

Roy Fox

Assistant Professor, UC Irvine

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

✉ roy.d.fox@gmail.com
🌐 royf.org
👤 [author:Roy Fox](#)
🐦 [@roydfox](#)
🔗 [royf](#)



Appointment

- 2019 **Assistant Professor**, *Department of Computer Science*, Donald Bren School of Information & Computer Science, University of California, Irvine.
Affiliations: Intelligent Dynamics Lab (indylab) ▪ Center for Machine Learning and Intelligent Systems (CML) [↗](#) ▪ Algorithms, Combinatorics and Optimization Center (ACO) [↗](#)

Research Interests

Reinforcement learning ▪ Dynamical systems ▪ Information theory ▪ Robotics

Education

- 2017 **PhD**, *School of Computer Science and Engineering*, Hebrew University.
Advisor: Prof. Naftali Tishby
Thesis: Information-Theoretic Methods for Planning and Learning in Partially Observable Markov Decision Processes [↗](#)
- 2008 **MS**, *Computer Science Department*, Technion IIT.
Advisor: Prof. 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.

Teaching

- Winter 2020 **Optimal Control and Reinforcement Learning** (CS 295) [↗](#)
- Winter 2020 **Project in Artificial Intelligence** (CS 175) [↗](#)
- 2014–2015 **Reinforcement Learning**, *Center for Brain Sciences*, Hebrew University [↗](#)

Research Experience

- 2016–2019 **Postdoctoral Scholar**, *Department of Electrical Engineering and Computer Sciences*, University of California, Berkeley.
Advisors: Profs. 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: Profs. Larry Abbott and Liam Paninski
- 2011 **Research Intern**, Microsoft Research.
- 2009–2011 **Head of Research**, Sentrigo.

- 2008 **Researcher**, Safend.
 2001–2007 **Researcher and Team Leader**, IDF.

Entrepreneurial Experience

- 2013–2015 **Founder & CTO**, AudLog.
 2008–2009 **Founder & CTO**, Serendipity.
 2007–2008 **Founder & CTO**, QuiXilver.

Publications

Conferences

- SPLASH 2019 AutoPandas: Neural-Backed Generators for Program Synthesis [↗](#)
 OOPSLA 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 Demon-
 strations [↗](#)
 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 Demon-
 strations [↗](#)
 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**
- 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 & 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 & 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

Preprints

- 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

Advising and Mentoring

Undergraduate Students

- 2019–current Catherine Yuan
2019–current Carleton Zhao
2019–current Bryon Tjanaka
2019–current Mingen Li
2019–current Dailin Hu
2018–2019 Andy Yan (MS student @ UC Berkeley)
2018–2019 Matthew Trepte (BS @ UC Berkeley)
2018–2019 William Paul (BS @ UC Berkeley → Johns Hopkins University Applied Physics Lab)
2018–2019 Yitian Zou (BS @ UC Berkeley → Uber)
2017–2018 David Tseng (MS @ UC Berkeley → Microsoft)

2014–2015 Erez Peterfreund (PhD @ Hebrew University → Postdoc @ Yale University)

2014–2015 Itay Fried (BS @ Hebrew University → Applied Materials)

Service

Organizer

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+

Senior Program Committee Member / Meta-Reviewer

IJCAI conference (2020) ▪ UAI conference (2019)

Program Committee Member / Reviewer

AAAI conference (2020) ▪ NeurIPS workshops (OPTRL 2019, Infer2Control 2018, HRL 2017) ▪ NeurIPS conference (2016–2019) ▪ JAAMAS journal (2018–2019) ▪ ICML workshop (HILL 2019) ▪ CASE conference (2019) ▪ IJRR journal (2018–2019) ▪ IJCAI conference (2019) ▪ RLDM conference (2019) ▪ ICLR workshop (SPiRL 2019) ▪ ICML conference (2018–2019) ▪ ICRA conference (2019) ▪ RSS conference (2017–2018) ▪ ICLR conference (2017–2018) ▪ Adaptive Behavior journal (2018)

Volunteer

2019 Technical judge, IBM Watson AI XPRIZE competition

2016 International Workshop on the Algorithmic Foundations of Robotics (WAFR)

2007 Conference of the Association for the Advancement of Artificial Intelligence (AAAI)

Talks and Presentations

Contributed Presentations

2019 Multi-Task Hierarchical Imitation Learning for Home Automation

CASE, talk, Vancouver, Canada, August 25

2018 Panel on Inference for Control

Infer2Control Workshop @ NeurIPS, moderator, Montréal, Canada, December 8

2018 Hierarchical Imitation Learning via Variational Inference of Control Programs

Infer2Control Workshop @ NeurIPS, poster, Montréal, Canada, December 8

2018 Imitation Learning of Hierarchical Programs via Variational Inference

NAMPI Workshop @ ICML, poster, Stockholm, Sweden, July 15

2018 Task-Relevant Embeddings for Robust Perception in Reinforcement Learning

PGMRL Workshop @ ICML, poster, Stockholm, Sweden, July 15

2018 Parametrized Hierarchical Procedures for Neural Programming

ICLR, poster, Vancouver, Canada, April 30

2017 Panel on Hierarchical Reinforcement Learning

HRL Workshop @ NeurIPS, panel moderator, Long Beach, California, December 9

2017 Ray RLLib: A Composable and Scalable Reinforcement Learning Library

DeepRL Symposium @ NeurIPS, poster, Long Beach, California, December 7

- 2016 Minimum-Information LQG Control
CDC, double talk, Las Vegas, Nevada, December 14
- 2016 Taming the Noise in Reinforcement Learning via Soft Updates
UAI, talk and poster, Jersey City, New Jersey, June 26
- 2013 A Multi-Agent Control Framework for Co-Adaptation in Brain-Computer Interfaces
NeurIPS, poster, Lake Tahoe, Nevada, December 8
- 2012 Bounded Planning in Passive POMDPs
ICML, spotlight and poster, Edinburgh, UK, June 27
- 2007 A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs
AAAI, talk, Vancouver, Canada, July 24
- 2007 A Reinforcement Learning Algorithm with Polynomial Interaction Complexity for Only-Costly-Observable MDPs
BISFAI, talk, Ramat Gan, Israel, June 21
- Invited Talks**
- 2020 Structured Control as Inference
 Information Theory and Applications workshop, San Diego, February 7
- 2019 Learning Hierarchical Control Programs
 - University of Michigan, Ann Arbor, April 11
 - University of California, Irvine, March 19
 - Université de Montréal, Canada, February 22
- 2018 Hierarchical Imitation Learning of Robot Skills and Control Programs
 - Google Research, Tel Aviv, Israel, November 15
 - Department of Electrical Engineering, Technion IIT, Haifa, Israel, November 14
- 2017 Deep Hierarchical Reinforcement Learning Made Fast and Easy with Ray
 Data Science Summit Europe, Jerusalem, Israel, May 29
- 2016 An Information-Theoretic Bag of Tools for Optimal Control and Reinforcement Learning
 Google DeepMind, London, UK, June 1
- 2016 Minimum-Information Planning in Partially-Observable Decision Problems
 - Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany, May 10
 - Marc Toussaint's group, University of Stuttgart, Germany, May 9
- 2015 Soft-Optimal Decisions under Uncertainty
 Information and Decision Making workshop, University of Hertfordshire, UK, September 3
- 2014 Optimal Selective Attention and Action in Reactive Agents
 - Pieter Abbeel's group, University of California, Berkeley, October 8
 - Richard Murray's group, California Institute of Technology, Los Angeles, October 2
 - Stefano Soatto's group, University of California, Los Angeles, October 1
 - Artificial Intelligence group, University of California, San Diego, September 29
 - Liam Paninski's group, Columbia University, New York, September 24
 - Microsoft Research, New York, September 16
 - GRASP lab, University of Pennsylvania, Philadelphia, September 15
- 2012 Information Theoretic Methods for POMDP Planning
 DARPA MSEE program review meeting, University of California, Los Angeles, May 2

Departmental Talks

- 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
 - Liam Paninski's group, 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