A Bayesian Model of Pronoun Production and Interpretation

Andrew Kehler
University of California, San Diego

Abstract

A standard assumption in linguistic, psycholinguistic, and computational research on pronoun use is that there is a unified notion of entity prominence that mediates between interpretation and production. Kehler et al. (2008) and Kehler & Rohde (2013) have argued instead for a Bayesian model, one in which pronoun production is insensitive to a class of semantically- and pragmatically-driven contextual biases that have been shown to influence pronoun interpretation. We begin by surveying a series of experiments that led to the formulation of the model, followed by more recent studies designed to evaluate the predictions it makes. The results support the claim that pronoun interpretation biases, but not production biases, are sensitive to a set of semantic and pragmatic factors, revealing precisely the asymmetry predicted by the Bayesian analysis. Correlation analyses further establish that the model provides better estimates of measured pronoun interpretation biases than two prominent competing models from the literature. The results open up new possibilities for utilizing unlabeled data to improve the state-of-the-art in computational work.