Michael Jones - Abstract

Growing Semantic Representations: The Emergence of Semantic Structure in Continuous Vector Space Models

Meaning develops over a lifetime of linguistic and perceptual experience. Laboratory experiments can probe the end product of the process of semantic abstraction, or can study the process at a small scale using well-controlled stimuli. But a full understanding of the mechanisms that drive semantic abstraction requires the study of models that learn over realistic data at a scale that humans do. Indeed, the models that perform best at small scales do not readily scale up to human-scale amounts of data, and relatively "dumb but scalable" models end up generating impressively complex representations when trained on sufficient data. I will present some recent work from my lab exploring vector accumulation models based on theories of associative memory and reinforcement learning. These models learn incrementally from linguistic corpora, making surprising predictions about semantic development, and can also integrate large-scale perceptual data from our crowdsourcing project, the *NSF Semantic Pictionary Project*.