

Specialized rankers and global models for coreference resolution

Coreference resolution is the task of mapping linguistic expressions to the discourse entities they evoke, e.g. determining whether a textual mention like "the man" refers to the same entity evoked by another mention like "John". It is an important aspect of natural language understanding that has great relevance for practical applications such as information retrieval and text summarization. Great progress has been made in coreference resolution through the use of machine learning techniques, but state-of-the-art performance still leaves much room for improvement. Even though a large part of the performance bottleneck stems from the need for (currently out-of-reach) deep understanding and reasoning about the content of the texts, there are significant opportunities to provide better models and use richer information sources for the task.

In this talk, I will discuss two strategies for overcoming deficiencies of previous approaches. The first strategy is to use specialized ranking models that target specific types of referential expressions and that are a better fit for the task than more commonly used classification models. The second is to use integer linear programming to create joint, global models that assume less independence between individual coreference decisions and that can cleanly integrate multiple information sources---such as discourse status and named-entity classification---with coreference determination. Both of these strategies lead to significant performance improvements, as measured according to three different metrics, and they open the way toward augmenting systems with further information, such as discourse structure. A running sub-theme of this talk will be the importance of evaluating coreference resolution systems with multiple scoring metrics.

This talk represents joint work with Jason Baldridge (UT Austin).