Jonathan N. Lee

Stanford University
Computer Science Department

jnl@stanford.edu
https://www.jonathannlee.com/

Education

Ph.D. Computer Science (Machine Learning), 2019—
Stanford University

B.S. (Highest Honors) Electrical Engineering & Computer Science, 2019
University of California, Berkeley
Honors Concentration: Mathematics & Statistics
GPA: 3.96

Research

Stanford AI Lab, 2020—
PhD Student
Advisor: Emma Brunskill
Area: reinforcement learning, foundation models for decision-making

Google Research, Learning Theory Team, 2022
Student Researcher
Hosts: Christoph Dann, Alekh Agarwal, Tong Zhang
Area: reinforcement learning

Google Brain (now Google DeepMind), 2021–2022
Research Intern & Student Researcher
Hosts: George Tucker, Ofir Nachum, Bo Dai
Area: offline reinforcement learning

UC Berkeley AUTOLab, 2015—2019
Undergraduate Researcher
Advisor: Ken Goldberg
Area: robot learning, imitation learning, optimization

Honors

NSF Graduate Research Fellowship
The Ravi Family Graduate Fellowship
CRA Outstanding Undergraduate Researcher Award 2019 Finalist
NeurIPS (2021) Outstanding Reviewer Award
UC Berkeley EECS Honors Degree Program
Eta Kappa Nu EECS Honor Society
Tau Beta Pi Engineering Honor Society
Ford Oval Scholarship
UC Berkeley Kraft Award
UC Berkeley College of Engineering Dean's List
2nd Place Siemens FutureMakers Challenge UC Berkeley.

Industry

Consulting, Machine Learning at Berkeley, 2016
GitHub: Built a pipeline (Lexicon) to classify languages for GitHub's code repositories.
H2O.ai: Built platform to demonstrate H2O machine learning algorithms.

Software Engineering Intern, DreameGGs Funding Club, 2015
Developed machine learning models to predict success of independent movies. Increased accuracy by 5.6%.

Software Engineering Intern, Salient Technology Intl. LLC, 2014
Built www.vstudy.co, a WebRTC-driven video client for student collaboration.

Teaching

Stanford Computer Science, 2022
Advanced Survey of Reinforcement Learning (CS 332) Course Assistant, 2022

UC Berkeley Dept. of EECS, 2017—2018
Machine Learning (CS 189) Teaching Assistant, 2018
Signals and Systems (EE 120) Teaching Assistant, 2017
Machine Learning (CS 189) Tutor/Academic Intern, 2017
Eta Kappa Nu EECS Honor Society Tutor, 2017

Coursework

Graduate Courses: Machine Learning, Information Theory & Statistics, Optimization Theory, Convex Optimization, Nonlinear Systems, Linear System Theory, Optimal and Learning-Based Control, Learning and Optimization, Randomized Algorithms, Theoretical Statistics, Probability Theory, Reinforcement Learning, Information-Theoretic Lower Bounds, Data Mining.
Undergraduate Courses: Machine Learning, Artificial Intelligence, Convex Optimization, Probability and Random Processes, Signals and Systems, Feedback Control, Real Analysis, Abstract Algebra, Linear Algebra, Multivariate Calculus, Data Structures and Algorithms, Operating Systems.

Service

Stanford Machine Learning Lunch Organizer, Stanford CS PhD Advisory Council. Machine Learning at Berkeley. The Triple Helix at Berkeley. Reviewer for NeurIPS, ICML, ICLR, AISTATS, JMLR, TMLR, IEEE Transactions on Information Theory.

Preprints

Information-Directed Search for Formal Reasoning with Large Language Models.
Yash Chandak, **Jonathan Lee**, Emma Brunskill.
In preparation.

Estimating Optimal Policy Value in General Linear Contextual Bandits.
Jonathan Lee, Weihao Kong, Aldo Pacchiano, Vidya Muthukumar, Emma Brunskill.
To appear in Transactions on Machine Learning Research (TMLR), 2024.

Publications

Supervised Pretraining Can Learn In-Context Reinforcement Learning.
Jonathan Lee*, Annie Xie*, Aldo Pacchiano, Yash Chandak, Chelsea Finn, Ofir Nachum, Emma Brunskill.
Conference on Neural Information Processing Systems (NeurIPS), 2023.
Accepted for spotlight.

Experiment Planning with Function Approximation.
Aldo Pacchiano, **Jonathan Lee**, Emma Brunskill.
Conference on Neural Information Processing Systems (NeurIPS), 2023.

Learning in POMDPs is Sample-Efficient with Hindsight Observability.
Jonathan Lee, Alekh Agarwal, Christoph Dann, Tong Zhang.
International Conference on Machine Learning (ICML), 2023.

Dueling RL: Reinforcement Learning with Trajectory Preferences.
Aldo Pacchiano, Aadirupa Saha, **Jonathan Lee**.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2023.

Oracle Inequalities for Model Selection in Offline Reinforcement Learning
Jonathan Lee, George Tucker, Ofir Nachum, Bo Dai, Emma Brunskill.
Conference on Neural Information Processing Systems (NeurIPS), 2022.

Model Selection in Batch Policy Optimization
Jonathan Lee, George Tucker, Ofir Nachum, Bo Dai.
International Conference on Machine Learning (ICML), 2022.

Design of Experiments for Stochastic Contextual Linear Bandits.
Andrea Zanette*, Kefan Dong*, **Jonathan Lee***, Emma Brunskill.
Conference on Neural Information Processing Systems (NeurIPS), 2021.

Near Optimal Policy Optimization via REPS.
Aldo Pacchiano, **Jonathan Lee**, Peter Bartlett, Ofir Nachum.
Conference on Neural Information Processing Systems (NeurIPS), 2021.

Online Model Selection for Reinforcement Learning with Function Approximation.
Jonathan Lee, Aldo Pacchiano, Vidya Muthukumar, Weihao Kong, Emma Brunskill.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2021.

Dynamic Regret Convergence Analysis and an Adaptive Regularization Algorithm for On-Policy Robot Imitation Learning.
Jonathan Lee, Michael Laskey, Ajay Tanwani, Anil Aswani, Ken Goldberg.
International Journal of Robotics Research (IJRR), 2021.

Accelerated Message Passing for Entropy-Regularized MAP Inference.
Jonathan Lee, Aldo Pacchiano, Peter Bartlett, Michael I. Jordan.
International Conference on Machine Learning (ICML), 2020.

Online Learning with Continuous Variations: Dynamic Regret and Reductions.
Ching-An Cheng*, **Jonathan Lee***, Ken Goldberg, Byron Boots.
International Conference on Artificial Intelligence and Statistics (AISTATS), 2020.

Convergence Rates of Smooth Message Passing with Rounding in Entropy-Regularized MAP Inference.

Jonathan Lee*, Aldo Pacchiano*, Michael I. Jordan
International Conference on Artificial Intelligence and Statistics (AISTATS), 2020.

On-Policy Robot Imitation Learning from a Converging Supervisor.
Ashwin Balakrishna*, Brijen Thananjeyan*, **Jonathan Lee**, Arsh Zahed, Felix Li, Joseph Gonzalez, Ken Goldberg.
Conference on Robot Learning (CoRL), 2019.
Accepted for long oral (5% of papers).

A Dynamic Regret Analysis and Adaptive Regularization Algorithm for On-Policy Robot Imitation Learning.
Jonathan Lee, Michael Laskey, Ajay Tanwani, Anil Aswani, Ken Goldberg.
Algorithmic Foundations of Robotics, 2018.
Invited to special issue in the International Journal of Robotics Research.

Generalizing Robot Imitation Learning with Invariant Hidden Semi-Markov Models.
Ajay Tanwani, **Jonathan Lee**, Brijen Thananjeyan, Michael Laskey, Sanjay Krishnan, Roy Fox, Ken Goldberg, Sylvain Calinon.
Algorithmic Foundations of Robotics, 2018.
Invited to special issue in the International Journal of Robotics Research.

Constraint Estimation and Derivative-Free Recovery for Robot Learning from Demonstrations.
Jonathan Lee, Michael Laskey, Roy Fox, Ken Goldberg.
Conference on Automation Science and Engineering (CASE), 2018

DART: Noise Injection for Robust Imitation Learning.
Michael Laskey, **Jonathan Lee**, Roy Fox, Anca Dragan, Ken Goldberg.
Conference on Robot Learning (CoRL), 2017.

Comparing Human-Centric and Robot-Centric Sample Efficiency for Robot Deep Learning from Demonstrations.
Michael Laskey, Caleb Chuck, **Jonathan Lee**, Jeffrey Mahler, Sanjay Krishnan, Kevin Jamieson, Anca Dragan, Ken Goldberg.
International Conference on Robotics and Automation (ICRA), 2017.

Robot Grasping in Clutter: Using a Hierarchy of Supervisors for Learning from Demonstrations.
Michael Laskey, **Jonathan Lee**, Caleb Chuck, David Gealy, Wesley Hsieh, Florian Pokorny, Anca Dragan, Ken Goldberg.
Conference on Automation Science and Engineering (CASE), 2016.

Workshop papers

Improved Estimator Selection for Off-Policy Evaluation
George Tucker, **Jonathan Lee**.
ICML Workshop on Reinforcement Learning Theory, 2021.

Continuous Online Learning and New Insights to Online Imitation Learning

Jonathan Lee*, Ching-An Cheng*, Ken Goldberg, Byron Boots.
NeurIPS Optimization in Reinforcement Learning Workshop, 2019.
Best paper award.

Stability Analysis of On-Policy Imitation Learning Algorithms Using Dynamic Regret.
Jonathan Lee, Michael Laskey, Ajay Tanwani, Ken Goldberg.
Robotics: Science and Systems (RSS) Workshop on Imitation and Causality, 2018.
Accepted for spotlight.

Iterative Noise Injection for Scalable Imitation Learning.
Michael Laskey, **Jonathan Lee**, Wesley Hsieh, Richard Liaw, Jeffrey Mahler, Roy Fox, Ken Goldberg.
arXiv, 2016.