

Bradly C. Stadie
Email: bstadie@ttic.edu

Education

University of Chicago

BA in Mathematics (Honors)

Chicago, IL
September 2010 - June 2014

University of California, Berkeley

Ph.D in Statistics

Berkeley, CA
August 2014 - August 2018

Employment

TTIC

Research Assistant Professor

Chicago, USA
September 2020 - Present

- Research Assistant Professor at the Toyota Technological Institute at Chicago.

Vector Institute

Postdoctoral Scholar

Toronto, Canada
August 2018 - July 2020

- Postdoctoral researcher at the Vector Institute at the University of Toronto

Open AI

Research Scientist

San Francisco, USA
March 2016 - November 2017

- Full-time research scientist at Open AI
- Advised by Ilya Sutskever

AI Expert at Y Combinator

Part-time Partner

San Francisco, USA
January 2017 - January 2018

- Advised Y Combinator on the latest trends in AI
- Advised over 50 startups on AI and ML problems

Young Scholars Program/SESAME

Counselor

Chicago, USA
June 2011 - May 2013

- Taught mathematics to Chicagoland high school students and CPS teachers
- Full-time during summers. Every Saturday and Sunday during school year
- Head counselor from June 2012 - May 2013

Conference Publications

1. *Third-Person Imitation Learning*

Bradly C. Stadie, Pieter Abbeel, Ilya Sutskever

In the proceedings of the International Conference on Learning Representations (ICLR), Toulon, France, April 2017

arXiv: <https://arxiv.org/pdf/1703.01703.pdf>

2. *One-Shot Imitation Learning*

Yan (Rocky) Duan, Marcin Andrychowicz, Bradly C. Stadie, Jonathan Ho, Jonas Schneider, Ilya Sutskever, Pieter Abbeel, Wojciech Zaremba

In Neural Information Processing Systems (NeurIPS), Long Beach, California, December 2017

arXiv: <https://arxiv.org/pdf/1703.07326.pdf>

3. *Some Considerations on Learning to Explore via Meta-Reinforcement Learning*

Bradly C. Stadie, Ge Yang, Rein Houthoofd, Xi Chen, Yan Duan, Yuhuai Wu, Pieter Abbeel, Ilya Sutskever

In Neural Information Processing Systems (NeurIPS), Montreal, Canada, December 2018

arXiv: <https://arxiv.org/pdf/1803.01118.pdf>

4. *Evolved Policy Gradients*

Rein Houthoofd, Richard Y. Chen, Phillip Isola, Bradly C. Stadie, Filip Wolski, Jonathan Ho, Pieter Abbeel
In Neural Information Processing Systems (NeurIPS) [Spotlight], Montreal, Canada, December 2018
arXiv: <https://arxiv.org/pdf/1802.04821.pdf>

5. *One-Shot Pruning of Recurrent Neural Networks by Jacobian Spectrum Evaluation*

Matthew Zhang, Bradly C. Stadie
In International Conference on Learning Representations (ICLR), Addis Ababa, Ethiopia, April 2020
ArXiv: <https://arxiv.org/pdf/1912.00120.pdf>

6. *Learning Intrinsic Rewards as a Bi-Level Optimization Problem*

Lunjun Zhang, Bradly C. Stadie, Jimmy Ba
Conference on Uncertainty in Artificial Intelligence (UAI), July 2020. Available at: https://www.auai.org/uai2020/proceedings/66_main_paper.pdf

7. *Maximum Entropy Gain Exploration for Long Horizon Multi-goal Reinforcement Learning*

Harris Chan, Silviu Pitis, Stephen Zhao, Bradly C. Stadie, Jimmy Ba.
In International Conference on Machine Learning (ICML), July 2020. ArXiv: <https://arxiv.org/abs/2007.02832>

8. *World Model as a Graph: Learning Latent Landmarks for Planning*

Lunjun Zhang, Ge Yang, Bradly C. Stadie
In International Conference on Machine Learning (ICML) [Long Presentation], July 2021. <https://arxiv.org/pdf/2011.12491.pdf>.

Workshop Publications

1. *Incentivizing Exploration in Reinforcement Learning with Deep Predictive Models*

Bradly C. Stadie, Sergey Levine, Pieter Abbeel
In Neural Information Processing Systems (NeurIPS) Deep RL Workshop, Montreal, Canada, December 2015
arXiv: <https://arxiv.org/pdf/1507.00814.pdf>

2. *Transfer Learning for Estimating Causal Effects Using Neural Networks.*

Bradly C. Stadie, Soeren R. Kuenzel, Nikita Vemuri, Varsha Ramakrishnan, Jasjeet S. Sekhon, Pieter Abbeel
INFORMS Annual Meeting ML and causal inference workshop (2019)
arXiv: <https://arxiv.org/pdf/1808.07804.pdf>

3. *Learning to Learn from Flawed, Failed, and Figurative Demonstrations*

Ge Yang, Bradly C. Stadie, Roberto Calandra, Pieter Abbeel, Sergey Levine, Chelsea Finn
In Neural Information Processing Systems (NeurIPS) Deep RL workshop [Spotlight], Montreal, Canada, December 2018.
Available at: <https://github.com/bstadie/All-Bradly-Stadie-Papers/blob/master/leaf.pdf>

Preprints

1. *Simulating the Stochastic Dynamics and Cascade Failure of Power Networks*

Charles Matthews, Bradly C. Stadie, Jonathan Weare, Mihai Anitescu, Christopher Demarco
Submitted to IEEE Transactions on Power Systems. Under Review.
arXiv: <https://arxiv.org/pdf/1806.02420.pdf>

2. *One-Demonstration Imitation Learning*

Bradly C. Stadie, Siyan Zhao, Qiqi Xu, Bonnie Li, Lunjun Zhang
Preprint. Available at: <https://github.com/bstadie/All-Bradly-Stadie-Papers/blob/master/one-demo.pdf>

Invited Talks

- **Northwestern University invited speaker (October 2019):** Learning reward functions.
- **INFORMS panel on ML in causal inference (October 2019):** Transfer learning causal effects.
- **University of Illinois, Urbana Champaign (October 2019):** Learning reward functions.
- **Toyota Technological Institute, Chicago (May 2019):** Learning from imperfect data.
- **University of Wisconsin, Madison (April 2019):** Learning from imperfect data.
- **Purdue University (April 2019):** Learning from imperfect data.
- **University of Montreal (February 2018):** The relationship between sampling and meta learning.
- **University of Toronto (February 2018):** The relationship between sampling and meta learning.
- **Harvard University 2017 Big Data Conference (August 2017):** Frontiers of meta learning.
- **Open AI (March 2016):** Exploration in deep reinforcement learning.

Teaching

- **Statistics 133: Concepts in Computing with Data.** Graduate Student Instructor. *Jan - May 2015*
- **Statistics 088: Introduction to Data Science.** Graduate Student Instructor. *Jan - May 2016*
- **TTIC 31170 - Planning, Learning, and Estimation for Robotics and Artificial Intelligence.** Co-Instructor. *March 2021 - June 2021*