University of Zurich[UZH]

Department of Economics – Neuroeconomics Seminar

October 1st, 2026 - 17:00 - 18:00

Maria Eckstein
Google DeepMind, San Francisco

Understanding Human Reward Learning Using Data-Driven Models and Procedurally-Generated Task Spaces

Classical work in the study of human reward-based learning involves designing an experimental task---often inspired by the Reinforcement Learning (RL) framework---and fitting a small set of computational models---often inspired by RL algorithms---to that dataset. For example, researchers often model human behavior on bandit tasks using variants of Q-learning. While this approach has been highly productive, including landmark discoveries like dopamine reward prediction errors, it also has limitations. This talk focuses on the lack of generalizability of classical, hand-crafted RL models: Models derived from the one-task-one-model paradigm often predict behavior on held-out tasks rather poorly. I argue that this lack of generalizability is a fundamental problem for the cognitive sciences: intuitively, models should be robust to superficial task differences such as variations in the numbers of choice options, presence of states, or features of the reward function. I will show recent work that addresses this problem from two angles: modeling and data. Regarding models, I will show how we moved beyond the limitations of classic hand-crafted models by using flexible, data-driven methods, which allowed us to compare classes of Q-learning models instead of individual model instances, allowing us to exhaustively cover a space of possible models, and efficiently innovate cognitive mechanisms. Regarding data, I will show how we moved from single learning tasks to a comprehensive task space encompassing most existing paradigms in the reward-based learning literature, while closing the gaps between tasks in a near-continuous fashion. Our results suggest that more general models in conjunction with broader datasets can pave the road toward increasingly general models of human reward-based learning and decision making, and surface novel cognitive models that show a striking departure from classic, hand-crafted RL-based ones.