

Max Simchowitz

CONTACT 8004 Gates Hillman Complex msimchow@andrew.cmu.edu
4902 Forbes Ave, Pittsburgh, PA 15213 <https://msimchowitz.github.io/>

APPOINTMENTS **Carnegie Mellon University**

Assistant Professor, *Machine Learning Department* Jan 2025 - Current

Amazon

Amazon Scholar, *Frontier AI & Robotics (FAR)* 2026 - Current

Massachusetts Institute of Technology

Postdoctoral Fellow, *Robot Locomotion Group, CSAIL* Sept 2021 - Sept 2024

EDUCATION **University of California, Berkeley**, Berkeley, CA, USA **2015 - 2021**

Ph.D., *Electrical Engineering and Computer Science*

Advisors: Benjamin Recht and Michael I. Jordan

Princeton University, Princeton, NJ, USA **June 2015**

A.B., *Mathematics*, Summa Cum Laude.

SELECTED **Best Paper Finalist**, *International Conference on Robotics and Automation (ICRA)*, 2024.

AWARDS **Best Paper Finalist**, *Robotics Science and Systems (RSS)*, 2023.

Outstanding Paper Award, *International Conference on Machine Learning (ICML)*, 2022.

Open Philanthropy AI Fellowship, Good Ventures Foundation, 2019.

Best Paper Award, *International Conference on Machine Learning (ICML)*, 2018.

Berkeley Fellowship University of California, Berkeley and by the Rose Hill Foundation, 2018.

Tong Leong Lim Pre-Doctoral Prize, University of California, Berkeley, 2017.

NSF GRFP Fellowship, NSF, 2015-2017.

George B. Covington Prize for Excellence in Mathematics, Princeton University, 2015.

Phi Beta Kappa, Early Induction (Top 2% of Graduating Class), Princeton University, 2014.

Shapiro Prize for Academic Excellence, Princeton University, 2012 & 2013.

Manfred Pyka Memorial Prize for Physics, Princeton University, 2012.

Quin Morton '36 Prize for Writing, Princeton University, 2012.

FUNDING AND **Jane Street Award Renewal**, Jane Street Trading, 2025.

GIFTS **Jones Interdisciplinary Seed Fund**, Mellon College of Sciences, 2026

University 3.0 Funding, Toyota Research Institute, 2025

Jane Street Award, Jane Street Trading, 2025.

University 2.0 Funding, Toyota Research TInstitute, 2024.

Google Robotics Award, Google, 2024.

SELECTED

PUBLICATIONS “OGPO: Sample Efficient Full-Finetuning of Generative Control Policies,” Sarvesh Patil, Mitsuhiro Nakamoto, Manan Agarwal, Shashwat Saxena, Jesse Zhang, Giri Anantharaman, Cleah Winston, Chaoyi Pan, Douglas Chen, Nai-Chieh Huang, Zeynep Temel, Oliver Kroemer, Sergey

Levine, Abhishek Gupta, Hongkai Dai, Paarth Shah, **Max Simchowitz**. *International Conference on Machine Learning (ICML)*, 2026. ArXiv: 2605.03065.

“Workspace Models: Lightweight Robotic Memory via Saliency-Driven Supervision,” Nitish Dashora, Douglas Chen, Idan Shenfeld, John Marangola, Pulkit Agrawal, **Max Simchowitz**. Preprint, 2026. ArXiv: 2609.20820.

“Much Ado About Noising: Dispelling the Myths of Generative Robotic Control.” Chaoyi Pan, Giri Anantharaman, Nai-Chieh Huang, Claire Jin, Daniel Pfrommer, Chenyang Yuan, Frank Permenter, Guannan Qu, Nicholas Boffi, Guanya Shi, **Max Simchowitz**. *International Conference on Representation Learning (ICLR)*, 2025. ArXiv: 2512.01809

“Diffusion Policy Policy Optimization,” Allen Z. Ren, Justin Lidard, Lars L. Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, **Max Simchowitz**. *International Conference on Representation Learning (ICLR)*, 2024. ArXiv: 2409.00588.

“Diffusion Forcing: Next-token Prediction Meets Full-Sequence Diffusion,” Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, Vincent Sitzmann. *Conference on Neural Informtion Processing Systems (NeurIPS)*, 2024. ArXiv: 2310.01362

“Do Differentiable Simulators Give Better Policy Gradients?” H.J. Terry Suh, **Max Simchowitz**, Kaiqing Zhang, Russ Tedrake. *International Conference on Machine Learning (ICML)*, 2022. **Out-standing Paper Award**.

“Naive Exploration is Optimal for Online LQR”, **Max Simchowitz**, Dylan Foster. *International Conference on Machine Learning (ICML)*, 2020. ArXiv: 2001.09576.

“Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification,” **Max Simchowitz**, Horia Mania, Stephen Tu, Benjamin Recht, Michael I. Jordan. *Conference on Learning Theory (COLT)*, 2018. ArXiv: 1802.08334

AWARD-

WINNING

PUBLICATIONS

“Constrained Bimanual Planning with Analytic Inverse Kinematics,” Thomas Cohn, Seiji Shaw, **Max Simchowitz**, Russ Tedrake. *International Conference on Robotics and Automation (ICRA)*, 2024. **Best Paper Finalist**. ArXiv: <https://arxiv.org/abs/2310.01362>

“Non-Euclidean Motion Planning with Graphs of Geodesically-Convex Sets,” Thomas Cohn, Mark Petersen, **Max Simchowitz**, Russ Tedrake. *Robotics Science and Systems (RSS)*, 2023. ArXiv: 2305.06341. **Best Paper Finalist**.

“Do Differentiable Simulators Give Better Policy Gradients?” H.J. Terry Suh, **Max Simchowitz**, Kaiqing Zhang, Russ Tedrake. *International Conference on Machine Learning (ICML)*, 2022. **Out-standing Paper Award**.

“Balancing Competing Objectives for Welfare-Aware Machine Learning with Imperfect Data,” Esther Rolf, **Max Simchowitz**, Sarah Dean, Lydia T. Liu, Daniel Björkegren, Moritz Hardt, Joshua Blumenstock. *The NeurIPS Joint Workshop on AI for Social Good*., 2019. **Best paper award**.

“Delayed Impact of Fair Machine Learning,” Lydia T. Liu, Sarah Dean, Esther Rolf, **Max Simchowitz**, Moritz Hardt. *International Conference on Machine Learning (ICML)*, 2018. **Best Paper Award**.

JOURNAL

PUBLICATIONS

“Non-Euclidean Motion Planning with Graphs of Geodesically-Convex Sets,” Thomas Cohn, Mark Petersen, **Max Simchowitz**, Russ Tedrake. *The International Journal of Robotics Research (IJRR)* (invited paper), 2025.

“Corruption Robust Exploration in Episodic Reinforcement Learning,” Thodoris Lykouris, **Max Simchowitz**, Aleksandrs Slivkins, Wen Sun, (alphabetical order). *Mathematics of Operations Research*. [ArXiv: 1911.08689](#).

“Do Differentiable Simulators Give Better Policy Gradients?” H.J. Terry Suh, **Max Simchowitz**, Kaiqing Zhang, Russ Tedrake. *Journal of Machine Learning Research (JMLR), Special Edition*, 2023 (to appear).

“Exploration and Incentives in Reinforcement Learning,” **Max Simchowitz**, Aleksandrs Slivkins. *Operations Research*, 2023 (to appear). [ArXiv: 2103.00360](#)

“First-order Methods Almost Always Avoid Saddle Points,” Jason D. Lee, Ioannis Panageas, Georgios Piliouras, **Max Simchowitz**, Michael I. Jordan, Benjamin Recht”. *Mathematical Programming* (pg. 1-27), 2019. [ArXiv: 1710.07406](#)

CONFERENCE PUBLICATIONS **2026**

“OGPO: Sample Efficient Full-Finetuning of Generative Control Policies,” Sarvesh Patil, Mitsuhiro Nakamoto, Manan Agarwal, Shashwat Saxena, Jesse Zhang, Giri Anantharaman, Cleah Winston, Chaoyi Pan, Douglas Chen, Nai-Chieh Huang, Zeynep Temel, Oliver Kroemer, Sergey Levine, Abhishek Gupta, Hongkai Dai, Paarth Shah, **Max Simchowitz**. *International Conference on Machine Learning (ICML)*, 2026. [ArXiv: 2605.03065](#).

“Diamond Maps: Efficient Reward Alignment via Stochastic Flow Maps,” Peter Holderrieth, Douglas Chen, Luca Eyring, Ishin Shah, Giri Anantharaman, Yutong He, Zeynep Akata, Tommi Jaakkola, Nicholas Matthew Boffi, **Max Simchowitz**. *International Conference on Machine Learning (ICML)*, 2026. [ArXiv: 2602.05993](#).

2025

“Much Ado About Noising: Dispelling the Myths of Generative Robotic Control.” Chaoyi Pan, Giri Anantharaman, Nai-Chieh Huang, Claire Jin, Daniel Pfrommer, Chenyang Yuan, Frank Permenter, Guannan Qu, Nicholas Boffi, Guanya Shi, **Max Simchowitz**. *International Conference on Representation Learning (ICLR)*, 2025. [ArXiv: 2512.01809](#)

“Action Chunking and Exploratory Data Collection Yield Exponential Improvements in Behavior Cloning for Continuous Control.” Thomas T. Zhang, Daniel Pfrommer, Chaoyi Pan, Nikolai Matni, **Max Simchowitz**, *International Conference on Representation Learning (ICLR)*, 2025. [ArXiv: 2507.09061](#)

“e3: Learning to Explore Enables Extrapolation of Test-Time Compute for LLMs.” Amrith Setlur, Matthew Y. R. Yang, Charlie Snell, Jeremy Greer, Ian Wu, Virginia Smith, **Max Simchowitz**, Aviral Kumar. *International Conference on Representation Learning (ICLR)*, 2025. [ArXiv: 2506.09026](#)

“Joint Distillation for Fast Likelihood Evaluation and Sampling in Flow-based Models.” Xinyue Ai, Yutong He, Albert Gu, Ruslan Salakhutdinov, J Zico Kolter, Nicholas Matthew Boffi, **Max Simchowitz**. *International Conference on Representation Learning (ICLR)*, 2025. [ArXiv: 2512.02636](#)

“Is Your Diffusion Model Actually Denoising?” Daniel Pfrommer, Zehao Dou, Christopher Scarvelis, **Max Simchowitz**, Ali Jadbabaie. *Conference on Neural Information Processing Systems (NeurIPS)*, 2025. [ArXiv: 2512.18736](#).

“A Test Function Approach to Incremental Stability.” Daniel Pfrommer, **Max Simchowitz**, Ali Jadbabaie. *Conference on Decision and Control (CDC)*, 2025. [ArXiv: 2507.00695](#)

“The Pitfalls of Imitation Learning when Actions are Continuous” **Max Simchowitz**, Daniel Pfrommer, Ali Jadbabaie. *Conference on Learning Theory (COLT)*, 2025. [ArXiv: 2503.09722](#).

“History-Guided Video Diffusion.” Kiwhan Song, Boyuan Chen, **Max Simchowitz**, Yilun Du, Russ Tedrake, Vincent Sitzmann. *International Conference on Machine Learning (ICML)*, 2025. [ArXiv: 2502.06764](#)

“Diffusion Policy Policy Optimization,” Allen Z. Ren, Justin Lidard, Lars L. Ankile, Anthony Simeonov, Pulkit Agrawal, Anirudha Majumdar, Benjamin Burchfiel, Hongkai Dai, **Max Simchowitz**. *International Conference on Representation Learning (ICLR)*, 2025. [ArXiv: 2409.00588](#).

“Self-Improvement in Language Models: The Sharpening Mechanism” Audrey Huang, Adam Block, Dylan J. Foster, Dhruv Rohatgi, Cyril Zhang, **Max Simchowitz**, Jordan T. Ash, Akshay Krishnamurthy. *International Conference on Representation Learning (ICLR)*, **Oral**, 2025. [ArXiv: 2412.01951](#).

“Is Linear Feedback on Smoothed Dynamics Sufficient for Stabilizing Contact-Rich Plans?” Yuki Shirai, Tong Zhao, H.J. Terry Suh, Huaijiang Zhu, Xinpei Ni, Jiuguang Wang, **Max Simchowitz**, Tao Pang. *International Conference on Robotics and Automation (ICRA)*, 2025. [ArXiv: 2411.06542](#).

2024

“Faster Algorithms for Growing Collision-Free Convex Polytopes in Robot Configuration Space,” Peter Werner*, Thomas Cohn, Rebecca H. Jiang, Tim Niklas Seyde, **Max Simchowitz**, Russ Tedrake, Daniela Rus. *International Symposium on Robotics Research*, to appear. [ArXiv: 2407.01392](#)

“Diffusion Forcing: Next-token Prediction Meets Full-Sequence Diffusion,” Boyuan Chen, Diego Marti Monso, Yilun Du, Max Simchowitz, Russ Tedrake, Vincent Sitzmann. *Conference on Neural Information Processing Systems (NeurIPS)*, 2024. [ArXiv: 2310.01362](#)

“Constrained Bimanual Planning with Analytic Inverse Kinematics,” Thomas Cohn, Seiji Shaw, **Max Simchowitz**, Russ Tedrake. *International Conference on Robotics and Automation (ICRA)*, 2024. **Best Paper Finalist**. [ArXiv: 2310.01362](#)

“Butterfly Effects of SGD Noise: Error Amplification in Behavior Cloning and Autoregression.” Adam Block, Dylan J. Foster, Akshay Krishnamurthy, **Max Simchowitz**, Cyril Zhang, (alphabetical order). 2023. [ArXiv: 2310.11428](#)

“Fleet Policy Learning via Weight Merging and An Application to Robotic Tool-Use.” Lirui Wang, Kaiqing Zhang, Allan Zhou, **Max Simchowitz**, Russ Tedrake. 2023. [ArXiv: 2310.01362](#)

2023

“Provable Guarantees for Generative Behavior Cloning: Bridging Low-Level Stability and High-Level Behavior.” Adam Block, Ali Jadbabie, Daniel Pfrommer, **Max Simchowitz**, Russ Tedrake, (alphabetical order). *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. [ArXiv: 2307.14619](#).

“Smoothed Online Learning for Prediction in Piecewise Affine Systems.” Adam Block, **Max Simchowitz**, Russ Tedrake, (alphabetical order). *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. [ArXiv: 2301.11187](#). **Spotlight**.

“RePo: Resilient Model-Based Reinforcement Learning by Regularizing Posterior Predictability.”

Chuning Zhu, Max Simchowitz, Siri Gadipudi, Abhishek Gupta. *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. [ArXiv: 2309.00082](#). **Spotlight**.

“Non-Euclidean Motion Planning with Graphs of Geodesically-Convex Sets,” Thomas Cohn, Mark Petersen, **Max Simchowitz**, Russ Tedrake. *Robotics Science and Systems (RSS)*, 2023. [ArXiv: 2305.06341](#). **Best Paper Finalist**.

“Oracle-Efficient Smoothed Online Learning for Piecewise Continuous Decision Making,” Adam Block, **Max Simchowitz**, Sasha Rakhlin, (alphabetical order). *Conference on Learning Theory (COLT)*, 2023. [ArXiv: 2302.05430](#).

“Tackling Combinatorial Distribution Shift: A Matrix Completion Perspective,” **Max Simchowitz**, Abhishek Gupta, Kaiqing Zhang. *Conference on Learning Theory (COLT)*, 2023. [ArXiv: 2307.06457](#)

“Statistical Learning under Heterogenous Distribution Shift,” **Max Simchowitz**, Anurag Ajay, Pulkit Agrawal, Akshay Krishnamurthy. *International Conference on Machine Learning (ICML)*, 2023. [ArXiv: 2302.13934](#).

“The Power of Learned Locally Linear Models for Nonlinear Policy Optimization,” Daniel Pfrommer, **Max Simchowitz**, Tyler Westenbroek, Nikolai Matni, Stephen Tu. *International Conference on Machine Learning (ICML)*, 2023. [ArXiv: 2305.09619](#).

“Learning to Extrapolate: A Transductive Approach,” Aviv Netanyahu, Abhishek Gupta, **Max Simchowitz**, Kaiqing Zhang, Pulkit Agrawal. *International Conference on Representation Learning (ICLR)*, 2023. [ArXiv: 2304.14329](#).

2022

“Efficient and Near-Optimal Smoothed Online Learning for Generalized Linear Functions,” Adam Block, **Max Simchowitz**, (alphabetical order). *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. [ArXiv: 2205.13056](#)

“Globally Convergent Policy Search over Dynamic Filters for Output Estimation,” Jack Umenberger, **Max Simchowitz**, Juan C. Perdomo, Kaiqing Zhang, Russ Tedrake. *Conference on Neural Information Processing Systems (NeurIPS)*, 2022. [ArXiv: 2202.11659](#)

“Do Differentiable Simulators Give Better Policy Gradients?” H.J. Terry Suh, **Max Simchowitz**, Kaiqing Zhang, Russ Tedrake. *International Conference on Machine Learning (ICML)*, 2022. [ArXiv: 2202.00817](#). **Outstanding Paper Award**.

“Reward-Free RL is No Harder Than Reward-Aware RL in Linear Markov Decision Processes,” Andrew Wagenmaker, Yifang Chen, **Max Simchowitz**, Simon S. Du, Kevin Jamieson. *International Conference on Machine Learning (ICML)*, 2022. [ArXiv: 2201.11206](#). **Spotlight**.

“First-Order Regret in Reinforcement Learning with Linear Function Approximation: A Robust Estimation Approach,” Andrew Wagenmaker, Yifang Chen, **Max Simchowitz**, Simon S. Du, Kevin Jamieson. *International Conference on Machine Learning (ICML)*, 2022. [ArXiv: 2112.03432](#). **Oral Presentation**.

“Beyond No Regret: Instance-Dependent PAC Reinforcement Learning,” Andrew Wagenmaker, **Max Simchowitz**, Kevin Jamieson. *Conference on Learning Theory (COLT)*, 2022. [ArXiv: 2108.02717](#).

2021

“Bayesian decision-making under misspecified priors with applications to meta-learning,” **Max Simchowitz**, Christopher Tosh, Akshay Krishnamurthy, Daniel Hsu, Thodoris Lykouris, Miroslav Dudik, Robert E. Schapire. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.

ArXiv: 2103.00360. **Spotlight.**

“Stabilizing Dynamical Systems via Policy Gradient Methods,” Juan C. Perdomo, Jack Umenberger, **Max Simchowitz**. *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.

“Online Control of Unknown Time-Varying Dynamical Systems,” Edgar Minasyan, Paula Gradu, **Max Simchowitz**, Elad Hazan. *Conference on Neural Information Processing Systems (NeurIPS)*. Neurips, 2021. ArXiv: 2202.07890.

“On Stability of Nonlinear Receding Horizon Control: A Geometric Perspective,” Tyler Westbroek, **Max Simchowitz**, Michael I. Jordan, S. Shankar Sastry. *Conference on Decision and Control (CDC)*, 2021. ArXiv: 2103.15010.

“Task-Optimal Exploration in Linear Dynamical Systems,” Andrew Wagenmaker, **Max Simchowitz**, Kevin Jamieson. *International Conference on Machine Learning (ICML)*, 2021. ArXiv: 2102.05214

“Towards a Dimension-Free Understanding of Adaptive Linear Control,” Juan C. Perdomo, **Max Simchowitz**, Alekh Agarwal, Peter Bartlett. *Conference on Learning Theory (COLT)*, 2021. ArXiv: 2103.10620

“Corruption Robust Exploration in Episodic Reinforcement Learning,” Thodoris Lykouris, **Max Simchowitz**, Aleksandrs Slivkins, Wen Sun, (alphabetical order). *Conference on Learning Theory (COLT)*, 2021 (extended abstract). ArXiv: 1911.08689.

2020

“Making Non-Stochastic Control (Almost) as Easy as Stochastic,” **Max Simchowitz**. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. ArXiv: 2006.05910.

“Learning a Linear Quadratic Regular from Nonlinear Observations,” Zak Mhammedi, Dylan Foster, **Max Simchowitz**, Dipendra Misra, Akshay Krishnamurthy, Sasha Rakhlin, John Langford. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. ArXiv: 2010.03799.

“Constrained Episodic Reinforcement Learning in Concave-Convex and Knapsack Setting,” Kianté Brantley, Miroslav Dudik, Thodoris Lykouris, Sobhan Miryoosefi, **Max Simchowitz**, Aleksandrs Slivkins, Wen Sun. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020. ArXiv: 2006.05051.

“Naive Exploration is Optimal for Online LQR”, **Max Simchowitz**, Dylan Foster. *International Conference on Machine Learning (ICML)*, 2020. ArXiv: 2001.09576.

“Logarithmic Regret for Online Control with Adversarial Disturbances”, Dylan Foster, **Max Simchowitz**. *International Conference on Machine Learning (ICML)*, 2020. ArXiv: 2003.00189

“Reward-Free Exploration for Reinforcement Learning,” Chi Jin, Akshay Krishnamurthy, **Max Simchowitz**, Tiancheng Yu (alphabetical order). *International Conference on Machine Learning (ICML)*, 2020. ArXiv: 2002.02794.

“Improper Learning for Nonstochastic Control,” **Max Simchowitz**, Karan Singh, Elad Hazan. *Conference on Learning Theory (COLT)*, 2020. ArXiv: 2001.09254.

“The Gradient Complexity of Linear Regression,” Mark Braverman, Elad Hazan, **Max Simchowitz**, Blake Woodworth, (alphabetical order). *Conference on Learning Theory (COLT)*, 2020. ArXiv: 1911.02212.

“Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning,” Esther Rolf, **Max Simchowitz**, Sarah Dean, Lydia T. Liu, Daniel Björkegren,

Moritz Hardt, Joshua Blumenstock. *International Conference on Machine Learning (ICML)*, 2020. ArXiv: 2003.06740.

“A Successive-Elimination Approach to Adaptive Robotic Source Seeking,” Esther Rolf, David Fridovich-Keil, **Max Simchowitz**, Benjamin Recht, Claire Tomlin. *IEEE Transactions on Robotics*, 2020. ArXiv: 1809.10611

2019

“Non-Asymptotic Gap-Dependent Regret Bounds for Tabular MDPs,” **Max Simchowitz**, Kevin Jamieson. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019. ArXiv: 1905.03814.

“Learning Linear Dynamical Systems with Semi-Parametric Least Squares,” **Max Simchowitz**, Ross Boczar, Benjamin Recht. *Conference on Learning Theory (COLT)*, 2019. ArXiv: 1902.00768

2018

“The Implicit Fairness Criterion of Unconstrained Learning,” Lydia Liu, **Max Simchowitz**, Moritz Hardt. *International Conference on Machine Learning (ICML)*, 2019. ArXiv: 1808.10013.

“Delayed Impact of Fair Machine Learning,” Lydia T. Liu, Sarah Dean, Esther Rolf, **Max Simchowitz**, Moritz Hardt. *International Conference on Machine Learning (ICML)*, 2018. ArXiv: 1803.04383. **Best Paper Award.**

“Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification,” **Max Simchowitz**, Horia Mania, Stephen Tu, Benjamin Recht, Michael I. Jordan. *Conference on Learning Theory (COLT)*, 2018. ArXiv: 1802.08334

“Tight Query Complexity Lower Bounds for PCA via Finite Sample Deformed Wigner Law,” **Max Simchowitz**, Ahmed El Alaoui, Benjamin Recht. *Symposium on the Theory of Computing (STOC)*, 2018. ArXiv: 1804.01221.

“Approximate Ranking from Pairwise Comparisons,” Reinhard Heckel, **Max Simchowitz**, Kannan Ramchandran, Martin J. Wainwright. *Artificial Intelligence and Statistics Conference (AISTATS)*, 2018. ArXiv: 1801.01253

2017 and before

“The Simulator: Towards a Richer Understanding of Adaptive Sampling in the Moderate-Confidence Regime,” **Max Simchowitz**, Kevin Jamieson, and Benjamin Recht. *Conference on Learning Theory (COLT)*, 2017. ArXiv: 1702.05186

“Best-of-K Bandits,” **Max Simchowitz**, Kevin Jamieson, and Benjamin Recht. *Conference on Learning Theory (COLT)*, 2016. ArXiv: 1603.02752.

“Gradient Descent Converges to Minimizers,” Jason D. Lee, **Max Simchowitz**, Michael I. Jordan, and Benjamin Recht. *Conference on Learning Theory (COLT)*, 2016. ArXiv: 1602.04915.

“Low-rank Solutions of Linear Matrix Equations via Procrustes Flow,” Stephen Tu, Ross Boczar, **Max Simchowitz**, Mahdi Soltanolkotabi, Benjamin Recht. *International Conference on Machine Learning (ICML)*, 2016. ArXiv: 1507.03566.

Block, Daniel Pfrommer, **Max Simchowitz**, (alphabetical order). *Frontiers4LCD Workshop, International Conference on Machine Learning (ICML)*, 2023.

“Pathologies and Challenges of Using Differentiable Simulators in Policy Optimization for Contact-Rich Manipulation,” H.J. Terry Suh, **Max Simchowitz**, Kaiqing Zhang, Russ Tedrake. *ICRA 2022 Workshop: Reinforcement Learning for Contact-Rich Manipulation*, 2022.

“Learning to Extrapolate: A Transductive Approach,”. Aviv Netanyahu, Abhishek Gupta, **Max Simchowitz**, Kaiqing Zhang, Pulkit Agrawal. *NeurIPS Workshop on Distribution Shifts (Dist-Shift)*, 2022.

“Balancing Competing Objectives for Welfare-Aware Machine Learning with Imperfect Data,” Esther Rolf, **Max Simchowitz**, Sarah Dean, Lydia T. Liu, Daniel Björkegren, Moritz Hardt, Joshua Blumenstock. *The NeurIPS Joint Workshop on AI for Social Good*, 2019. **Best paper award**.

PREPRINTS

2026

“Workspace Models: Lightweight Robotic Memory via Saliency-Driven Supervision,” Nitish Dashora, Douglas Chen, Idan Shenfeld, John Marangola, Pulkit Agrawal, **Max Simchowitz**. Preprint, 2026. [ArXiv: 2609.20820](https://arxiv.org/abs/2609.20820).

“Adaptation of Generalist Robot Policies with Minimal Data,” Shreyas Kowshik, Sreyas Venkataraman, Leo Wang, Niharika Pant, **Max Simchowitz**, Aviral Kumar. [ArXiv: 2608.11363](https://arxiv.org/abs/2608.11363).

“Few-Step Boltzmann Generators via Scalable Likelihood Flow Maps,” RuiKang OuYang, Hanlin Yu, Xinyue Ai, Yutong He, Nicholas M. Boffi, Pradeep Ravikumar, José Miguel Hernández-Lobato, **Max Simchowitz**, Benjamin Kurt Miller, Omar Chehab. [ArXiv: 2606.29110](https://arxiv.org/abs/2606.29110).

“FADA: Few-Shot Domain Adaptation via Dynamics Alignment for Humanoid Control,” Angchen Xie, Nikhil Sobanbabu, Ishayu Shikhare, Alan Wang, **Max Simchowitz**, Guanya Shi. [ArXiv: 2606.28476](https://arxiv.org/abs/2606.28476).

“Double Preconditioning (DoPr): Optimization for Test-Time Performance, not Validation Loss,” Thomas T. Zhang, Alok Shah, Yifei Zhang, Vincent Zhang, Nikolai Matni, **Max Simchowitz**. [ArXiv: 2606.06418](https://arxiv.org/abs/2606.06418).

“Turning Video Models into Generalist Robot Policies,” Sizhe Lester Li, Evan Kim, Xingjian Bai, Tong Zhao, Tao Pang, **Max Simchowitz**, Vincent Sitzmann. [ArXiv: 2605.27817](https://arxiv.org/abs/2605.27817).

“Nano World Models: A Minimalist Implementation of Future Video Prediction,” Siqiao Huang, Partha Kaushik, Michael Chen, Hengkai Pan, Kaiwen Geng, Omar Chehab, Fernando Moreno-Pino, **Max Simchowitz**. [ArXiv: 2605.23993](https://arxiv.org/abs/2605.23993).

“MLS-Bench: A Holistic and Rigorous Assessment of AI Systems on Building Better AI,” Bohan Lyu, Yucheng Yang, Siqiao Huang, Jiaru Zhang, Qixin Xu, Xinghan Li, Xinyang Han, Yicheng Zhang, Huaqing Zhang, Runhan Huang, Kaicheng Yang, Zitao Chen, Wentao Guo, Junlin Yang, Xinyue Ai, Wenhao Chai, Yadi Cao, Ziran Yang, Kun Wang, Dapeng Jiang, Huan-ang Gao, Shange Tang, Chengshuai Shi, Simon S. Du, **Max Simchowitz**, Jiantao Jiao, Dawn Song, Chi Jin. [ArXiv: 2605.08678](https://arxiv.org/abs/2605.08678).

Additional preprints: <https://msimchowitz.github.io/publications.html>

TUTORIALS

“Online Nonstochastic Control,” *Conference on Decision and Control (CDC)*, December, 2025

“Foundations of Imitation Learning: From Language Modeling to Continuous Control,” *Conference on Neural Information Processing Systems (NeurIPS)*, December 2025.

INVITED
TALKS

“Do we need diffusion in robotics?” *Diffusion Generative Modeling: Progress and Next Steps*, *Simons Institute for the Theory of Computing*, August 2026.

“Diffusion: Do we need it in robot learning?” *Diffusion for Robot Learning Workshop*, *Robotics: Science and Systems*, July 2026.

“Generative Control, Action Chunking, and Moravec’s Paradox,” *Carnegie Mellon Robotics Seminar*, March 2026.

“Generative Control, Action Chunking, and Moravec’s Paradox,” *Stanford Robotics Seminar*, February 2026.

“A mathematical basis for Moravec’s paradox,” *Berkeley CLIMB Seminar*, February 2026.

“A mathematical basis for Moravec’s paradox,” *MIT Statistics Seminar*, February 2026.

“Why we chunk out actions,” *CoRL Workshop on Robotics Data*, September, 2025.

“Why AI is harder in the physical world, and what to do about it,” *Post-ICML Workshop @ Simon Fraser University*, July, 2025.

“Pitfalls of Imitation Learning when the Action Space is Continuous,” *RL Theory Seminar*, April, 2025.

“Generalization in Robotic Behavior Cloning: Pitfalls and Promises,” *Emerging Generalization Settings*, *Simons Institute for the Theory of Computing*, September, 2024.

“Provable Guarantees for Generative Behavior Cloning,” *Workshop on Foundations of Reinforcement Learning and Control: Connections and Perspectives*, *International Conference on Machine Learning (ICML)*, July, 2024.

TEACHING

“ML 10726: Machine Learning Theory: The Art of Theorem Proving and Problem Formulation,” *Carnegie Mellon University*, Spring 2025. **Instructor.**

Instructor. “ML 10701: Introduction to Machine Learning,” *Carnegie Mellon University*, Spring 2025.

SERVICE AND
LEADERSHIP

Conference Area Chair.

International Conference on Representation Learning (ICLR) - 2025, 2026

Conference on Learning Theory (COLT) - 2026.

Neural Information Processing Systems (NeurIPS) - 2026.

Conference on Robot Learning (CoRL) - 2026.

Workshop Organization.

Post-Training for Robotics Foundation Models, RSS, 2026. *ICML 2020 Workshop on the Theoretical Foundations of Reinforcement Learning*, ICML, 2020

Conference Reviewing. Neural Information Processing Systems (NeurIPS) (Top 400 Reviewer, 2019), Conference on Learning Theory (COLT), International Conference on Machine Learning (ICML), Symposium on the Theory of Computing (STOC), Artificial Intelligence and Statistics (AISTATS), Conference on Decision and Control (CDC).

Journal Reviewing.

Journal of Machine Learning Research (JMLR), *Editorial Board*.

Workshop Volunteering.

Reviewing for first-time authors, *NeurIPS 2019*.

Career Mentoring Volunteer, *Women in Machine Learning Workshop, ICML 2022*.

Mentoring

BAIR Mentor, *University of California, Berkeley*, 2017-2018.

CALMentor, *University of California, Berkeley*, 2020.

Graduate Application Assistance Program (GAAP) Mentor, *MIT EECS Department*, 2023.

Student Mentor, *Learning Theory Alliance, Fall Workshop*, 2023.

REFERENCES

1. Benjamin Recht

Professor, Department of Electrical Engineering and Computer Science, *UC Berkeley*

brecht@berkeley.edu

2. Michael I. Jordan

Pehong Chen Distinguished Professor, Department of Electrical Engineering and Computer Science & Department of Statistics, *UC Berkeley*.

jordan@cs.berkeley.edu

3. Russ Tedrake

Toyota Professor of Electric Engineering and Computer Science, Aero/Astro, & Mechanical Engineering, *Massachusetts Institute of Technology*.

Vice President, Robotics Research, *Toyota Research Institute*.

russt@mit.edu

4. Elad Hazan

Professor of Computer Science, *Princeton University*.

Director and Co-Founder, *Google AI Princeton*.

mitraforhazan@gmail.com

5. Sasha Rakhlin

Professor of Brain & Cognitive Sciences, *MIT*.

Statistics and Data Science Center, IDSS.

rakhlin@mit.edu