

RL at Harvard 2023

August 29th, Harvard University main campus, Geological Lecture Hall

RL at Harvard is a one-day workshop on reinforcement learning broadly construed.

Our goal is to provide a forum for researchers interested in topics related to reinforcement learning—from building agents that beat humans at Atari games, to building robots that can efficiently navigate complex environments, to the neurobiology of dopamine and behavior—to learn about foundational concepts, discuss the latest advances, and make connections between disparate parts of the field.

This year, we will hear from world experts in reinforcement learning like Na Li, Mitsuko Watabe-Uchida, Stephanie Gil, and Samuel Gershman. Our hope is that the roster of speakers showcases the staggering variety of reinforcement-learning-related work that happens at Harvard, and helps strengthen and enrich the connections among Harvard's reinforcement learning researchers.

The latter half of the workshop will be dedicated to a symposium on the **neurobiology of causality and contingency**. As animals learn to act effectively in complex environments, they face the problem of learning **cause-and-effect relationships**: *if A happens, then B will happen*. They must also learn that some environmental transitions and rewards are **contingent** on other circumstances being true: *if A does not happen, then B will not happen*. New research from the Nambodiri Lab, Janak Lab, and Uchida Lab suggests that dopamine may provide a neurobiological mechanism for learning these relationships, although it is not completely clear how this role relates to other roles dopamine is thought to have in helping animals learn from experience.

With support from:



Organizing committee: John Vastola, Paul Masset, Jacob Zavatone-Veth, Kenneth Blum

Special thanks: Sam Gershman, Barbara Perlo, Dean Christopher Stubbs

Full Program (Talks will be held in Geological Lecture Hall):

8.15-8.45	Registration and coffee break (in Tozzler courtyard)
8.45-9	Introductory remarks
9-10.30	Session 1: Big picture RL (chair: John Vastola, HMS Neurobiology) Na Li (SEAS) Wilka Carvalho (Kempner) Adam Lowet (MCB) Mitsuko Watabe-Uchida (MCB)
10.30-11	Coffee break
11-12.30	Session 2: RL and applications (chair: Jacob Zavatone-Veth, Physics) Stephanie Gil (Computer Science) Changmin Yu (University College London - Gatsby Unit) Ann Huang (Program in Neuroscience) Agnese Seminara (University of Genoa)
12.30-1.30	Lunch
1.30-3.15	Session 3: Causality and contingency I (chair: Paul Masset, MCB) Vijay K. Namboodiri (UCSF - Neurology) Selina Qian (MCB) Eric Garr (Johns Hopkins University - Psychological and brain sciences)
3.15-3.45	Coffee break
3.45-5	Session 4: Causality and contingency II (moderator: Sam Gershman, Psychology) Introduction by Sam Gershman Debate/discussion: Gershman, Uchida, Namboodiri, Garr, Qian
5-6.30	Poster Session and drinks (Northwest Cafe and Patio)
6.30-8	Dinner + Poster Session (Northwest Cafe and Patio)

Poster Session (5 pm - 8 pm):

Jennifer J. Bussell, **Ryan P. Badman**, Christian D. Márton, Ethan S. Bromberg-Martin, Larry Abbott, Kanaka Rajan, Richard Axel

Representations of information value in the mouse orbitofrontal cortex during information seeking

Daniel Barabasi

Complex Computations from Developmental Priors

Aakash Basu, Jen-Hau Yang, Abigail Yu, Samira Glaeser-Khan, Jiesi Feng, John H. Krystal, Yulong Li, Alfred P. Kaye

Norepinephrine represents a threat prediction error under uncertainty

Léo Benac, Sonali Parbhoo, Finale Doshi-Velez

Bayesian Inverse Transition Learning for Offline Settings

Cheshta Bhatia, Kiah Hardcastle, Kevin Mizes, Lucy Lai, Sam Gershman, Bence P.O. Olveczky

Associative striatum as a substrate for sensory cue guided action selection

Ryan Boldi, Lee Spector

Lexicase Selection & Learning

Blake Bordelon, Paul Masset, Henry Kuo, Cengiz Pehlevan

Dynamics of Temporal Difference Reinforcement Learning

Mark Burrell, Naoshige Uchida

Testing the Temporal Difference Dopamine Hypothesis against ANCCR

Alexandra Decker, Steven L. Meisler, Nick Hubbard, Clemens Bauer, Julia Leonard, Andrea Imhoff, Melissa Giebler, Hannah Grotzinger, Yesi Camacho Torres, Rachel Romeo, and John D. E. Gabrieli

Lower socioeconomic status is associated with reduced striatal and behavioral sensitivity to reward in adolescence

Dan Dillon

Using Diffusion Modeling to Gain Insight into RL in Depression

Noémi Éltető and Peter Dayan

Habits of Mind: Reusing Action Sequences for Efficient Planning

Sam Hall-McMaster, Momchil Tomov, Nicolas W. Schuck* and Samuel J. Gershman*
Successor Features as Neural and Computational Building Blocks of Generalisation

M Ganesh Kumar, Cheston Tan, Camilo Libedinsky, Shih-Cheng Yen, Andrew Tan
One-shot learning of paired association navigation with biologically plausible schemas

Bo Liu
One nose but two nostrils: learning with sparse connection between two hemispheres aligns bilateral odor representations

Paul Masset, Pablo Tano, Athar N. Malik, HyungGoo R. Kim, Alexandre Pouget, Naoshige Uchida
Multi-timescale reinforcement learning in the brain

Haritosh Patel, Benjamin Bar-Ziv, **Ryan Rezai**, William Kim, Anna Shneidman, Joanna Aizenberg
Taught to Sniff: Reinforcement Learning to Bioinspired Sensing

Kimberly Reinhold, Marci Iadarola, Shi Tang, Bernardo Sabatini
Timescale of contribution of posterior dorsomedial striatum tail to trial and error learning


Tankred Saanum, Noemi Elteto, Peter Dayan, Marcel Binz, Eric Schulz
Reinforcement Learning with Simple Sequence Priors

Satpreet H. Singh, Floris van Breugel, Rajesh P. N. Rao & Bingni W. Brunton
Emergent behaviour and neural dynamics in artificial agents tracking odour plumes

Hugh Zhang, David Parkes
Chain of Thought Reasoning is a Policy Improvement Operator

Workshop location:

The workshop will be held **in-person**. Talks will be held in **Geological Lecture Hall**, (24 Oxford street, Cambridge MA 02138). Please find here further information on accessibility (<https://nw.fas.harvard.edu/accessibility>) and parking (<https://nw.fas.harvard.edu/parking>).



* To find the **Geological Lecture Hall** in the red brick Museum of Comparative Zoology building, paste this plus code into Google Maps: **“9VHM+7R Cambridge, Massachusetts.”** There are 2 entrances: from the south end of Tozzer Courtyard and from the south set of stairs up from the Oxford Street side (see photos).



[Tozzer courtyard](#) is right next to Geological Lecture Hall. Look for a large white tent as in the below picture (taken from the door of Geological Lecture Hall):



[Northwest Building](#), the location of the poster session, is close to Geological Lecture Hall.

It looks like this:

