

# Roy Fox

Postdoctoral Scholar, UC Berkeley

RISELab  
UC Berkeley  
465 Soda Hall  
Berkeley, CA 94720-1776


✉ [roy.d.fox@gmail.com](mailto:roy.d.fox@gmail.com)

🌐 [roydfox.com](http://roydfox.com)

🎓 Google Scholar

🐦 [@roydfox](https://twitter.com/roydfox)

in [royfox](https://www.linkedin.com/in/royfox)



## Appointment

2016–current **Postdoctoral Scholar**, *Berkeley AI Research (BAIR), RISELab, and AUTOLAB, Department of Electrical Engineering and Computer Sciences*, University of California, Berkeley.  
Advisors: Profs. Ion Stoica and Ken Goldberg

## Research Interests

Reinforcement learning • Dynamical systems • Information theory • Robotics

## Education

- 2016 **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 **MSc**, *Computer Science Department*, Technion IIT.  
Advisor: Prof. Moshe Tennenholtz  
Thesis: Reinforcement Learning in Partially Observable Decision Processes [🔗](#)
- 2001 **BSc** (cum laude), *School of Computer Science*, Hebrew University.
- 2000 **BSc**, *Medical Sciences*, Hebrew University Hadassah Medical School.

## Research Experience

- 2012–2013 **Exchange PhD Student**, *Theoretical Neuroscience*, Columbia University.  
Hosts: Profs. Larry Abbott and Liam Paninski
- 2011 **Research Intern**, Microsoft Research, Israel.  
Investigated and implemented data mining for auction pricing
- 2009–2011 **Head of Research**, Sentrigo.  
Led a team of security researchers to discover database vulnerabilities and detect attacks
- 2008 **Researcher**, Safend.  
Advised on using machine learning for securing end-point devices
- 2001–2007 **Researcher and Team Leader**, IDF.  
Performed algorithmic research and development; led a team of Math and CS teachers

## Teaching Experience

- 2014–2015 **Reinforcement learning**, *teacher*, Center for Brain Sciences, Hebrew University.  
Created and taught a new graduate class
- 2011 **Computability and Complexity**, *teaching assistant*, School of Computer Science and Engineering, Hebrew University.
- 2007 **Digital Systems**, *teaching assistant*, Department of Electrical Engineering, Technion IIT.

## Entrepreneurial Experience

- 2013–2015 **Founder & CTO**, *AudLog*.  
Developed a mobile app for recording meetings and summarizing topics and action items
- 2008–2009 **Founder & CTO**, *Serendipity*.  
Developed a web widget for recommended content discovery by collaborative filtering
- 2007–2008 **Founder & CTO**, *QuiXilver*.  
Developed middleware for seamlessly utilizing configurable hardware to accelerate software

## Honors and Awards

- 2012–2013 **Research scholarship**, *The Gatsby Charitable Foundation*, Columbia University.
- 2011 **First place**, *HUJICode*, Hebrew University's algorithmic programming competition.
- 2007 **Travel grant**, *AAAI*, Vancouver, Canada.
- 1998 **Rector award for excellence**, Hebrew University.
- 1997 **First place**, *The International Mathematics Tournament of the Towns*.

## Publications

### Conferences

- 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 (WAFR)
- 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 (CASE)
- 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 (ICML)
- 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)
- 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 (ICRA)
- 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 (ICLR)
- 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)
- 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 (CoRL)

- 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)
- 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 (CASE)
- CDC 2016 Minimum-Information LQG Control — Part I: Memoryless Controllers [↗](#)  
**Roy Fox** and Naftali Tishby  
 55th IEEE Conference on Decision and Control (CDC)
- CDC 2016 Minimum-Information LQG Control — Part II: Retentive Controllers [↗](#)  
**Roy Fox** and Naftali Tishby  
 55th IEEE Conference on Decision and Control (CDC)
- EWRL 2016 Principled Option Learning in Markov Decision Processes [↗](#)  
**Roy Fox\***, Michal Moshkovitz\*, and Naftali Tishby  
 13th European Workshop on Reinforcement Learning (EWRL)
- 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 (UAI)
- NIPS 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 (NIPS)
- ICML 2012 Bounded Planning in Passive POMDPs [↗](#)  
**Roy Fox** and Naftali Tishby  
 29th International Conference on Machine Learning (ICML)
- 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 (AAAI)
- Symposia and Workshops**
- Infer2Control Hierarchical Imitation Learning via Variational Inference of Control Programs [↗](#)  
 @NeurIPS2018 **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 (Infer2Control @ NeurIPS)
- DLT An Empirical Exploration of Gradient Correlations in Deep Learning [↗](#)  
 @NeurIPS2018 Daniel Rothchild, **Roy Fox**, Noah Golmant, Joseph Gonzalez, Michael Mahoney, Kai Rothauge, Ion Stoica, and Zhewei Yao  
 Integration of Deep Learning Theories workshop (DLT @ NeurIPS)
- ML for Sys Neural Inference of API Functions from Input–Output Examples [↗](#)  
 @NeurIPS2018 Rohan Bavishi, Caroline Lemieux, Neel Kant, **Roy Fox**, Koushik Sen, and Ion Stoica  
 Machine Learning for Systems workshop (ML for Sys @ NeurIPS)
- NAMPI Imitation Learning of Hierarchical Programs via Variational Inference [↗](#)  
 @ICML2018 **Roy Fox\***, Richard Shin\*, Pieter Abbeel, Ken Goldberg, Dawn Song, and Ion Stoica  
 Neural Abstract Machines & Program Induction workshop (NAMPI @ ICML)

- PGMRL Task-Relevant Embeddings for Robust Perception in Reinforcement Learning [↗](#)
- @ICML2018 Eric Liang, **Roy Fox**, Joseph Gonzalez, and Ion Stoica  
Prediction and Generative Modeling in Reinforcement Learning workshop (PGMRL @ ICML)
- Causal Imit. Robot Learning with Invariant Hidden Semi-Markov Models [↗](#)
- @RSS2018 Ajay Kumar Tanwani, Jonathon Lee, Michael Laskey, Sanjay Krishnan, **Roy Fox**, and Ken Goldberg  
Perspectives on Robot Learning: Imitation and Causality workshop (Causal Imit. @ RSS)
- DeepRL Ray RLLib: A Composable and Scalable Reinforcement Learning Library [↗](#)
- @NIPS2017 Eric Liang\*, Richard Liaw\*, Robert Nishihara, Philipp Moritz, **Roy Fox**, Joseph Gonzalez, Ken Goldberg, and Ion Stoica  
Deep Reinforcement Learning symposium (DeepRL @ NIPS)

### Preprints

- 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:1609.05524

## Talks and Presentations

### Contributed Presentations

- 2018 Panel on Inference for Control  
*Infer2Control Workshop @ NeurIPS*, panel 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 @ NIPS*, panel moderator, Long Beach, California, December 9
- 2017 Ray RLLib: A Composable and Scalable Reinforcement Learning Library  
*DeepRL Symposium @ NIPS*, 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  
*NIPS*, 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

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
- Vision lab, 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, 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

## Service

### Workshop Organizer

- 2018 Infer to Control: Probabilistic Reinforcement Learning and Structured Control Workshop (Infer2Control @ NeurIPS) [↗](#)  
Sponsors: Intel AI, Google Research, and DeepMind; estimated attendance: 650+
- 2017 Hierarchical Reinforcement Learning Workshop (HRL @ NIPS) [↗](#)  
Sponsors: Intel AI and DeepMind; estimated attendance: 950+

### Program Committee Member / Reviewer

- ICRA (2019)
- Infer2Control workshop @ NeurIPS (2018)
- NeurIPS (formerly NIPS) (2016–2018)
- JAAMAS journal (2018)
- ICML (2018)
- RSS (2017–2018)
- ICLR (2017–2018)
- Adaptive Behavior journal (2018)
- HRL Workshop @ NIPS (2017)

### Volunteer

- 2016 International Workshop on the Algorithmic Foundations of Robotics (WAFR)
- 2007 Conference of the Association for the Advancement of Artificial Intelligence (AAAI)

## Advising and Mentoring

### Master's Student

- David Tseng (2017–current)

### Undergraduate Students

- Andy Yan (2018–current)
- Matthew Trepte (2018–current)
- Yitian Zou (2018–current)
- William Paul (2018–current)
- Erez Peterfreund (2014–2015)
- Itay Fried (2014–2015)