

Dylan Foster

MIT Institute for Data, Systems, and Society (IDSS)
50 Ames Street, Building E17-481C
Cambridge, MA, 02142

(480) 316-9934
dylanf@mit.edu
<http://dylanfoster.net>

Research Interests

Machine learning theory: Online machine learning and sequential decision making (reinforcement learning, contextual bandits, online learning), statistical learning, optimization (non-convex, stochastic), deep learning, high-dimensional statistics. Concentration inequalities, martingale inequalities, empirical process theory.

Current Position

Massachusetts Institute of Technology Postdoctoral Fellow, MIT Institute for Foundations of Data Science (MIFODS)	Cambridge, MA <i>Jan. 2019 - Present</i>
---	---

Education

Cornell University Ph.D. in Computer Science <i>Advisor: Professor Karthik Sridharan</i> <i>Committee: Karthik Sridharan (chair), Robert Kleinberg, Eva Tardos, Kilian Weinberger</i> <i>Thesis: Adaptive Learning: Algorithms and Complexity</i>	Ithaca, NY <i>Aug. 2014 - Jan. 2019</i>
University of Southern California Joint B.S./M.S. in Electrical Engineering Minor in Mathematics <i>Magna Cum Laude</i>	Los Angeles, CA <i>Aug. 2010 - May 2014</i>

Publications

- Learning the Linear Quadratic Regulator from Nonlinear Observations***
Zakari Mhammedi, Dylan J. Foster, Max Simchowitz, Dipendra Misra, Wen Sun, Akshay Krishnamurthy, Alexander Rakhlin, and John Langford
Neural Information Processing Systems (NeurIPS), 2020.
Short version at Theoretical Foundations of Reinforcement Learning workshop at ICML 2020.
- Adapting to Misspecification in Contextual Bandits**
Dylan J. Foster, Claudio Gentile, Mehryar Mohri, and Julian Zimmert
Neural Information Processing Systems (NeurIPS), 2020.
- Decoupled Policy Gradient Methods for Competitive Reinforcement Learning**
Constantinos Daskalakis, Dylan J. Foster, and Noah Golowich
Neural Information Processing Systems (NeurIPS), 2020.
- Beyond UCB: Optimal and Efficient Contextual Bandits with Regression Oracles**
Dylan J. Foster and Alexander Rakhlin
International Conference on Machine Learning (ICML), 2020
- Logarithmic Regret for Adversarial Online Control***
Dylan J. Foster and Max Simchowitz
International Conference on Machine Learning (ICML), 2020

6. **Naive Exploration is Optimal for Online LQR***
Max Simchowitz and Dylan J. Foster
International Conference on Machine Learning (ICML), 2020
7. **Tight Bounds on Minimax Regret under Logarithmic Loss via Self-Concordance**
Blair Bilodeau, Dylan J. Foster, and Daniel Roy
International Conference on Machine Learning (ICML), 2020
8. **Second-Order Information in Non-Convex Stochastic Optimization: Power and Limitations**
Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Ayush Sekhari, and Karthik Sridharan
Conference on Learning Theory (COLT), 2020
9. **Open Problem: Model Selection for Contextual Bandits**
Dylan J. Foster, Akshay Krishnamurthy, and Haipeng Luo
Open Problem, Conference on Learning Theory (COLT), 2020
10. **Learning Nonlinear Dynamical Systems from a Single Trajectory**
Dylan J. Foster, Tuhin Sarkar, and Alexander Rakhlin
Learning for Dynamics and Control (LADC), 2020
Full oral presentation (top 14/131 submissions).
11. **Model Selection for Contextual Bandits**
Dylan J. Foster, Akshay Krishnamurthy, and Haipeng Luo
Neural Information Processing Systems (NeurIPS), 2019
Spotlight presentation (top 2.43% of submissions).
12. **Hypothesis Set Stability and Generalization**
Dylan J. Foster, Spencer Greenberg, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan
Neural Information Processing Systems (NeurIPS), 2019
13. **Distributed Learning with Sublinear Communication**
Jayadev Acharya, Christopher De Sa, Dylan J. Foster, and Karthik Sridharan
International Conference on Machine Learning (ICML), 2019
Long talk (top 4.5% of submissions).
14. **Sum-of-Squares Meets Square Loss: Fast Rates for Agnostic Tensor Completion**
Dylan J. Foster and Andrej Risteski
Conference on Learning Theory (COLT), 2019
15. **Statistical Learning with a Nuisance Component**
Dylan J. Foster and Vasilis Syrgkanis
Conference on Learning Theory (COLT), 2019
Best paper award (1 of 393 submissions).
16. **The Complexity of Making the Gradient Small in Stochastic Convex Optimization**
Dylan J. Foster, Ayush Sekhari, Ohad Shamir, Nathan Srebro, Karthik Sridharan, and Blake Woodworth
Conference on Learning Theory (COLT), 2019
Best student paper award (2 of 393 submissions).
17. **Contextual Bandits with Surrogate Losses: Margin Bounds and Efficient Algorithms**
Dylan J. Foster and Akshay Krishnamurthy
Neural Information Processing Systems (NeurIPS), 2018
18. **Uniform Convergence of Gradients for Non-Convex Learning and Optimization**
Dylan J. Foster, Ayush Sekhari, and Karthik Sridharan
Neural Information Processing Systems (NeurIPS), 2018
Short version at Nonconvex Optimization workshop at ICML 2018.
19. **Logistic Regression: The Importance of Being Improper**
Dylan J. Foster, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan
Conference on Learning Theory (COLT), 2018
Best student paper award (2 of 335 submissions).
20. **Online Learning: Sufficient Statistics and the Burkholder Method**
Dylan J. Foster, Alexander Rakhlin, and Karthik Sridharan
Conference on Learning Theory (COLT), 2018
21. **Practical Contextual Bandits with Regression Oracles***
Dylan J. Foster, Alekh Agarwal, Miroslav Dudik, Haipeng Luo, and Robert Schapire
International Conference on Machine Learning (ICML), 2018
Long talk (top 6% of submissions).

22. **Inference in Sparse Graphs with Pairwise Measurements and Side Information**
Dylan J. Foster, Daniel Reichman, and Karthik Sridharan
International Conference on Artificial Intelligence and Statistics (AISTATS), 2018.
23. **Spectrally-Normalized Margin Bounds for Neural Networks**
Peter Bartlett, Dylan J. Foster, and Matus Telgarsky
Neural Information Processing Systems (NIPS), 2017
Spotlight presentation (top 3.45% of submissions).
24. **Parameter-Free Online Learning via Model Selection**
Dylan J. Foster, Satyen Kale, Mehryar Mohri, and Karthik Sridharan
Neural Information Processing Systems (NIPS), 2017
Spotlight presentation (top 3.45% of submissions).
25. **ZigZag: A New Approach to Adaptive Online Learning**
Dylan J. Foster, Alexander Rakhlin, and Karthik Sridharan
Conference on Learning Theory (COLT), 2017
26. **Learning in Games: Robustness of Fast Convergence**
Dylan J. Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos
Neural Information Processing Systems (NIPS), 2016
Short version at Ad Auctions Workshop at EC 2016.
27. **Adaptive Online Learning**
Dylan J. Foster, Alexander Rakhlin, and Karthik Sridharan
Neural Information Processing Systems (NIPS), 2015
Spotlight presentation (top 4.46% of submissions).

* denotes contribution order

In Submission

1. **Orthogonal Statistical Learning**
Dylan J. Foster and Vasilis Syrgkanis
Under review at *Annals of Statistics*, 2020.
Preprint available at <https://arxiv.org/abs/1901.09036>
2. **Lower Bounds for Non-Convex Stochastic Optimization**
Yossi Arjevani, Yair Carmon, John C. Duchi, Dylan J. Foster, Nathan Srebro, and Blake Woodworth
Under review at *Mathematical Programming*, 2019.
Preprint available at <https://arxiv.org/abs/1912.02365>

In Preparation

1. **Instance-Dependent Complexity of Contextual Bandits and Reinforcement Learning: A Disagreement-Based Perspective**
Dylan J. Foster, Alexander Rakhlin, David Simchi-Levi, and Yunzong Xu
2020. To be submitted to *Mathematics of Operations Research*.
Preprint available upon request.

Honors and Awards

- Best Spotlight Talk Award, 14th Annual New York Academy of Science Machine Learning Symposium, 2020.
- Nominated for ACM Doctoral Dissertation Award by Cornell University (one per institution), 2019
- **Best Paper Award, Conference on Learning Theory, 2019**
- **Best Student Paper Award, Conference on Learning Theory, 2019**
- Best Spotlight Talk Award, 13th Annual New York Academy of Science Machine Learning Symposium, 2019.
- **Best Student Paper Award, Conference on Learning Theory, 2018**
- **Facebook Fellowship, 2018**
- **NDSEG Fellowship, 2016**
- Outstanding Teaching Award, Computer Science, Cornell University, 2015
- Stone Endowed Scholarship, USC, 2013 - 2014

- Tau Beta Pi, 2012
- Eta Kappa Nu, USC, 2012
- Provost's Undergraduate Research Fellowship, USC, 2012 - 2014
- Jerome Linn Endowed Scholarship, USC, 2011 - 2013
- Academic Achievement Award, USC, 2011 - 2013
- Dean's List of Distinguished Students, USC, 2010 - 2014

Research Positions

- **Visiting Postdoc, Simons Institute for the Theory of Computing, UC Berkeley**, August - Dec. 2020.
Theory of Reinforcement Learning Program
- **Long-term Visitor, Simons Institute for the Theory of Computing, UC Berkeley**, July - August 2019.
Foundations of Deep Learning Program
- **Intern, Microsoft Research New England**, June 2018 - Aug. 2018.
Host: Vasilis Syrgkanis
Developed theory for statistical learning with nuisance parameters, including treatment effect estimation, policy learning, and related problems.
- **Intern, Microsoft Research NYC**, June 2017 - Aug. 2017.
Hosts: Robert Schapire and Miroslav Dudik
Developed algorithms and analysis for contextual bandits in new model where learner has access to regression oracle.
- **Visiting Graduate Student, Simons Institute for the Theory of Computing, UC Berkeley**, Jan. - May 2017.
Foundations of Machine Learning Program
- **Intern, Google Research NYC**, June 2016 - Aug. 2016.
Host: Sanjiv Kumar
Researched generalization tradeoffs in optimization for deep learning and adaptive algorithms for online optimization.
- **Research Intern, MIT Lincoln Laboratory**, June 2014 - Aug. 2014.
Host: Michael Boulet
Researched obstacle avoidance approaches for small UAVs. Extended techniques for probabilistic optical flow and scene reconstruction to meet project needs. Developed statistical model for custom imaging device.
- **Research Intern, NASA Jet Propulsion Laboratory**, June 2013 - Aug. 2013.
Principal Investigator: Dr. Jeff Norris
Helped develop system utilizing motion capture and the Oculus Rift to allow users to explore Mars terrain point clouds captured by the Curiosity rover.
- **Undergraduate Researcher, Interaction Lab, USC**, Feb. 2012 - May 2014.
Principal Investigator: Prof. Maja Matarić
Developed optimization approach to retargeting abstract gestures between robot platforms and successfully implemented this algorithm for the PR2 and Bandit robot platforms.

Further Employment

- **Guidance, Navigation, and Controls Intern, SpaceX**, May 2012 - Aug. 2012. Developed software and algorithms to facilitate the development of custom sensing solutions for spacecraft navigation. This work included an algorithm for LiDAR calibration, a network-based star tracker camera driver, and software for testing camera firmware and star tracker algorithms.

Invited Talks

1. **Beyond UCB: Optimal and Efficient Contextual Bandits with Regression Oracles:** Stochastics and Statistics Seminar, MIT, September 2020; International Conference on Machine Learning, June 2020; Learning Theory Seminar, Google NYC, March 2020; Spotlight talk, New York Academy of Science Machine Learning Symposium, March 2020; Young Researcher Seminar Series, Toyota Technological Institute at Chicago, Feb. 2020.

2. **Logarithmic Regret for Adversarial Online Control:** International Conference on Machine Learning, June 2020.
3. **The Complexity of Non-Convex Stochastic Optimization:** Foundations of Deep Learning Reunion Workshop, Simons Institute, UC Berkeley, Aug. 2020; Machine Learning Seminar, University of Washington, December 2019.
4. **Model Selection for Contextual Bandits:** Open Problem Session, COLT 2020; Spotlight Presentation, NeurIPS 2019.
5. **The Complexity of Making the Gradient Small in Stochastic Convex Optimization:** Google ML Theory Workshop, NYC, September 2019.
6. **Statistical Learning with a Nuisance Component:** Best paper track, IJCAI 2020, June 2020; Conference on Learning Theory, June 2019.
7. **Sum-of-Squares Meets Square Loss: Fast Rates for Agnostic Tensor Completion:** Conference on Learning Theory, June 2019.
8. **Distributed Learning with Sublinear Communication:** Long talk, International Conference on Machine Learning, June 2019; Spotlight talk, New York Academy of Science Machine Learning Symposium, March 2019.
9. **Contextual Bandits with Surrogate Losses: Margin Bounds and Efficient Algorithms:** Reinforcement Learning Reading Group, Microsoft Research, August 2018.
10. **Practical Contextual Bandits with Regression Oracles:** Long talk, International Conference on Machine Learning, July 2018.
11. **Logistic Regression: The Importance of Being Improper:** Stochastics and Statistics Seminar, MIT, April 2019; Machine Learning and Friends Lunch, UMass Amherst, April 2019; Information Theory and its Applications workshop (ITA 2019), February 2019; Cornell Computer Science Colloquium, November, 2018; MIFODS Meeting, MIT, October 2018; Microsoft Research AI Breakthroughs Workshop, September 2018; Conference on Learning Theory, July 2018.
12. **Online Learning: Sufficient Statistics and the Burkholder Method:** Algorithms and Complexity Seminar, MIT, October 2018; StatML Reading Group, Carnegie Mellon University, September 2018; Conference on Learning Theory, July 2018; Theory Seminar, Cornell University, April 2018.
13. **Parameter-Free Online Learning via Model Selection:** Spotlight talk, New York of Science Machine Learning Symposium, March 2018; Spotlight presentation, NIPS 2017, December 2017; AI Seminar, Cornell University, November 2017.
14. **Spectrally-Normalized Margin Bounds for Neural Networks:** Machine Learning Reading Group, Microsoft Research New England, June 2018; Spotlight talk, New York of Science Machine Learning Symposium, March 2018; Machine Learning Seminar, Cornell University, October 2017; Machine Learning Seminar, Georgia Tech, October 2017.
15. **ZigZag: A New Approach to Adaptive Online Learning:** Spotlight talk, New York Academy of Science Machine Learning Symposium, March 2018; Conference on Learning Theory, July 2017; Machine Learning Reading Group, MSR NYC, July 2017; Simons Institute for Theory of Computing, March 2017.
16. **Inference in Sparse Graphs with Pairwise Measurements and Side Information:** Theory Lunch, Stanford University, March 2017; Theory Seminar, Cornell University, October 2016; Machine Learning Seminar, New York University, October 2016; Algorithms Seminar, Google Research NYC, August 2016.
17. **Adaptive Online Learning:** Machine Learning Seminar, Google Research NYC, June 2016; Learning from Easy Data II Workshop, NIPS 2015, December 2015; Spotlight Presentation, NIPS 2015, December 2015; Computer Science Theory Retreat, Cornell University, May 2015.

Professional Service

Conference reviewing: COLT, NeurIPS, ICML, STOC, FOCS, SODA, AAAI, AISTATS, ALT, ISIT

Program Committee: COLT 2020, 2021 (Senior PC), NeurIPS 2020 (Area chair), ALT (2019, 2020, 2021), Learning for Dynamics and Control (L4DC) (2020, 2021)

Journal reviewing: JMLR (Editorial Board), Journal of the ACM, Mathematics of Operations Research, Biometrika, Operations Research

Organization:

- **Program on Learning in Games**, Simons Institute for Theory of Computing (occurring Spring 2022), *Organizer*. with Constantinos Daskalakis, Michael Jordan, Christos Papadimitriou, Georgios Piliouras, Vasilis Syrgkanis, and Eva Tardos.

- RL with Function Approximation Reading Group, Simons Institute (RL Theory Program), Fall 2020, *Organizer*. with Lin Chen and Vidya Muthukumar.
- Reinforcement Learning Theory Reading Group, MIT, 2019-2020, *Organizer*. with Alexander Rakhlin.
- Computer Science PhD Student Visit Day, Cornell University, 2015, *Organizer*.

Teaching

Teaching Assistant Machine Learning Theory Professor Karthik Sridharan	Cornell University Spring 2018
---	-----------------------------------

Teaching Assistant Introduction to Analysis of Algorithms Professors David Steurer and Eva Tardos <i>Received outstanding teaching award.</i>	Cornell University Spring 2015
---	-----------------------------------

Teaching Assistant Foundations of Artificial Intelligence Professor Bart Selman	Cornell University Fall 2014
--	---------------------------------

Guest Lectures

- Mirror Descent, Machine Learning Theory (course by Karthik Sridharan), Cornell University, 2018;
- Rademacher Complexity, Machine Learning Theory (course by Karthik Sridharan), Cornell University, 2018;
- VC Dimension, Graduate Machine Learning (course by Jacob Abernethy), Georgia Tech, October 2017.