CCG Parsing and Multiword Expressions

Miryam de Lhoneux
Crowdynews, The University of Edinburgh
COST ACTION IC1207 PARSEME Meeting, Malta, WG2 and WG4

Introduction

Foreword This poster presents work carried out for my MSc dissertation (de Lhoneux, 2014) at the University of Edinburgh under the supervision of Mark Steedman and Omri Abend.

Objectives To find out whether or not information about Multiword Expressions (MWEs) can improve statistical parsing with Combinatory Categorial Grammar (CCG).

Motivation

- Information about MWEs is useful for syntactic parsing.
- Syntactic parsing is central to NLP.
- MWEs forced a non-modular view of grammar in Linguistics theory (e.g. Construction Grammar): Syntax and the Lexicon are not entirely modular.
- CCG adopts a grammar architecture in which syntactic information is partly encoded in the Lexicon.

Background studies

Nivre and Nilsson (2004) manually created two versions of a Treebank, one in which MWE units are joined to form a token (commonly called the 'words-with-spaces' approach) and one in which they are separate. They tested whether this 'perfect MWE recognition' could help parsing accuracy. Korkontzelos and Manandhar (2010) automatically created two versions of an unannotated corpus based on a list of MWEs randomly selected. They observed a gain in parsing accuracy when the *test* data contained MWEs joined as one token. Both studies limited the types of MWEs dealt with. Two questions remained unanswered:

- Can parsing benefit from MWE information obtained by automatic MWE recognition?
- Can the representation of MWEs as one unit in a parsing model improve the parsing model when used with other MWE types?

Methodology

Data CCGbank a translation of the PTB into CCG derivations. Sections 01-22 : training; 00 : development; 23: testing.

MWE recognition Use of the jMWE library (Finlayson and Kulkarni, 2011) to detect MWEs in a sentence. Experiments with library tools.

Treebank conversion

Collapsing trees

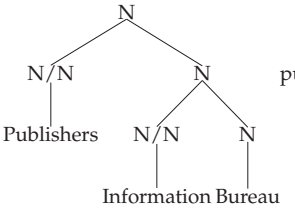

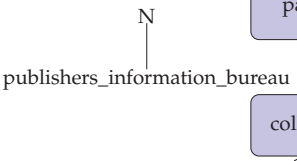

Collapsing dependencies

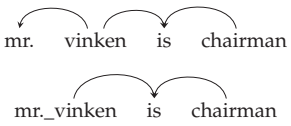

Comparing models

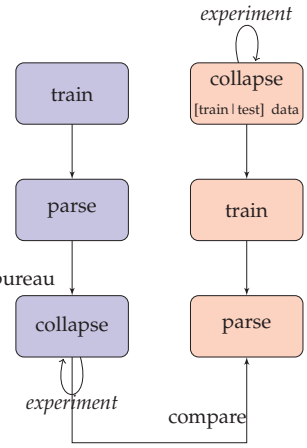

Figure 1: Pipeline

Terminology

Unchanged = A, changed = B.
Training effect = parser learns something useful.
Parsing effect = collapsed test data help parser.

Evaluation Precision (P), recall (R) and F1 (F_1) of unlabelled dependencies against gold standard.
Is there a training effect? Compare outA with outB (with training and test data changed) on goldB.
Is there a parsing effect? Compare outA with outB (with test data changed) on goldB.
Does recognition method influence the results? Experiment with different MWE recognition methods and compare outBs decollapsed on goldA.

Results

Data collapsed	P	R	F_1
test	**84.53**	84.76	84.64
training and test	84.48	**85.28**	**84.88**

Table 1: Training effect

Data collapsed	P	R	F_1
test (before parsing)	**79.83**	79.54	**79.69**
test (after parsing)	79.38	**79.60**	79.49

Table 2: Parsing effect

Data collapsed	decollapsed	MWE types handled	F_1
None	None		85.15
training and test	out	all	85.02
training and test	out	Proper Nouns	**85.28**
training and test	out	Length 2	85.07
training and test	out	Stop words	85.19

Table 3: MWE recognition experimentation

Due to a shortcoming in the methodology (the algorithm we used to modify MWEs in the Treebank is only capable of dealing with MWEs that do not cross constituent boundaries), adjustments had to be made to obtain a fair comparison of the models (the differences between Table 1 and Table 2 are not meaningful). The tables are representative of the tendencies found. Training (Table 1) and parsing effects (Table 2) were found with changed models performing slightly but significantly ($p>.05$) better than unchanged models. Differences in results obtained with different recognizers are also small but significant (Table 3).

Conclusion

The main contributions of our work are:

- Improvements on CCG parsing with automatic MWE recognition
- Significant results despite limited settings
- Encouraging results on a hard task
- An algorithm to automatically collapse MWEs in a Treebank
- Techniques for distinguishing training from parsing effects
- Empirical support that there is both training and parsing effects
- Interesting differences in results when using different recognizers

Future Research

- Extending the collapsing algorithm to the non-sibling case
- Testing more MWE recognition methods with more data
- Further integrating MWE recognition and syntactic parsing
- Conducting error analysis

Contact Information

Web http://linkd.in/1ybG9kI

Email miryam@crowdynews.com

Phone +31 616 22 62 10

Address Crowdynews
Friesestraatweg, 213B
Groningen, 9743 AD
The Netherlands

References

de Lhoneux, M. (2014). CCG Parsing and Multiword Expressions. Master's thesis, The University of Edinburgh.

Finlayson, M. A. and Kulkarni, N. (2011). Detecting Multiword Expressions Improves Word Sense Disambiguation. In *Proceedings of the Workshop on Multiword Expressions: From Parsing and Generation to the Real World*, MWE '11, pages 20–24, Stroudsburg, PA, USA. Association for Computational Linguistics.

Korkontzelos, I. and Manandhar, S. (2010). Can recognising multiword expressions improve shallow parsing? In *HLT-NAACL*, pages 636–644.

Nivre, J. and Nilsson, J. (2004). Multiword units in syntactic parsing. In *Workshop on Methodologies and Evaluation of Multiword Units in Real-World Applications*, pages 39–46.