Controlling the Dynamics of Recurrent Neural Networks by Maximizing Action-Path Entropy

Chiara Mastrogiuseppe
Department of Communication and Information Technologies
Universitat Pompeu Fabra
Barcelona, Spain
chiara.mastrogiuseppe@upf.edu

Rubén Moreno-Bote
Department of Communication and Information Technologies,
Serra Húnter Fellow Programme
Universitat Pompeu Fabra
Barcelona, Spain
ruben.moreno@upf.edu

Abstract

Cortical activity is characterized by highly variable and complex pattern of activation. Recurrent neural networks (RNNs) have proven to be a useful framework to model neural activity, showing chaotic activity for a wide range of the parameters. A key ingredient is a saturating non-linearity activation of the units, which prevents the explosion of the activities of the RNN while making all units continuously reach saturation. Saturation is not a realistic feature of sound activity, as it limits information transmission. Indeed, neurons have proven to show adaptation of the tuning curve to exploit the encoding. Here, we study how to prevent a network to reach saturation while showing the observed neural variability. We model a reinforcement learning agent acting on a RNN through external actions (i.e. currents). Its only goal is to maximize the entropy of the action sequences (path-action entropy). By defining activity saturation as terminal states where actions cannot be longer make, we find that the agent learns how to avoid them while generating variability. The algorithm generalizes well to networks with non-saturating activation functions, showing a highly variable activity confined within arbitrary boundaries. Our work provides a novel theory of neural variability, as well as a framework to study collective dynamics in brain disorders and motor disfunctions.

Preprint. Under review.