

Semantic Case Role Detection for Information Extraction

Rik De Busser, Roxana Angheluta and Marie-Francine Moens



Introduction

Very roughly, information extraction could be defined as a discipline in which it is attempted to extract semantically relevant elements from a text, using only shallow analysis. This has very often much to do with identifying semantic case roles, i.e. with the detection of the events, states, and their participants as they are mentioned in a text.

Unfortunately, case role detection as a goal in itself has very often been treated in a rather trivial way. Using notions from systemic-functional linguistics, we will try to build a model for extracting generic semantic case roles, i.e. case roles that are not specialized to any particular domain. These roles are learned from a tagged and hand-annotated corpus and can be reassigned to previously unseen text.

Our research will focus on two major questions:

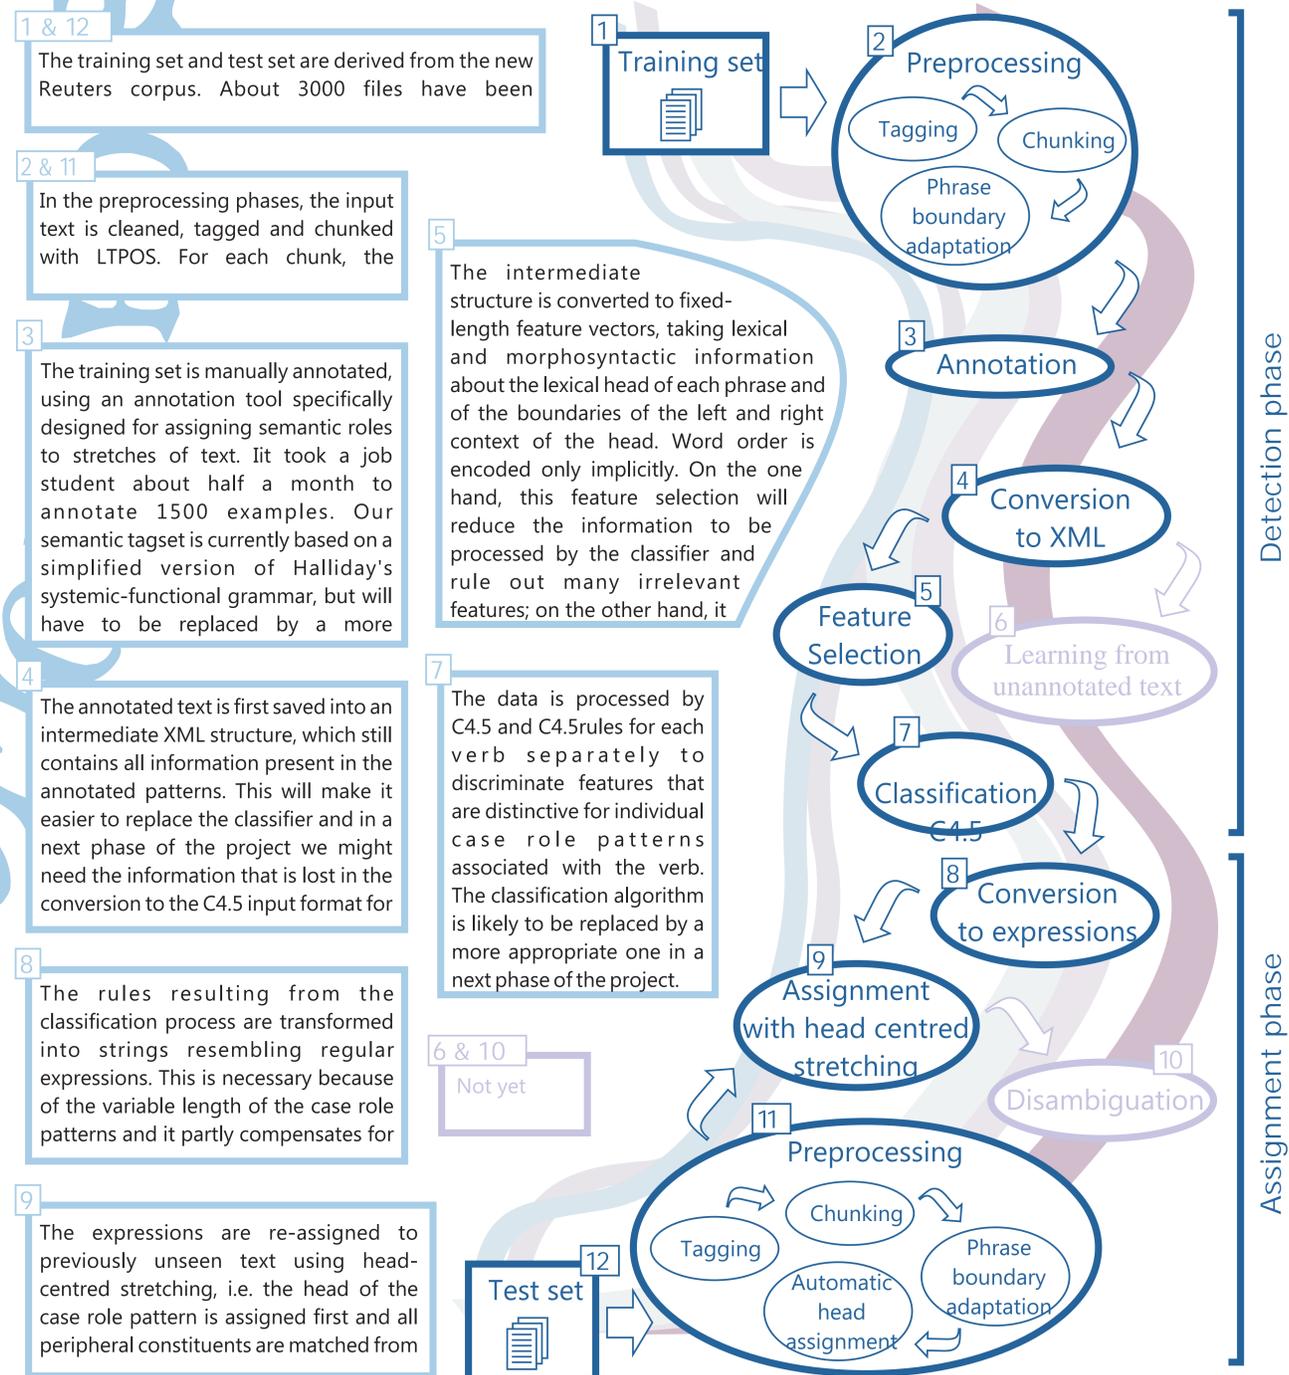
- How shallow can you be? To what extent is there a relationship between the generic semantic plane and morphosyntactic and lexical surface phenomena? What is the minimal amount of information you need?
- Will a shallow approach be useful for information extraction? Can the mapping between morphosyntactic properties and functional-semantic case roles be unambiguously determined?

Case role detection will provide a way to identify individual semantic entities in a text. For instance, the sentence 'Invesco in merger talks with AIM Management' will be labelled with the following roles:

Roles like the ones above constitute the semantic tagset. In the case of our algorithm, they are based on a simplified version of Halliday's systemic-functional grammar. The main advantage of our approach over previous, domain-specific

Overview

We consider case role detection to be a standard classification task. Each class is a concatenation of individual case roles into a meaningful pattern; the features that are used for training the classifier are lexical and morphosyntactic properties of the chunks corresponding to the case roles. Distinctive features that are



Evaluation

Future improvements

The current setup is only a first try and therefore it should not be a surprise that initial results are far from impressive. However, many improvements are still possible:

- Replacing C4.5 by a more adequate classifier will boost results and is likely to reduce the annotation effort.
- Encoding relative word order as an explicit feature will make it easier to match corresponding case roles and will enable us to treat circumstances (which do not have a fixed position) more accurately.
- Using analytic tags instead of synthetic tags will improve rule generalization.
- Lexicons might also be employed to match partially correct patterns, to deal with idiomatic expressions and to use...
- Analogous reasoning to match partially correct patterns and to expand the pattern base.
- Designing a preference order for empty roles will match individual chunks correctly.
- Pattern selection, based on parameters such as pattern length,

We would like to thank IWT Flanders, Roularta Media Group, Kluwer Belgium, Land & Computing and ICMS for the research funding and