Evaluating systematicity in neural networks with natural language inference

Emily Goodwin
McGill University, Montreal, Quebec, Canada

Koustuv Sinha
McGill University, Montreal, Quebec, Canada

Timothy ODonnell
McGill University, Montreal, Quebec, Canada

Abstract

Compositionality makes linguistic creativity possible. By combining words, we can express uncountably many thoughts; by learning new words, we can extend the system and express a vast number of new thoughts. Recently, a number of studies have questioned the ability of neural networks to generalize compositionally (Dasgupta, Guo, Gershman & Goodman, 2018). We extend this line of work by systematically investigating the way in which these systems generalize novel words.

In the setting of a simple system for natural language inference, natural logic (McCartney & Manning, 2007), we systematically explore the generalization capabilities of various neural network architectures. We identify several key properties of a compositional system, and develop metrics to test them. We show that these architectures do not generalize in human-like ways, lacking inductive leaps characteristic of human learning.