Nathan Lambert

nathan@natolambert.com • natolambert.com • natolambert
natolambert • natolambert • Last updated on February 2, 2022

I research **data-driven decision making**, including progressing **model-based reinforcmement learning** algorithms, applying them to real-world problems such as **robotics**, and planning for the **societal implications** of these technologies.

Education

Ph.D. in Computer Science, University of California, Berkeley (4.0/4.0) 2017 – 2022
Advisors: Kristofer S.J. Pister, Roberto Calandra

B.S. in Electrical and Computer Engineering, Cornell University (4.0/4.0) 2013 – 2017

Industry Experience

DeepMind, London (*Virtual*) | Research Intern, (Host: Martin Riedmiller) 2021
Facebook AI, Menlo Park | Research Intern and Student Researcher (Host: Roberto Calandra) 2019 – 2020
Tesla, Palo Alto Test Engineering Intern 2015

Honors & Awards

Best Student Paper Finalist; IEEE Symposium on Multi-Robot and Multi-Agent Systems 2021
Berkeley EECS Demetri Angelakos Memorial Achievement Award 2021
Heart to Humanity Eternal (H2H8) Pioneer 2021
NDSEG Graduate Research Fellowship Program Top 200 2018
NSF Graduate Research Fellowship Program Honorable Mention 2017, 2018
Berkeley EECS Department Fellowship 2017
Eight undergraduate scholarships 2013 – 2017
Cornell Rowing Charles E. Courtney Award, Tau Beta Pi Scholarship, Southeastern New England Defense Industry Alliance STEM Scholarship 2016, 2017, Cornell Athletics 400 Club Induction, Beta Pi Induction, Eta Kappu Nu Induction, American Society of Engineering Education SMART Scholar Award

Publications [Google Scholar]

Representative publications that I am a primary author on are highlighted.

2022

1. *Reward Reports for Reinforcement Learning*
   Thomas Gilbert, Sarah Dean, **Nathan Lambert**, Tom Zick, and Aaron Snoswell
   Under Review 2022

2. *Choices, Risks, and Reward Reports: Charting Public Policy for Reinforcement Learning Systems*
   Thomas Gilbert, Sarah Dean, Tom Zick, and **Nathan Lambert**
   Center for Long-Term Cybersecurity Whitepaper Series 2022

3. *Investigating Compounding Prediction Errors in One-step Dynamics Models*
   Nathan Lambert, Roberto Calandra, and Kristofer Pister
   Under Review 2022

4. *Understanding the Challenges of Exploration for Offline Reinforcement Learning*
   Nathan Lambert, Markus Wulfmeier, Arunkumar Byravan, Michael Bloesch, William Whitney, Vibhavari Dasagi, Tim Hertweck, and Martin Riedmiller
   arXiv Preprint 2022

5. *BLISS: Interplanetary Exploration with Swarms of Low-Cost Spacecraft*
   Alexander Alvara*, **Nathan Lambert***, Emmanuel Sin*, Lydia Lee*, Beau Kuhn, Andrew Westphal, and Kristofer Pister
   Under Review 2022 (*co-lead authors)

2021

6. *MBRL-Lib: A Modular Library for Model-based Reinforcement Learning* [code]
Luis Pineda, Brandon Amos, Amy Zhang, **Nathan Lambert**, and Roberto Calandra
arXiv Preprint 2021

7. *BotNet: A Simulator for Studying the Effects of Accurate Communication Models on High-agent-count Multi-agent Control* [code]
Mark Selden, Felipe Campos, Jason Zhou, **Nathan Lambert**, Daniel Drew, and Kristofer Pister
Symposium on Multi-Agent and Multi-Robot Systems 2021 (Best Student Paper Finalist)

8. *Axes for Sociotechnical Inquiry in AI Research*
Sarah Dean, Thomas Krendl Gilbert, **Nathan Lambert**, and Tom Zick
Transactions on Technology and Society (TTS) 2021 (Authors arranged alphabetically)

9. *On the Importance of Hyperparameter Optimization for Model-based Reinforcement Learning*
Baohe Zhang, Raghu Rajan, Luis Pineda, **Nathan Lambert**, André Biedenkapp, Kurtland Chua, Frank Hutter, and Roberto Calandra
International Conference on Artificial Intelligence and Statistics (AISTATS) 2021

10. *Learning Accurate Long-term Dynamics for Model-based Reinforcement Learning* [code]
Nathan Lambert, Albert Wilcox, Howard Zhang, Kristofer SJ Pister, and Roberto Calandra
International Conference on Decision and Control (CDC) 2021

11. *Nonholonomic Yaw Control of an Underactuated Flying Robot With Model-Based Reinforcement Learning*
Nathan Lambert, Craig Schindler, Daniel Drew, and Kristofer Pister
Robotics and Automation Letters (RAL) 2021

2020

12. *Objective Mismatch in Model-based Reinforcement Learning*
Nathan Lambert, Brandon Amos, Omry Yadan, and Roberto Calandra
Conference on Learning for Decision and Control (L4DC) 2020

13. *AI Development for the Public Interest: From Abstraction Traps to Sociotechnical Risks*
McKane Andrus, Sarah Dean, Thomas Gilbert, **Nathan Lambert**, and Tom Zick
International Symposium on Technology and Society (ISTATS) 2020 (Authors arranged alphabetically)

14. *Learning for Microrobot Exploration: Model-based Locomotion, Robust Navigation, and Low-Power Deep Classification*
Nathan Lambert, Fahran Toddywala, Brian Liao, Eric Zhu, Lydia Lee, and Kristofer Pister
International Conference on Manipulation, Automation and Robotics at Small Scales (MARSS) 2020

15. *Learning Generalizable Locomotion Skills with HierarchicalReinforcement Learning*
Tianyu Li, **Nathan Lambert**, Roberto Calandra, Akshara Rai, and Franziska Meier
International Conference on Robotics and Automation (ICRA) 2020

2019

16. *Low-Level Control of a Quadrotor With Deep Model-Based Reinforcement Learning* [code]
Nathan Lambert, Daniel Drew, Joseph Yaconelli, Sergey Levine, Roberto Calandra, and Kristofer Pister
Robotics and Automation Letters (RAL) 2019

2018

17. *Toward Controlled Flight of the Ionocraft: A Flying Microrobot Using Electrohydrodynamic Thrust With Onboard Sensing and No Moving Parts*
Daniel S Drew, **Nathan Lambert**, Craig B Schindler, and Kristofer Pister
Robotics and Automation Letters (RAL) 2018

2017

18. *Enhanced lithium niobate pyroelectric ionizer for chip-scale ion mobility-based gas sensing*
K. B. Vinayakumar, V. Gund, **Nathan Lambert**, S. Lodha, and A. Lal
Sensors 2017

Repositories

facebookresearch/mbrl-lib | ★564 | *Model-based reinforcement learning library* 2021

Invited Talks

Cornell Robotics Seminar (Talk available, slides) March 2021
Improving Model Predictive Control in Model-based Reinforcement Learning

UC Berkeley Semiautonomous Seminar (Talk available, slides) April 2020
Model Learning for Low-level Control in Robotics

UC Berkeley Semiautonomous Seminar May 2019

Mentorship

Mark Selden (UC Berkeley BS '22) 2020
Albert Wilcox (UC Berkeley BS '22) 2019
Jason Zhou (UC Berkeley BS, MS '21 to Matician) 2019
Felipe Campos (UC Berkeley BS '20 to Armstrong Robotics) 2018
Howard Zhang (UC Berkeley BS, MS'21 to UCLA PhD) 2018

Peer Review

Conference on Machine Learning (ICML) 2020, 2022
Conference on Learning Representations (ICLR) (*Outstanding Reviewer) 2021*, 2022
Conference on Robot Learning (CORL) 2020
Conference on Robotics and Automation (ICRA) 2020, 2021, 2022
Conference on Intelligent Robots and Systems (IROS) 2021
Conference on Decision and Control (CDC) 2021
Robotics and Automation Letters (RA-L) 2019, 2020
Transactions on Cybernetics 2020

Professional Activities

Member of Well-Being in Machine Learning 2021 – cont.
NeurIPs Workshop on Robot Learning Organizor 2021
Tapia Panel on Student Mental Health Organizor 2021
Founder of UC Berkeley EECS Equal Access to Application Assistance (EAAA) Program 2020 – 2022
Wellness Coordinator for UC Berkeley Electrical Engineering Graduate Student Assembly (EEGSA) 2020 – 2022
Bay Area Teachers in Schools 2017

Teaching

Introduction to Artificial Intelligence (UCB CS188), TA Su2020, Fa2020
Introduction to Artificial Intelligence (UCB CS188), Instructor *lectured to 800+ students* Sp2020
Designing Information Devices and Systems II (UCB EE16B), TA Fa2019
Integrated Micro Sensors and Actuators (Cornell ECE4320), Grader Sp2017
Mathematics of Signal and System Analysis (Cornell ECE 3250), TA Fa2016

Skills

Programming	Python
Frameworks	JAX, NumPy, Pandas, PyTorch, SciPy
Tools	Linux, vim, git, tmux