

A Discriminative Model for Tree-to-Tree Translation

Michael Collins,
MIT

(Joint work with Brooke Cowan and Ivona Kucerova)

Abstract

In this talk I will describe a framework for tree-to-tree based statistical translation. Our goal is to learn a model that maps parse trees in the source language to parse trees in the target language. Our translation framework involves a process where the target-language parse tree is broken down into a sequence of clauses, and each clause is then translated separately. A central concept we introduce in the translation of clauses is that of an **aligned extended projection** (AEP). AEPs are derived from the concept of an extended projection in lexicalized tree adjoining grammars (LTAG) (Frank, 2002), with the addition of alignment information that is based on work in synchronous LTAG. A key contribution of this work is a method for learning to map German clauses to AEPs using a feature-based model with a perceptron learning algorithm. We describe experiments on translation from German to English on the Europarl data set. Evaluation in terms of both BLEU scores and human judgments shows that our system performs similarly to the phrase-based model of Koehn et al, 2003.

BioSketch

My research interests are in natural language processing, and machine learning. I completed a PhD in computer science from the University of Pennsylvania in December 1998. From January 1999 to November 2002 I was a researcher at AT&T Labs-Research. I joined MIT in January 2003.