

Adam Block

Updated September 15, 2025

Email: adam.block@columbia.edu **Address:** 500 W 120th St, New York, NY 10027
Office: CSB 454 **Homepage:** <https://abblock.github.io/>

Research interests Reinforcement Learning, Language Models, Sequential Decision Making, On-line Learning

Employment **Assistant Professor at Columbia University** New York, NY
07/2025 – Present

Postdoc at Microsoft Research NYC New York, NY
07/2024 – 07/2025

Adjunct Professor at Columbia University New York, NY
01/2025 – 07/2025

Education **MIT** Cambridge, MA
PhD in Mathematics and Statistics 09/2019 Year – 05/2024
Advisor: Alexander Rakhlin

Columbia University New York, New York
BA in Mathematics (Summa cum Laude) 09/2015 Year – 05/2019

Honors Finalist for Jane Street Graduate Research Fellowship 2023
NeurIPS 2022 Scholar Award 2022
COLT 2022 Deepmind Student Travel Grant 2022
SIGIR 2022 Student Travel Award 2022
NSF Graduate Research Fellowship 2019-Present
Phi Beta Kappa (Early designation) 2019
I.I. Rabi Scholar (Columbia University) 2015-2019

Publications **GaussMark: A Practical Approach for Structural Watermarking of Language Models**
Adam Block, Ayush Sekhari, and Alexander Rakhlin
ICML 2025

Is Best-of-N the Best of Them? Coverage, Scaling, and Optimality in Inference-Time Alignment
Audrey Huang, Adam Block, Qinghua Liu, Nan Jiang, Dylan J. Foster, and Akshay Krishnamurthy

ICML 2025

A Theory of Learning with Autoregressive Chain of Thought

Nirmit Joshi, Gal Vardi, Adam Block, Surbhi Goel, Zhiyuan Li, Theodor Misiakiewicz, Nathan Srebro

COLT 2025

Computational-Statistical Tradeoffs at the Next-Token Prediction Barrier: Autoregressive and Imitation Learning under Misspecification

Dhruv Rohatgi, Adam Block, Audrey Huang, Akshay Krishnamurthy, and Dylan J. Foster

COLT 2025

Self-Improvement in Language Models: The Sharpening Mechanism

Audrey Huang, Adam Block, Dylan J. Foster, Dhruv Rohatgi, Cyril Zhang, Max Simchowitz, Jordan T. Ash, and Akshay Krishnamurthy

ICLR 2025 (Oral)

Is Behavior Cloning All You Need? Understanding Horizon in Imitation Learning

Dylan J. Foster, Adam Block, and Dipendra Misra

NeurIPS 2024 (Spotlight)

Oracle-Efficient Differentially Private Learning with Public Data

Adam Block, Mark Bun, Rathin Desai, Abhishek Shetty, and Zhiwei Steven Wu

NeurIPS 2024

On the Performance of Empirical Risk Minimization with Smoothed Data

Adam Block, Alexander Rakhlin, and Abhishek Shetty

COLT 2024

Butterfly Effects of SGD Noise: Error Amplification in Behavior Cloning and Autoregression

Adam Block, Dylan J. Foster, Akshay Krishnamurthy, Max Simchowitz, and Cyril Zhang

ICLR 2024

Smoothed Online Learning for Prediction in Piecewise Affine Systems

Adam Block, Max Simchowitz, and Russ Tedrake

NeurIPS 2023 (Spotlight)

On the Imitation of Non-Markovian Demonstrations: From Low-Level Stability to High Level Planning

Adam Block, Ali Jadbabaie, Daniel Pfrommer, Max Simchowitz, and Russ Tedrake

ICML 2023, Workshop on Optimal Transport in Learning, Control, and Dynamical Systems

NeurIPS 2023

Efficient Model-Free Exploration in Low-Rank MDPs

Zak Mhammedi, Adam Block, Dylan Foster, and Alexander Rakhlin

NeurIPS 2023

Oracle-Efficient Smoothed Online Learning for Piecewise Continuous Decision Making

Adam Block, Max Simchowitz, and Alexander Rakhlin

COLT 2023

The Sample Complexity of Approximate Rejection Sampling With Applications to Smoothed Online Learning

Adam Block and Yury Polyanskiy

COLT 2023

Efficient and Near-Optimal Smoothed Online Learning for Generalized Linear Functions

Adam Block and Max Simchowitz

NeurIPS 2022

Intrinsic Dimension Estimation using Wasserstein Distance

Adam Block, Zeyu Jia, Yury Polyanskiy, and Alexander Rakhlin.

Journal of Machine Learning Research (Accepted 2022)

Smoothed Online Learning is as Easy as Statistical Learning

Adam Block, Yuval Dagan, Noah Golowich, and Alexander Rakhlin

COLT 2022

Counterfactual Learning To Rank for Utility-Maximizing Query Auto-completion

Adam Block, Rahul Kidambi, Thorsten Joachims, Daniel N. Hill, and Inderjit S. Dhillon

SIGIR 2022

Majorizing Measures, Sequential Complexities, and Online Learning

Adam Block, Yuval Dagan, and Alexander Rakhlin.

COLT 2021

Probability Theory and Related Fields (in revision)

Preprints

EMA Without the Lag: Bias-Corrected Iterate Averaging Schemes

Adam Block and Cyril Zhang
arxiv preprint arXiv:2508.00180

Small Loss Bounds for Online Learning Separated Function Classes: A Gaussian Process Perspective

Adam Block and Abhishek Shetty
arxiv preprint arXiv:2502.10292

Rate of Convergence of the Smoothed empirical Wasserstein Distance

Adam Block, Zeyu Jia, Yury Polyanskiy, and Alexander Rakhlin
arxiv preprint arXiv:2205.02128

Fast mixing of multi-scale langevin dynamics under the manifold hypothesis

Adam Block, Youssef Mroueh, Jerret Ross, and Alexander Rakhlin.
arXiv preprint arXiv:2006.11166

Generative modeling with denoising auto-encoders and Langevin sampling

Adam Block, Youssef Mroueh, and Alexander Rakhlin.
arXiv preprint arXiv:2002.00107, 2020.

Research experience (Math)

Cohomology of Schur Functors Indexed by Hook Partitions

Mentor: Daniel Litt (Columbia University) 09/2015 – 05/2019
Bounding the Frobenius Amplitude of vector bundles on schemes over positive characteristic. Used SageMath and Python for computations.

Persistent Homology of Point Cloud Data

Mentor: Benjamin Antieau (University of Illinois at Chicago) 09/14 – 05/15
Explored the algebro-topological invariants of random point-cloud data and began working towards creating an efficient package in SageMath for calculating said invariants.

Work experience

Research Intern at Microsoft Research NYC

Mentors: Cyril Zhang, Akshay Krishnamurthy, Dylan Foster
Worked on understanding the inductive biases of different training approaches to deep networks for behavior cloning and NLP.

Applied Science Intern at Amazon Science

Mentors: Rahul Kidambi (Amazon), Thorsten Joachims (Cornell) 05/21 – 10/21

Talks

Improved Amazon's auto-complete feature with a new method relying on counterfactual estimation in ranking as a member of the MIDAS team.

GaussMark: A Practical Approach for Watermarking Language Models

Amazon AWS Responsible AI Seminar 05/2025

GaussMark: A Practical Approach for Watermarking Language Models

Google Research Learning Seminar 05/2025

Scaling Inference Time Compute: From Self-Improvement to Pessimism

Universitat Pompeu Fabra 04/2025

Scaling Inference Time Compute: From Self-Improvement to Pessimism

Flatiron Institute 03/2025

Self-Improvement in Language Models: The Sharpening Mechanism

Google Research Learning Theory Seminar 12/2024

Empirical Risk Minimization with Dependent Data

10/2024

INFORMS Invited Session *The Power of Optimization in Sequential Decision Making*

Understanding the Role of Horizon in Imitation Learning

09/2024

Berkeley Theory Seminar

Smoothed Online Learning: Theory and Applications

11/2023

Columbia University Theory Student Seminar

Smoothed Online Learning: Theory and Applications

10/2023

Google Research Learning Theory Seminar

Tackling Combinatorial Distribution Shift: A Matrix Completion Perspective

07/2023

COLT 2023 (on behalf of Max Simchowitz, Kaiqing Zhang, and Abhishek Gupta)

Entropic characterization of optimal rates for learning Gaussian mixtures

07/2023

COLT 2023 (on behalf of Zeyu Jia, Yury Polyanskiy, and Yihong Wu)

Minimax optimal testing by classification

07/2023

COLT 2023 (on behalf of Patrik Gerber, Yanjun Han, and Yury Polyanskiy)

Imitation of Non-Markovian Demonstrations 06/2023
Microsoft Research Reinforcement Learning Seminar

Intrinsic Dimension Estimation with Wasserstein Distances 11/2022
McMaster University Applied Probability Seminar

Intrinsic Dimension Estimation with Wasserstein Distances 02/2022
MIT LIDS & Stats Tea Talk

Generative Modeling with Langevin Dynamics 12/2020
Methods of Deep Learning Workshop

Sampling with Langevin Dynamics 06/2020
MIT-IBM Machine Learning Reading Group (Host: Youssef Mroueh)

Generative Modeling with De-Noising Auto-Encoders and Langevin Dynamics 04/2020
Carnegie-Mellon ML Reading Group (Host: Andrej Risteski)

Organizer

Service and Organization

Foundations of Pre-training Workshop (COLT 2025) 2025

Reviewer

COLT 2020, ALT 2020, NeurIPS 2020, COLT 2021, *Annals of Statistics*, *Bernoulli*, COLT 2022, MIT SDSCON 2022, IEEE Transactions on Information Theory, *Journal of Machine Learning Research*, COLT 2023 PC, ALT 2023 PC, COLT 2024 PC, NeurIPS 2024 Reviewer, Foundations of RL and Control (Workshop at ICML 2024), *Journal of Management Sciences*, ICLR 2025 Reviewer, ICML 2025 Reviewer, COLT 2025 Senior PC

Programming

Skills

Proficient in: PyTorch, Python, Julia, \LaTeX , PySpark
Familiar with: TensorFlow