a large human-annotated corpus of semantic verb-argument relations, has opened doors for researchers to apply machine learning techniques to this task.

The focus of the research has been on improving the performance of the SRL system by using, in addition to raw text, various syntactic and semantic information, e.g. *Part of Speech (POS) tags, chunks, clauses, syntactic parse tree, and named entities*, which is found crucial to the SRL system (Punyakanok et al., 2005).

In order to support a real world application such as an interactive question-answering system, the ability of an SRL system to analyze text in real time is a necessity. However, in previous research, the overall efficiency of the SRL system has not been considered. At best, the efficiency of an SRL system may be reported in an experiment assuming that all the necessary information has already been provided, which is not realistic. A real world scenario requires the SRL system to perform all necessary preprocessing steps in real time. The overall efficiency of SRL systems that include the preprocessors is not known.

Our demonstration aims to address this issue. We present an interactive system that performs the SRL task from raw text in real time. Its architecture is based on the top system in the 2005 CoNLL shared task (Koomen et al., 2005), modified to process raw text using lower level processors but maintaining

good real time performance.

2 The SRL System Architecture

Our system begins preprocessing raw text by using sentence segmentation tools (available at <http://l2r.cs.uiuc.edu/~cogcomp/tools.php>). Next, sentences are analyzed by a state-of-the-art syntactic parser (Charniak, 2000) the output of which provides useful information for the main SRL module.

The main SRL module consists of four stages: *pruning*, *argument identification*, *argument classification*, and *inference*. The following is the overview of these four stages. Details of them can be found in (Koomen et al., 2005).

Pruning The goal of pruning is to filter out unlikely argument candidates using simple heuristic rules. Only the constituents in the parse tree are considered as argument candidates. In addition, our system exploits a heuristic modified from that introduced by (Xue and Palmer, 2004) to filter out very unlikely constituents.

Argument Identification The argument identification stage uses binary classification to identify whether a candidate is an argument or not. We train and apply the binary classifiers on the constituents supplied by the pruning stage.

Argument Classification This stage assigns the final argument labels to the argument candidates supplied from the previous stage. A multi-class classifier is trained to classify the types of the arguments supplied by the argument identification stage.

Inference The purpose of this stage is to incorporate some prior linguistic and structural knowledge, such as “arguments do not overlap” and “each verb takes at most one argument of each type.” This knowledge is used to resolve any inconsistencies in argument classification in order to generate legitimate final predictions. The process is formulated as an integer linear programming problem that takes as input confidence values for each argument type supplied by the argument classifier for each constituent, and outputs the optimal solution subject to the constraints that encode the domain knowledge.

The system in this demonstration, however, differs from its original version in several aspects.

First, all syntactic information is extracted from the output of the full parser, where the original version used different information obtained from different processors. Second, the named-entity information is discarded. Finally, no combination of different parse tree outputs is performed. These alterations aim to enhance the efficiency of the system while maintaining strong performance.

Currently the system runs at the average speed of 1.25 seconds/predicate. Its performance is 77.88 and 65.87 F1-score on WSJ and Brown test sets (Carreras and Màrquez, 2005) while the original system achieves 77.11 and 65.6 on the same test sets without the combination of multiple parser outputs and 79.44 and 67.75 with the combination.

3 Goal of Demonstration

The goal of the demonstration is to present the system’s ability to perform the SRL task on raw text in real time. An interactive interface allows users to input free form text and to receive the SRL analysis from our system. This demonstration can be found at <http://l2r.cs.uiuc.edu/~cogcomp/srl-demo.php>.

Acknowledgments

We are grateful to Dash Optimization for the free academic use of Xpress-MP. This research is supported by ARDA’s AQUAINT Program, DOI’s Reflex program, and an ONR MURI Award.

References

- X. Carreras and L. Màrquez. 2005. Introduction to the conll-2005 shared tasks: Semantic role labeling. In *Proc. of CoNLL-2005*.
- E. Charniak. 2000. A maximum-entropy-inspired parser. In *Proc. of NAACL 2000*.
- P. Kingsbury and M. Palmer. 2002. From Treebank to PropBank. In *Proc. of LREC-2002*, Spain.
- P. Koomen, V. Punyakanok, D. Roth, and W. Yih. 2005. Generalized Inference with Multiple Semantic Role Labeling Systems. In *Proceedings of CoNLL-2005*.
- V. Punyakanok, D. Roth, and W. Yih. 2005. The necessity of syntactic parsing for semantic role labeling. In *Proc. of IJCAI-2005*.
- N. Xue and M. Palmer. 2004. Calibrating features for semantic role labeling. In *Proc. of the EMNLP-2004*.