Learning Language from Perceptual Context

Raymond Mooney
University of Texas at Austin
mooney@cs.utexas.edu

Abstract

Machine learning has become the dominant approach to building natural-language processing systems. However, current approaches generally require a great deal of laboriously constructed human-annotated training data. Ideally, a computer would be able to acquire language like a child by being exposed to linguistic input in the context of a relevant but ambiguous perceptual environment. As a step in this direction, we have developed systems that learn to sportscast simulated robot soccer games and to follow navigation instructions in virtual environments by simply observing sample human linguistic behavior in context. This work builds on our earlier work on supervised learning of semantic parsers that map natural language into a formal meaning representation. In order to apply such methods to learning from observation, we have developed methods that estimate the meaning of sentences given just their ambiguous perceptual context.

Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, page 602,
Avignon, France, April 23 - 27 2012.