

Learning to reinforcement learn

Jane Wang

DeepMind

Zeb Kurth-Nelson

DeepMind, Max Planck-UCL Centre for Computational Psychiatry

Hubert Soyer

DeepMind

Joel Leibo

DeepMind

Dhruva Tirumala

DeepMind

Remi Munos

DeepMind

Charles Blundell

DeepMind

Dharshan Kumaran

DeepMind, Institute of Cognitive Neuroscience, UCL

Matt Botvinick

DeepMind, Gatsby Computational Neuroscience Unit, UCL

Abstract: In recent years deep reinforcement learning (RL) systems have attained superhuman performance in a number of challenging task domains, but are constrained by a demand for large training sets. A critical present objective is thus to develop deep RL methods that can adapt rapidly to new tasks. In the present work we introduce a novel approach to this challenge, which we refer to as deep meta-reinforcement learning. Previous work has shown that recurrent networks can support meta-learning in a fully supervised context. We extend this approach to the RL setting. What emerges is a system that is trained using one RL algorithm, but whose recurrent dynamics implement a second, quite separate RL procedure. This second, learned RL algorithm can differ from the original one in arbitrary ways and exploit structure in the training domain. We unpack these points in five proof-of-concept experiments to examine key aspects of deep meta-RL.