

What Do Computers Know About Semantics Anyway? Testing Distributional Semantics Models Against a Broad Range of Relatedness Ratings

Kevin Brown

Oregon State University, Corvallis, Oregon, United States

Eiling Yee

University of Connecticut, Storrs, Connecticut, United States

Elliot Saltzman

Boston University, Boston, Massachusetts, United States

James Magnuson

University of Connecticut, Storrs, Connecticut, United States

Ken McRae

University of Western Ontario, London, Ontario, Canada

Abstract

Distributional Semantics Models (DSMs) are a primary method for distilling semantic information from corpora. However, a key question remains: What types of semantic relations do DSMs detect? Prior work has addressed this question using a limited set of ratings that typically are either amorphous (association norms) or restricted to semantic similarity (SimLex, SimVerb). We tested four DSMs (SkipGram, CBOW, GloVe, PPMI) using multiple hyperparameters on a theoretically-motivated, rich set of relations involving words from multiple syntactic classes spanning the abstract-concrete continuum (21 sets of ratings). Results show wide variation in the DSMs' ability to account for the ratings, and that hyperparameter tuning buys comparatively little for improving correlations. For CBOW and SkipGram, we included word and context embeddings. For SkipGram, there was a marked improvement in simulating the human data by averaging them. Our results yield important insights into the types of semantic relations that are captured by DSMs.