

Extending the Rogers and McClelland Model of Semantic Cognition (2003) to work with Raw Pixel Information

Arihant Jain

New York University, New York City, New York, United States

Brenden Lake

NYU, New York, New York, United States

Todd Gureckis

New York University, New York, New York, United States

Abstract

Understanding how we acquire semantic knowledge is a central topic in cognitive science. In a now classic paper, Rogers and McClelland (2003) explored how a parallel distributed processing (PDP) model could recreate several important phenomena in semantic memory including how concepts are acquired, lost due to semantic dementia, and become organized hierarchically. One well known limitation of this model, which was acknowledged by the original authors, is that the features used in the model were largely hand coded. In this project we revisit this classic PDP account in light of modern advances in neural network techniques. In particular, we show that we can recreate several of the predictions of the Rogers and McClelland (2003) model in a network trained directly on raw pixel information from category exemplars. These results add realism to the original model while also showing how the principles of the model generalize to higher dimensional input spaces.