University of Zurich UZH

Department of Economics – Neuroeconomics Seminar

December 12, 2019 - 17:00 - 18:00

Matthew Botvinick
DeepMind

A distributional code for value in dopamine-based reinforcement learning

Twenty years ago, a link was discovered between the neurotransmitter dopamine and the computational framework of reinforcement learning. Since then, it has become well established that dopamine release reflects a reward prediction error, a surprise signal that drives learning of reward predictions and shapes future behavior. 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. I'll present recent work in which we have proposed a novel account of dopamine-based reinforcement learning, and adduced experimental results which point to a significant modification of the standard reward prediction error theory. 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.