Symposia

[3S08m]The Annual Meeting - Organized Symposium

Interplays between theory and experiment in decision

neuroscience

Organizer: Naoshige Uchida (Harvard University) Fri. Jul 31, 2020 9:00 AM - 11:00 AM Room 8 Support contributed by Tianqiao and Chrissy Chen Institute

*Videos are available throughout the meeting period.

[3S08m-02]A distributional code for value in dopamine-based reinforcement learning

*Will Dabney¹, Zebulun Lloyd Kurth-Nelson¹, Naoshige Uchida², Clara K Starkweather², Demis Hassabis¹, Rémi Munos¹, Matthew Botvinick¹ (1. DeepMind, 2. Center for Brain Science, Department of Molecular and Cellular Biology, Harvard University)

Since its introduction, the reward prediction error (RPE) theory of dopamine has explained a wealth of empirical phenomena, providing a unifying framework for understanding the representation of reward and value in the brain. According to the now canonical theory, reward predictions are represented as a single scalar quantity, which supports learning about the expectation, or mean, of stochastic outcomes. In the present work, we propose a novel account of dopamine-based reinforcement learning. Inspired by recent artificial intelligence research on distributional reinforcement learning, we hypothesized that the brain represents possible future rewards not as a single mean, but instead as a probability distribution, effectively representing multiple future outcomes simultaneously and in parallel. This idea leads immediately to a set of empirical predictions, which we tested using single-unit recordings from mouse ventral tegmental area. Our findings provide strong evidence for a neural realization of distributional reinforcement learning.

The RPE theory of dopamine derives from work in the artificial intelligence (AI) field of reinforcement learning (RL). Since the link to neuroscience was first made, however, RL has made substantial advances, revealing factors that radically enhance the effectiveness of RL algorithms. In some cases, the relevant mechanisms invite comparison with neural function, suggesting new hypotheses concerning reward-based learning in the brain. Here, we examine one particularly promising recent development in AI research and investigate its potential neural correlates. Specifically, we consider a computational framework referred to as distributional reinforcement learning.