

Learning Language from Its Perceptual Context

Raymond J. Mooney

University of Texas at Austin, USA

Current systems that learn to process natural language require 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 present a system that learns to sportscast simulated robot soccer games by example. The training data consists of textual human commentaries on Robocup simulation games. A set of possible alternative meanings for each comment is automatically constructed from game event traces. Our previously developed systems for learning to parse and generate natural language (KRISP and WASP) were augmented to learn from this data and then commentate novel games. The system is evaluated based on its ability to parse sentences into correct meanings and generate accurate descriptions of game events. Human evaluation was also conducted on the overall quality of the generated sportscasts and compared to human-generated commentaries.