

# Roy Fox

Assistant Professor of Computer Science  
University of California, Irvine

Intelligent Dynamics Lab  
4064 Donald Bren Hall  
Irvine, CA 92697-3435

✉ [royf@uci.edu](mailto:royf@uci.edu)

🌐 [royf.org](http://royf.org)

👤 [Roy Fox](#)

🐦 [@roydfox](#)

🎧 [royf](#)

## Appointment

2019–current **Assistant Professor**, *Department of Computer Science*, Donald Bren School of Information and Computer Science, University of California, Irvine

Affiliations: Intelligent Dynamics Lab ([indylab](#)) • Center for Machine Learning and Intelligent Systems (CML) • Department of Electrical Engineering and Computer Science (EECS) • HPI Research Center in Machine Learning and Data Science • Center for Theoretical Behavioral Sciences (CTBS) • Algorithms, Combinatorics, and Optimization Center (ACO)

## Research Positions

2016–2019 **Postdoctoral Scholar**, *Department of Electrical Engineering and Computer Sciences*, University of California, Berkeley

Advisors: Ion Stoica and Ken Goldberg

Affiliations: Real-time Intelligent Secure Explainable Systems (RISELab) • Automation Lab (AUTOLAB) • Berkeley Artificial Intelligent Research (BAIR)

2012–2013 **Exchange PhD Student**, *Center for Theoretical Neuroscience*, Columbia University  
Hosts: Larry Abbott and Liam Paninski

2011 **Research Intern**, Microsoft Research

## Education

2017 **PhD**, *School of Computer Science and Engineering*, Hebrew University

Advisor: Naftali Tishby

Thesis: Information-Theoretic Methods for Planning and Learning in Partially Observable Markov Decision Processes

2008 **MS**, *Computer Science Department*, Technion IIT

Advisor: Moshe Tennenholtz

Thesis: Reinforcement Learning in Partially Observable Decision Processes

2001 **BS** (cum laude), *School of Computer Science*, Hebrew University

2000 **BS**, *Medical Sciences*, Hebrew University Hadassah Medical School

## Grants and Awards

2024–current **NSF CISE: Large**, *Causal Foundations for Decision Making and Learning*  
Role: Senior Personnel; PI: Elias Bareinboim; award: 2321786; amount: \$5,000,000

2024 **ICS Research Award: Exploration**, *Safe Transfer from Simulation to Deployment*  
Role: Principal Investigator; amount: \$75,000

2022 **Chancellor's Award for Excellence in Undergraduate Research**, UCI

2020–current **HPI Research Center**, *Machine Learning and Data Science*

Role: Supervisor; Director: Erik Sudderth; three fellowships support my students

- 2020 **ICS Research Award: Endeavor**, *AI for High-Level Computational Science*  
Role: Key Personnel; PI: Eric Mjolsness; amount: \$299,849
- 2018 **Workshop sponsorship**, *Infer to Control workshop at NeurIPS*  
Sponsors: Intel and DeepMind; amount: \$6500
- 2017 **Workshop sponsorship**, *Hierarchical Reinforcement Learning workshop at NeurIPS*  
Sponsor: Intel; amount: \$5000
- 2012–2013 **Research scholarship**, *The Gatsby Charitable Foundation*, Columbia University
- 2011 **First place**, *HUJICode* algorithmic programming competition, Hebrew University
- 1998 **Rector award for excellence**, Hebrew University

## Publications

### Conferences

- RLC 2024 Reinforcement Learning from Delayed Observations via World Models [↗](#)  
Armin Karamzade, Kyungmin Kim, Montek Kalsi, and **Roy Fox**  
1st Reinforcement Learning Conference
- RLC 2024 Verification-Guided Shielding for Deep Reinforcement Learning [↗](#)  
Davide Corsi, Guy Amir, Andoni Rodríguez, Guy Katz, César Sánchez, and **Roy Fox**  
1st Reinforcement Learning Conference
- ICML 2024 Skill Set Optimization: Reinforcing Language Model Behavior via Transferable Skills [↗](#)  
Kolby Nottingham, Bodhisattwa Majumder, Bhavana Dalvi, Sameer Singh, Peter Clark, and **Roy Fox**  
41st International Conference on Machine Learning
- NAACL 2024 Selective Perception: Learning Concise State Descriptions for Language Model Actors [↗](#)  
Kolby Nottingham, Yasaman Razeghi, Kyungmin Kim, JB Lanier, Pierre Baldi, **Roy Fox**, and Sameer Singh  
2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics
- ICLR 2024 Toward Optimal Policy Population Growth in Two-Player Zero-Sum Games [↗](#)  
Stephen McAleer, JB Lanier, Kevin Wang, Pierre Baldi, Tuomas Sandholm, and **Roy Fox**  
12th International Conference on Learning Representations
- ICML 2023 Do Embodied Agents Dream of Pixelated Sheep?: Embodied Decision Making using Language Guided World Modelling [↗](#)  
Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi, Sameer Singh, and **Roy Fox**  
40th International Conference on Machine Learning
- ICML 2023 Learning to Design Analog Circuits to Meet Threshold Specifications [↗](#)  
Dmitrii Krylov, Pooya Khajeh, Junhan Ouyang, Thomas Reeves, Tongkai Liu, Hiba Ajmal, Hamidreza Aghasi, and **Roy Fox**  
40th International Conference on Machine Learning

- ICML 2022 Reducing Variance in Temporal-Difference Value Estimation via Ensemble of Deep Networks [↗](#)  
Litian Liang, Yaosheng Xu, Stephen McAleer, Dailin Hu, Alexander Ihler, Pieter Abbeel, and **Roy Fox**  
39th International Conference on Machine Learning
- RLDM 2022 Learning to Query Internet Text for Informing Reinforcement Learning Agents [↗](#)  
Kolby Nottingham, Alekhya Pyla, Sameer Singh, and **Roy Fox**  
5th Multi-disciplinary Conference on Reinforcement Learning and Decision Making
- AISTATS 2022 Independent Natural Policy Gradient Always Converges in Markov Potential Games [↗](#)  
**Roy Fox**, Stephen McAleer, Will Overman, and Ioannis Panageas  
25th International Conference on Artificial Intelligence and Statistics
- NeurIPS 2021 XDO: A Double Oracle Algorithm for Extensive-Form Games [↗](#)  
Stephen McAleer, JB Lanier, Kevin Wang, Pierre Baldi, and **Roy Fox**  
35th Conference on Neural Information Processing Systems
- NeurIPS 2020 Pipeline PSRO: A Scalable Approach for Finding Approximate Nash Equilibria in Large Games [↗](#)  
Stephen McAleer, JB Lanier, **Roy Fox**, and Pierre Baldi  
34th Conference on Neural Information Processing Systems
- SPLASH AutoPandas: Neural-Backed Generators for Program Synthesis [↗](#)
- OOPSLA 2019 Rohan Bavishi, Caroline Lemieux, **Roy Fox**, Koushik Sen, and Ion Stoica  
10th ACM SIGPLAN Conference on Systems, Programming, Languages, and Applications: Software for Humanity
- CASE 2019 Multi-Task Hierarchical Imitation Learning for Home Automation [↗](#)  
**Roy Fox\***, Ron Berenstein\*, Ion Stoica, and Ken Goldberg  
15th IEEE Conference on Automation Science and Engineering
- WAFR 2018 Generalizing Robot Imitation Learning with Invariant Hidden Semi-Markov Models [↗](#)  
Ajay Kumar Tanwani, Jonathan Lee, Brijen Thananjeyan, Michael Laskey, Sanjay Krishnan, **Roy Fox**, Ken Goldberg, and Sylvain Calinon  
13th International Workshop on the Algorithmic Foundations of Robotics
- CASE 2018 Constraint Estimation and Derivative-Free Recovery for Robot Learning from Demonstrations [↗](#)  
Jonathan Lee, Michael Laskey, **Roy Fox**, and Ken Goldberg  
14th IEEE Conference on Automation Science and Engineering
- ICML 2018 RLlib: Abstractions for Distributed Reinforcement Learning [↗](#)  
Eric Liang\*, Richard Liaw\*, Robert Nishihara, Philipp Moritz, **Roy Fox**, Ken Goldberg, Joseph Gonzalez, Michael Jordan, and Ion Stoica  
35th International Conference on Machine Learning
- ICRA 2018 Fast and Reliable Autonomous Surgical Debridement with Cable-Driven Robots Using a Two-Phase Calibration Procedure [↗](#)  
Daniel Seita, Sanjay Krishnan, **Roy Fox**, Stephen McKinley, John Canny, and Ken Goldberg  
35th IEEE International Conference on Robotics and Automation
- ICRA 2018 Robustly Adjusting Indoor Drip Irrigation Emitters with the Toyota HSR Robot [↗](#)  
Ron Berenstein, **Roy Fox**, Stephen McKinley, Stefano Carpin, and Ken Goldberg  
35th IEEE International Conference on Robotics and Automation

- ICLR 2018 Parametrized Hierarchical Procedures for Neural Programming [↗](#)  
**Roy Fox**, Richard Shin, Sanjay Krishnan, Ken Goldberg, Dawn Song, and Ion Stoica  
6th International Conference on Learning Representations
- CoRL 2017 DDCO: Discovery of Deep Continuous Options for Robot Learning from Demonstrations [↗](#)  
Sanjay Krishnan\*, **Roy Fox**\*, Ion Stoica, and Ken Goldberg  
1st Conference on Robot Learning
- CoRL 2017 DART: Noise Injection for Robust Imitation Learning [↗](#)  
Michael Laskey, Jonathan Lee, **Roy Fox**, Anca Dragan, and Ken Goldberg  
1st Conference on Robot Learning
- CASE 2017 An Algorithm and User Study for Teaching Bilateral Manipulation via Iterated Best Response Demonstrations [↗](#)  
Carolyn Chen, Sanjay Krishnan, Michael Laskey, **Roy Fox**, and Ken Goldberg  
13th IEEE Conference on Automation Science and Engineering
- CASE 2017 Statistical Data Cleaning for Deep Learning of Automation Tasks from Demonstrations [↗](#)  
Caleb Chuck, Michael Laskey, Sanjay Krishnan, Ruta Joshi, **Roy Fox**, and Ken Goldberg  
13th IEEE Conference on Automation Science and Engineering
- CDC 2016 Minimum-Information LQG Control — Part I: Memoryless Controllers [↗](#)  
**Roy Fox** and Naftali Tishby  
55th IEEE Conference on Decision and Control
- CDC 2016 Minimum-Information LQG Control — Part II: Retentive Controllers [↗](#)  
**Roy Fox** and Naftali Tishby  
55th IEEE Conference on Decision and Control
- EWRL 2016 Principled Option Learning in Markov Decision Processes [↗](#)  
**Roy Fox**\*, Michal Moshkovitz\*, and Naftali Tishby  
13th European Workshop on Reinforcement Learning
- UAI 2016 Taming the Noise in Reinforcement Learning via Soft Updates [↗](#)  
**Roy Fox**\*, Ari Pakman\*, and Naftali Tishby  
32nd Conference on Uncertainty in Artificial Intelligence
- NeurIPS 2013 A Multi-Agent Control Framework for Co-Adaptation in Brain-Computer Interfaces [↗](#)  
Josh Merel\*, **Roy Fox**\*, Tony Jebara, and Liam Paninski  
27th Annual Conference on Neural Information Processing Systems
- ICML 2012 Bounded Planning in Passive POMDPs [↗](#)  
**Roy Fox** and Naftali Tishby  
29th International Conference on Machine Learning
- AAAI 2007 A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs [↗](#)  
**Roy Fox** and Moshe Tennenholtz  
22nd Conference on Artificial Intelligence

## Symposia and Workshops

- TAFM Make the Pertinent Salient: Task-Relevant Reconstruction for Visual Control with  
@ RLC 2024 Distraction [↗](#)  
Kyungmin Kim, Charless Fowlkes, and **Roy Fox**  
Training Agents with Foundation Models workshop
- PRL Q\* Search: Heuristic Search with Deep Q-Networks [↗](#)  
@ ICAPS 2024 Forest Agostinelli, Shahaf Shperberg, Alexander Shmakov, Stephen McAleer, **Roy Fox**, and  
Pierre Baldi  
Bridging the Gap Between AI Planning and Reinforcement Learning workshop
- FMDM Selective Perception: Learning Concise State Descriptions for Language Model  
@ NeurIPS 2023 Actors [↗](#)  
Kolby Nottingham, Yasaman Razeghi, Kyungmin Kim, JB Lanier, Pierre Baldi, **Roy Fox**,  
and Sameer Singh  
Foundation Models for Decision Making workshop
- RRL Do Embodied Agents Dream of Pixelated Sheep?: Embodied Decision Making using  
@ ICLR 2023 Language Guided World Modelling  
Kolby Nottingham, Prithviraj Ammanabrolu, Alane Suhr, Yejin Choi, Hannaneh Hajishirzi,  
Sameer Singh, and **Roy Fox**  
Reincarnating Reinforcement Learning workshop
- DeepRL Feasible Adversarial Robust Reinforcement Learning for Underspecified Environ-  
@ NeurIPS 2022 ments [↗](#)  
JB Lanier, Stephen McAleer, Pierre Baldi, and **Roy Fox**  
Deep Reinforcement Learning workshop
- RLG Anytime PSRO for Two-Player Zero-Sum Games [↗](#)  
@ AAAI 2022 Stephen McAleer, Kevin Wang, JB Lanier, Marc Lanctot, Pierre Baldi, Tuomas Sandholm,  
and **Roy Fox**  
Reinforcement Learning in Games workshop
- DeepRL Temporal-Difference Value Estimation via Uncertainty-Guided Soft Updates [↗](#)  
@ NeurIPS 2021 Litian Liang, Yaosheng Xu, Stephen McAleer, Dailin Hu, Alexander Ihler, Pieter Abbeel,  
and **Roy Fox**  
Deep Reinforcement Learning workshop
- DeepRL Count-Based Temperature Scheduling for Maximum Entropy Reinforcement Learn-  
@ NeurIPS 2021 ing [↗](#)  
Dailin Hu, Pieter Abbeel, and **Roy Fox**  
Deep Reinforcement Learning workshop
- DeepRL Target Entropy Annealing for Discrete Soft Actor–Critic [↗](#)  
@ NeurIPS 2021 Yaosheng Xu, Dailin Hu, Litian Liang, Stephen McAleer, Pieter Abbeel, and **Roy Fox**  
Deep Reinforcement Learning workshop
- PRL Obtaining Approximately Admissible Heuristic Functions through Deep Reinforce-  
@ ICAPS 2021 ment Learning and A\* Search [↗](#)  
Forest Agostinelli, Stephen McAleer, Alexander Shmakov, **Roy Fox**, Marco Valtorta, Biplav  
Srivastava, and Pierre Baldi  
Bridging the Gap Between AI Planning and Reinforcement Learning workshop

- EmbodiedAI Modular Framework for Visuomotor Language Grounding [↗](#)  
 @ CVPR 2021 Kolby Nottingham, Litian Liang, Daeyun Shin, Charless Fowlkes, **Roy Fox**, and Sameer Singh  
 Embodied AI workshop
- RLG CFR-DO: A Double Oracle Algorithm for Extensive-Form Games [↗](#)  
 @ AAI 2021 Stephen McAleer, John Lanier, Pierre Baldi, and **Roy Fox**  
 Reinforcement Learning in Games workshop
- OPTRL Toward Provably Unbiased Temporal-Difference Value Estimation [↗](#)  
 @ NeurIPS 2019 **Roy Fox**  
 Optimization Foundations for Reinforcement Learning workshop
- AMTL Multi-Task Learning via Task Multi-Clustering [↗](#)  
 @ ICML 2019 Andy Yan, Xin Wang, Ion Stoica, Joseph Gonzalez, and **Roy Fox**  
 Adaptive and Multitask Learning workshop
- Infer2Control Hierarchical Imitation Learning via Variational Inference of Control Programs [↗](#)  
 @ NeurIPS 2018 **Roy Fox**, Richard Shin, William Paul, Yitian Zou, Dawn Song, Ken Goldberg, Pieter Abbeel, and Ion Stoica  
 Infer to Control: Probabilistic Reinforcement Learning and Structured Control workshop
- DLT An Empirical Exploration of Gradient Correlations in Deep Learning [↗](#)  
 @ NeurIPS 2018 Daniel Rothchild, **Roy Fox**, Noah Golmant, Joseph Gonzalez, Michael Mahoney, Kai Rothauge, Ion Stoica, and Zhewei Yao  
 Integration of Deep Learning Theories workshop
- ML for Sys Neural Inference of API Functions from Input–Output Examples [↗](#)  
 @ NeurIPS 2018 Rohan Bavishi, Caroline Lemieux, Neel Kant, **Roy Fox**, Koushik Sen, and Ion Stoica  
 Machine Learning for Systems workshop
- NAMPI Imitation Learning of Hierarchical Programs via Variational Inference [↗](#)  
 @ ICML 2018 **Roy Fox\***, Richard Shin\*, Pieter Abbeel, Ken Goldberg, Dawn Song, and Ion Stoica  
 Neural Abstract Machines and Program Induction workshop
- PGMRL Task-Relevant Embeddings for Robust Perception in Reinforcement Learning [↗](#)  
 @ ICML 2018 Eric Liang, **Roy Fox**, Joseph Gonzalez, and Ion Stoica  
 Prediction and Generative Modeling in Reinforcement Learning workshop
- Causal Imit. Robot Learning with Invariant Hidden Semi-Markov Models [↗](#)  
 @ RSS 2018 Ajay Kumar Tanwani, Jonathon Lee, Michael Laskey, Sanjay Krishnan, **Roy Fox**, and Ken Goldberg  
 Perspectives on Robot Learning: Imitation and Causality workshop
- DeepRL Ray RLLib: A Composable and Scalable Reinforcement Learning Library [↗](#)  
 @ NeurIPS 2017 Eric Liang\*, Richard Liaw\*, Robert Nishihara, Philipp Moritz, **Roy Fox**, Joseph Gonzalez, Ken Goldberg, and Ion Stoica  
 Deep Reinforcement Learning symposium

### Theses

- 2016 Information-Theoretic Methods for Planning and Learning in Partially Observable Markov Decision Processes [↗](#)  
**Roy Fox**  
 PhD Thesis, Hebrew University

- 2008 Reinforcement Learning in Partially Observable Decision Processes [↗](#)  
**Roy Fox**  
 MS Thesis, Technion IIT
- [Preprints](#)
- 2021 Improving Social Welfare while Preserving Autonomy via a Pareto Mediator [↗](#)  
 Stephen McAleer, JB Lanier, Michael Dennis, Pierre Baldi, and **Roy Fox**  
 arXiv:2106.03927
- 2021 A\* Search Without Expansions: Learning Heuristic Functions with Deep Q-Networks [↗](#)  
 Forest Agostinelli, Alexander Shmakov, Stephen McAleer, **Roy Fox**, and Pierre Baldi  
 arXiv:2102.04518
- 2019 Hierarchical Variational Imitation Learning of Control Programs [↗](#)  
**Roy Fox**, Richard Shin, William Paul, Yitian Zou, Dawn Song, Ken Goldberg, Pieter Abbeel, and Ion Stoica  
 arXiv:1912.12612
- 2017 Multi-Level Discovery of Deep Options [↗](#)  
**Roy Fox\***, Sanjay Krishnan\*, Ion Stoica, and Ken Goldberg  
 arXiv:1703.08294
- 2015 Optimal Selective Attention in Reactive Agents [↗](#)  
**Roy Fox** and Naftali Tishby  
 arXiv:1512.08575

## Talks and Presentations

### Invited Talks

- 2024 Learning to Control with Vision–Language Models  
 University of Pennsylvania, June 2
- 2024 How Large Models of Image and Text Help Agents to Learn to Behave  
 The Institute for Learning-enabled Optimization at Scale (TILOS), University of California, San Diego, March 20
- 2023 Population-Based Competitive Multi-Agent Reinforcement Learning  
 Information Theory and Applications workshop (ITA), February 16
- 2023 Population-Based Methods for Single- and Multi-Agent Reinforcement Learning  
 Information Sciences Institute, University of Southern California, January 27
- 2022 Better Together: Population-Based Methods for Single- and Multi-Agent Reinforcement Learning  
 Department of Electrical and Computer Engineering, Technion IIT, December 29
- 2020 Structured Control as Inference  
 Information Theory and Applications workshop (ITA), February 7
- 2019 Learning Hierarchical Control Programs
- University of Michigan, April 11
  - University of California, Irvine, March 19
  - Université de Montréal, February 22

- 2018 Panel on Inference for Control, moderator  
Infer to Control: Probabilistic Reinforcement Learning and Structured Control workshop (Infer2Control @ NeurIPS 2018), December 8
- 2018 Hierarchical Imitation Learning of Robot Skills and Control Programs
  - Google Research, November 15
  - Department of Electrical Engineering, Technion IIT, November 14
- 2017 Panel on Hierarchical Reinforcement Learning, moderator  
Hierarchical Reinforcement Learning workshop (HRL @ NeurIPS), December 9
- 2017 Deep Hierarchical Reinforcement Learning Made Fast and Easy with Ray  
Data Science Summit (DSS) Europe, May 29
- 2016 An Information-Theoretic Bag of Tools for Optimal Control and Reinforcement Learning  
Google DeepMind, June 1
- 2016 Minimum-Information Planning in Partially-Observable Decision Problems
  - Max Planck Institute for Mathematics in the Sciences, May 10
  - University of Stuttgart, May 9
- 2015 Soft-Optimal Decisions under Uncertainty  
Information and Decision Making workshop (IDeM), University of Hertfordshire, September 3
- 2014 Optimal Selective Attention and Action in Reactive Agents
  - University of California, Berkeley, October 8
  - California Institute of Technology, October 2
  - University of California, Los Angeles, October 1
  - University of California, San Diego, September 29
  - Columbia University, September 24
  - Microsoft Research, September 16
  - University of Pennsylvania, September 15
- 2012 Information Theoretic Methods for POMDP Planning  
DARPA MSEE program review meeting, University of California, Los Angeles, May 2

### Contributed Talks

- 2022 MeanQ: Reducing Variance in Temporal-Difference Value Estimation via Ensemble of Deep Networks  
*ICML* spotlight, July 21
- 2019 Multi-Task Hierarchical Imitation Learning for Home Automation  
*CASE*, August 25
- 2016 Minimum-Information LQG Control  
*CDC* double talk, December 14
- 2016 Taming the Noise in Reinforcement Learning via Soft Updates  
*UAI*, June 26
- 2012 Bounded Planning in Passive POMDPs  
*ICML* spotlight, June 27
- 2007 A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs  
*AAAI*, July 24



- 2007 A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs  
*BISFAI*, June 21
- Departmental Talks**
- 2024 Learning to Control with Vision–Language Models  
 HPI@UCI Workshop, University of California, Irvine, April 3
- 2022 Curiously Effective Ensemble and Double-Oracle Reinforcement Learning Methods  
 CML Seminar, University of California, Irvine, January 10
- 2020 When “Optimal Control” Isn’t  
 ○ CS Seminar, University of California, Irvine, November 20  
 ○ IMBS Seminar, University of California, Irvine, April 17
- 2019 Hierarchical Imitation Learning of Robot Skills and Control Programs  
 RISE Retreat, University of California, Berkeley, January 17
- 2018 Multi-Task Hierarchical Imitation Learning of Robot Skills  
 ○ BAIR/CPAR/BDD Seminar, University of California, Berkeley, October 12  
 ○ RISE Seminar, University of California, Berkeley, October 8
- 2018 SkillHub: Shared Hierarchical Imitation Learning of Robot Control  
 RISE Retreat, University of California, Berkeley, May 23
- 2018 Parametrized Hierarchical Procedures for Neural Programming  
 RISE Retreat, University of California, Berkeley, January 10
- 2017 Discovery of Hierarchical Structure for Robot Learning and Neural Programming  
 BAIR/CPAR/BDD Seminar, University of California, Berkeley, November 3
- 2017 Reinforcement Learning Concepts  
 RISE Camp, University of California, Berkeley, September 7
- 2017 Multi-Level Discovery of Deep Options  
 ○ RISE–BDD Joint Mini-Retreat, University of California, Berkeley, May 2  
 ○ BAIR/CPAR/BDD Seminar, University of California, Berkeley, March 17
- 2017 Learning Control Hierarchies  
 RISE Retreat, University of California, Berkeley, January 11
- 2015 Soft-Optimal Decisions under Uncertainty  
 Machine Learning Club, Hebrew University, October 22
- 2013 Separating Inference and Control  
 The Center for Theoretical Neuroscience, Columbia University, May 13
- 2013 What to Notice and What to Remember: Finding Good Inference Policies  
 The Interdisciplinary Center for Neural Computation Retreat, Hebrew University, January 29
- 2012 Information Theory in Reinforcement Learning  
 ○ The Center for Theoretical Neuroscience, Columbia University, August 20  
 ○ Neural Data Analysis Seminar, Columbia University, May 22
- 2012 Bounded Planning in Passive POMDPs  
 Machine Learning Club, Hebrew University, March 15
- 2007 Reinforcement Learning in Partially Observable Decision Processes  
 Graduate Seminar, Technion IIT, November 4

## Advising and Mentoring

### Postdocs

2024–current Davide Corsi

### PhD Students

2023–current Armin Karamzade

2022–current Dmitry Krylov

2021–current Kyungmin Kim

2021–current Alexander Shmakov (advised by Pierre Baldi)

2020–current JB Lanier (jointly with Pierre Baldi)

2020–current Kolby Nottingham (jointly with Sameer Singh)

### Master's Students

2024–current Shubham Sharma

2024–current Zhiyan Tan

### Undergraduate Students

2023–current Montek Kalsi • Yanran Wang

### PhD Alumni

2019–2021 Stephen McAleer (advised by Pierre Baldi) → postdoc @ CMU → OpenAI

### Master's Alumni

2022–2023 Alekhya Pyla → Pinterest

2021–2022 Ankita Sinha → Intuit

2020–2022 Dailin Hu → Blizzard

2020–2022 Andrew Jiang → Cisco

2020–2022 Alex Konrad (thesis) → Spatial Genomics

### Past Visitors

2023–2024 Kat He → MS student @ Columbia

2022 Bryan Yu → ShieldAI

2021–2022 Kevin Wang → PhD student @ Brown

### Undergraduate Alumni

2023–2024 Tony Liu → MS student @ UCSD

2022–2023 Sushrut Borkar

2022–2023 Isabel Rosen → Microsoft

2022–2023 Zihao Zhu → MS student @ UCLA

2021–2023 Yifei Liu → MS student @ CMU

2021–2022 Xinyi Huang → Circle

2021–2022 Junhan Ouyang → DataVisor

2021–2022 Thomas Reeves → MS student @ USC → PhD student @ USC

2021–2022 Jiawen Zhang → MS student @ UPenn

2020–2022 Litian Liang → MS student @ UCSD

2020–2022 Yaosheng Xu → MS student @ Harvard

2020–2021 Yongxuan Fu → MS student @ Brown

2020–2021 Yao Luo → MS student @ Columbia  
2020–2021 Chinmay Tyagi → MS student @ GeorgiaTech  
2020–2021 Yinfei Wang → Huawei  
2020–2021 Ted Zadori → MS student @ UCLA → PhD student @ Princeton  
2019–2021 Carleton Zhao → Amazon  
2019–2020 Dailin Hu → MS student @ UCI  
2019–2020 Mingen Li → MS student @ UCSD  
2019–2020 Bryon Tjanaka → PhD student @ USC  
2019–2020 Catherine Yuan → General Atomics Aeronautical  
2018–2019 William Paul → Johns Hopkins University Applied Physics Lab  
2018–2019 Matthew Trepte → Nvidia  
2018–2019 Andy Yan → MS student @ Berkeley  
2018–2019 Yitian Zou → Uber  
2017–2018 David Tseng → Microsoft  
2014–2015 Itay Fried → Applied Materials  
2014–2015 Erez Peterfreund → MS student @ Hebrew University → ... → postdoc @ Yale

---

## Teaching

### Instruction

- CS 200S **Seminar in Computer Science Research**, *Computer Science*, UC Irvine [↗](#)  
Departmental seminar organizer; Fall 2021 – Spring 2024
- CS 277 **Control and Reinforcement Learning**, *Computer Science*, UC Irvine  
Winter 2020 [↗](#) • Winter 2021 [↗](#) • Winter 2022 [↗](#) • Winter 2024 [↗](#)
- CS 273A **Machine Learning**, *Computer Science*, UC Irvine  
Winter 2021 [↗](#) • Fall 2021 [↗](#)
- CS 175 **Project in Artificial Intelligence**, *Computer Science*, UC Irvine  
Winter 2020 [↗](#) • Fall 2021 [↗](#)
- 76929 **Reinforcement Learning**, *Center for Brain Sciences*, Hebrew University  
Winter 2014 [↗](#) • Winter 2015

### Guest Lectures

- Fall 2020 **AI/ML for Computational Science (CS 295)**, *Computer Science*, UC Irvine  
Fall 2019 **Honors Seminar (CS H197)**, *Information and Computer Science*, UC Irvine

---

## Professional Service

### Workshop Organizer

- 2024 Heuristics and Search for Domain-independent Planning (HSDIP @ ICAPS) [↗](#)  
Co-organizers: Imène Ait Abderrahim, Clemens Büchner, Daniel Gnad, Sofia Lemons
- 2018 Infer to Control: Probabilistic Reinforcement Learning and Structured Control workshop (Infer2Control @ NeurIPS) [↗](#)  
Co-organizers: Leslie Kaelbling, Martin Riedmiller, Marc Toussaint, Igor Mordatch, Tuomas Haarnoja; estimated attendance: 650+

2017 Hierarchical Reinforcement Learning workshop (HRL @ NeurIPS) [↗](#)  
Co-organizers: Andrew Barto, Doina Precup, Shie Mannor, Tom Schaul, Carlos Florensa;  
estimated attendance: 950+

### Outreach

2021 UCI End Racism Initiative  
Work group on *Representation: Recruitment and Outreach to Black Students*

### Senior Program Committee Member / Meta-Reviewer

2024 IJCAI conference • NeurIPS conference  
2023 IJCAI conference • NeurIPS conference  
2022 ICML conference • NeurIPS conference  
2021 IJCAI conference  
2020 IJCAI conference  
2019 UAI conference

### Session Chair / Moderator

2023 ITA workshop, Game Theory session chair  
2022 ICML conference, Reinforcement Learning session chair  
2022 NeurIPS conference, session chair  
2021 IJCAI conference, Deep Reinforcement Learning session chair  
2019 CASE conference, Service Robots session co-chair  
2018 Infer2Control @ NeurIPS workshop, panel moderator  
2017 HRL @ NeurIPS workshop, panel moderator  
2016 CDC conference, Information Theory and Control session co-chair

### Proposal Reviewer

2021 NSF CISE • NSF FRR  
2020 CEREC • ISF • NSF CISE  
2019 AI XPRIZE competition technical judge

### Program Committee Member / Reviewer

2024 **Workshops:** HSDIP @ ICAPS  
2022 **Journals:** IEEE-TAC • **Conferences:** DeepMath • RLDM  
2021 **Journals:** IEEE-TAC • JMLR • **Conferences:** ICRA • NeurIPS  
2020 **Journals:** Artificial Intelligence • JMLR • **Conferences:** AAAI • CDC • DeepMath  
• IROS • RSS  
2019 **Journals:** IJRR • JAAMAS • **Conferences:** CASE • ICML • ICRA • IJCAI •  
NeurIPS • RLDM • **Workshops:** HILL @ ICML • OPTRL @ NeurIPS • SPiRL @  
ICLR  
2018 **Journals:** Adaptive Behavior • IJRR • JAAMAS • **Conferences:** ICLR • ICML  
• NeurIPS • RSS • **Workshops:** Infer2Control @ NeurIPS  
2017 **Conferences:** ICLR • NeurIPS • RSS • **Workshops:** HRL @ NeurIPS  
2016 **Conferences:** NeurIPS