Modeling word learning as communicative inference

Michael C. Frank

Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139
mcfrank@mit.edu

Abstract

How do children learn their first words? I describe a model that makes joint inferences about what speakers are trying to talk about and the meanings of the words they use. This model provides a principled framework for integrating a wide variety of non-linguistic information sources into the process of word learning.

Talk Précis

How do children learn their first words? Much work in this field has focused on the social aspects of word learning: that children make use of speakers' intentions—as signaled by a wide range of non-linguistic cues such as their eye-gaze, what they are pointing at, or even what referents are new to them—to infer the meanings of words (Bloom, 2002). However, recent evidence has suggested that adults and children are able to learn words simply from the consistent co-occurrence of words and their referents, even across otherwise ambiguous situations and without explicit social cues as to which referent is being talked about (Yu & Smith 2007; Smith & Yu, 2008).

In this talk I describe work aiming to combine these two sets of evidence within a single probabilistic framework (Frank, Goodman, & Tenenbaum, 2009). We propose a model in which learners attempt to infer speakers' moment-to-moment communicative intentions jointly with the meanings of the words they have used to express these intentions. This process of joint inference allows our model to

explain away two major sources of noise in simpler statistical word learning proposals: the fact that speakers do not talk about every referent and that not all words that speakers utter are referential.

We find that our model outperforms associative models in learning words accurately from natural corpus data and is able to fit children's behavior in a number of experimental results from developmental psychology. In addition, we have used this basic framework to begin investigating how learners use the rich variety of non-linguistic information signaling speakers' intentions in service of word learning. As an example of this work, I will describe an extension of the model to use discourse continuity as a cue for speakers' intentions.

Acknowledgments

This work supported by a Jacob Javits Graduate Fellowship and NSF Doctoral Dissertation Research Improvement Grant #0746251.

References

- Paul Bloom. 2002. *How Children Learn the Meanings of Words*. Cambridge, MA: MIT Press.
- Michael C. Frank, Noah D. Goodman, and Joshua B. Tenenbaum. 2009. Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*.
- Linda Smith and Chen Yu. 2008. Infants rapidly learn word-referent mappings via cross-situational statistics. *Cognition*, 106, 1558-1568.
- Chen Yu and Linda Smith. 2007. Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18, 414-420.