A tradeoff between generalization and perceptual capacity in recurrent neural networks

Taylor Webb
Princeton University, Princeton, New Jersey, United States

Steven Frankland
Princeton University, Princeton, New Jersey, United States

Simon Segert
Princeton University, Princeton, New Jersey, United States

Alexander Petrov
Ohio State University, Columbus, Ohio, United States

Randall O'Reilly
University of Colorado Boulder, Boulder, Colorado, United States

Jonathan Cohen
Princeton University, Princeton, New Jersey, United States

Abstract

In a classic paper, Miller (1956) summarized findings showing that people can only identify a limited number of distinct stimuli at a time. One puzzling aspect of this capacity limitation is that it is approximately invariant to range. That is, the number of accurately identifiable stimuli is approximately the same regardless of how far apart the stimuli are spaced. Models of this phenomenon have suggested that people operate in a context-coding mode when performing these tasks, effectively carrying out a form of contextual normalization, but why such normalization might take place is unclear. Here, we propose an explanation by appealing to a tradeoff with generalization. Specifically, we implement contextual normalization in a recurrent neural network and show that this normalization enables stronger generalization in a relational reasoning task, but also results in a perceptual capacity limitation which captures many of these classic phenomena.