

A Two-Process Model of Semantic Development

Philip Huebner

University of Illinois, Urbana-Champaign, Urbana, Illinois, United States

Jon Willits

University of Illinois at Urbana-Champaign, Champaign, Illinois, United States

Abstract

How do children acquire semantic knowledge? In this work, we explore an old answer to this question: Semantic development is a hybrid of two distinct processes. The first process involves unsupervised learning of relations between objects, providing a representation of objects that is useful for a wide range of possible goals. The second process involves explicitly learning to put objects and their relations into categories. Critically, this second process uses the representations of the first process as its starting point. Here, we demonstrate this using a two-process model, where the first process is a distributional semantic model (e.g. HAL, Word2Vec, RNN), and the second process is a transformation of representations learned during process 1 into a task-specific target space. This approach improves performance on multiple semantic tasks, compared to using the representations learned by process 1 directly. We believe this model demonstrates that a task- or goal-oriented perspective of semantic cognition has promise for furthering our understanding of semantic development.