An Evaluation of Syntactic Dependency Parsers on Clinical Data

Stuart J. Taylor, BS[1], Travis R. Goodwin, MS[1], Sanda M. Harabagiu, PhD[1]
[1]University of Texas at Dallas, Richardson, Texas

Introduction

With electronic medical records becoming more widely adopted, the potential to extract information from the narratives in medical records is increasing. In clinical decision support systems, quality control, medication reconciliation, and other biomedical applications[1], Natural Language Processing (NLP) techniques can be used to automatically discern knowledge from text. The accuracy of NLP techniques often heavily depends on the type of data used to train or develop them. Advanced NLP techniques – like semantic role labeling – require accurate *syntactic parsing*. Syntactic parsing reveals the syntactic structures (e.g. trees) in which words participate and discovers the syntactic roles of the words in a sentence. In this work we evaluate whether existing annotated clinical text is sufficient to train state-of-the-art neural parsing methods.

Methodology

We used the MiPACQ[1] clinical text corpus with a random 17:1:2 split for training, development, and testing. We trained and evaluated two neural parsers: (1) Stanford's greedy transition-based parser[2]; (2) SyntaxNet[5], Google's globally normalized transition-based parser; and two non-neural parsers: (3) OpenNLP's[*] chunking shift-reduce parser; (4) ClearParser[3], a shift-pop transition-based dependency parser used by cTAKES[4]. Additionally, we evaluated the performance of each parser on clinical text when using their general (i.e., non-clinical) pre-trained models.

Results

We computed the *Unlabeled Attachment Score* (UAS) on the clinical test data to perform our evaluations. The UAS is the percentage of words which have the correct *syntactic head*. Syntactic heads are used to encode parent-child relations in a parsing constituent (e.g. in "stimulation increases flow" the noun "stimulation" is the child of the verb "increases" because the verb is the head of the phrase and there is a NOMINAL SUBJECT relation between the head and the noun). The results on the test set are shown in Table 1. The UAS-G column shows the UAS score for each parser when using the pre-trained general-purpose model, and the UAS-C column shows the UAS score after training each parser with the clinical training data.

Table 1: Evaluation scores on MiPACQ testing data.

Parser	UAS-G	UAS-C
Stanford	**78.47%**	80.66%
OpenNLP	74.04%	81.20%
ClearParser	39.11%	83.08%
SyntaxNet	63.76%	**85.19%**

Conclusion

While SyntaxNet performed the best after training on clinical text, all parsers showed improvement. These results suggest that the MiPACQ corpus is sufficient to train neural syntactic parsers. Interestingly, the parsers which performed the best with the pre-trained model performed the worst after being trained on clinical text. Directions for future research include: (1) tuning hyper-parameters; (2) incorporating more training data; and (3) evaluating on additional types of clinical data.

References

1. Albright D, Lanfranchi A, Fredriksen A, Styler WF, Warner C, Hwang JD, Choi JD, Dligach D, Nielsen RD, Martin J, Ward W. Towards comprehensive syntactic and semantic annotations of the clinical narrative. Journal of the American Medical Informatics Association. 2013 Sep 1;20(5):922-30.
2. Chen D, Manning CD. A Fast and Accurate Dependency Parser using Neural Networks. InEMNLP 2014 Oct 25 (pp. 740-750).
3. Choi JD, Palmer M. Getting the most out of transition-based dependency parsing. InProceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers-Volume 2 2011 Jun 19 (pp. 687-692). Association for Computational Linguistics.
4. Savova GK, Masanz JJ, Ogren PV, Zheng J, Sohn S, Kipper-Schuler KC, Chute CG. Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications. Journal of the American Medical Informatics Association. 2010 Sep 1;17(5):507-13.
5. Andor D, Alberti C, Weiss D, Severyn A, Presta A, Ganchev K, et al. Globally Normalized Transition-Based Neural Networks. Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). 2016 Aug.

[*] `https://opennlp.apache.org/`