

# JASON ALTSCHULER

alts@upenn.edu | <https://jasonaltschuler.github.io/>

## ACADEMIC POSITIONS

---

<b>UPenn</b> Assistant Professor, Statistics and Data Science (primary) Assistant Professor, Electrical and Systems Engineering (secondary) Assistant Professor, Computer and Information Science (secondary) Affiliated Faculty, Applied Mathematics and Computational Science Affiliated Faculty, PRiML Center for Machine Learning Affiliated Faculty, Warren Center for Networks and Data Science	<i>2023-current</i>
<b>NYU</b> Faculty Fellow	<i>2022-2023</i>

## EDUCATION

---

<b>Massachusetts Institute of Technology (PhD)</b> Electrical Engineering and Computer Science Thesis committee: Pablo Parrilo (advisor), Philippe Rigollet, Guy Bresler Thesis: Transport and beyond: efficient optimization over probability distributions. <i>Sprowls Award for Best PhD Thesis in Artificial Intelligence &amp; Decision Making (MIT)</i>	<i>2018-2022</i>
<b>Massachusetts Institute of Technology (MS)</b> Electrical Engineering and Computer Science Advisor: Pablo Parrilo Thesis: Greed, hedging, and acceleration in convex optimization. <i>GPA: 5.0/5.0</i>	<i>2016-2018</i>
<b>Princeton University (BSE)</b> Major: Computer Science Minors: Applied and Computational Mathematics; Statistics and Machine Learning CS thesis: Online learning with limited decision changes. Advisor: Elad Hazan. Math thesis: Probabilistic linear Boolean classification. Advisor: Emmanuel Abbe. <i>GPA 3.99/4.0, Major GPA: 4.0/4.0, highest honors</i>	<i>2012-2016</i>
<b>Oxford University (undergrad semester abroad)</b> Math and Computer Science <i>GPA 4.0/4.0 (converted to US)</i>	<i>Spring 2015</i>

## SELECTED ACADEMIC AWARDS

---

SIAM Early Career Prize for Optimization (awarded once every 3 years)	<i>2026</i>
AFOSR Young Investigator Program (YIP) Award in Mathematical Optimization	<i>2026-2029</i>
ICS Prize for best paper at the interface of Operations Research and Computer Science	<i>2025</i>
Sloan Research Fellowship in Mathematics	<i>2025</i>
Undergraduate Teaching Excellence Award x2	<i>2024, 2025</i>
Seed Grant Award from Apple	<i>2024</i>
Selected for National Academy of Sciences' Kavli Frontiers of Science symposium (declined due to scheduling conflict)	<i>2024</i>
A.W. Tucker Prize Finalist (Mathematical Optimization Society)	<i>2021-2024</i>
George M. Sprowls Award for Best MIT PhD Thesis in Artificial Intelligence & Decision Making	<i>2022-2023</i>
NYU Data Science Faculty Fellowship	<i>2022-2024</i>

IBM Herman Goldstine Fellowship (declined)	2022-2024
TwoSigma Ph.D. Fellowship	2020-2022
Best talk, MIT LIDS Student Conference	2021
NeurIPS Top Reviewers Award	2019
ICML Top Reviewers Award	2019
National Science Foundation Graduate Fellowship	2016-2021
ICML Travel Award	2016
James Hayes-Edgar Palmer Prize in Engineering (Princeton)	2016
Highest honors, summa cum laude (Princeton)	2016
Sigma Xi Honor Society (Princeton)	2016
Accenture Prize in Computer Science (Princeton)	2015
Phi Beta Kappa Honor Society, early induction (Princeton)	2015
Tau Beta Pi Engineering Honor Society, early induction (Princeton)	2014
Shapiro Award for Academic Excellence, all possible years (Princeton)	2013, 2014

## PUBLICATIONS

---

### Preprints

- P1 Henry Shugart, Shuyi Wang, Jason Altschuler. Negative Momentum in Convex-Concave Optimization. *Preprint at arXiv:2604.17145, 2026.*
- P2 Henry Shugart, Jason Altschuler. Min-Max Optimization Is Strictly Easier Than Variational Inequalities. *Preprint at arXiv:2511.03052, 2025.*
- P3 Henry Shugart, Jason Altschuler. Negative Stepsizes Make Gradient-Descent-Ascent Converge. *Preprint at arXiv:2505.01423, 2025.*
- P4 Matthew Zhang, Jason Altschuler, Sinho Chewi. Algorithmic warm starts for Hamiltonian Monte Carlo. *Preprint at arXiv:2603.22741, 2026.*
- P5 Jason Altschuler, Sinho Chewi, Matthew Zhang. Shifted Composition IV: Toward Ballistic Acceleration for Log-Concave Sampling. *Symposium on Theory of Computing (STOC), to appear in 2026.*

### Journal articles

- J1 Jinho Bok, Jason Altschuler. Optimized Methods for Composite Optimization: a Reduction Perspective. *Mathematical Programming, to appear in 2026+.* (Preliminary conference version in C1.)
- J2 Xufeng Cai, Jason Altschuler, Jelena Diakonikolas. Near-Linear Runtime for a Classical Matrix Preconditioning Algorithm. *Journal of the ACM, 2025.*
- J3 Jason Altschuler, Pablo Parrilo. Acceleration by Random Stepsizes: Hedging, Equalization, and the Arcsine Step Size Schedule. *Foundations of Computational Mathematics, to appear in 2026+.*
- J4 Jason Altschuler, Sinho Chewi. Shifted Composition III: Local Error Framework for KL Divergence. *Foundations of Computational Mathematics, to appear in 2026+.*
- J5 Jason Altschuler, Sinho Chewi. Shifted Composition II: Shift Harnack Inequalities and Curvature Upper Bounds. *IEEE Transactions on Information Theory, 2025.*
- J6 Jason Altschuler, Sinho Chewi. Shifted Composition I: Harnack and Reverse Transport Inequalities. *IEEE Transactions on Information Theory, 2025.*
- J7 Jason Altschuler, Kunal Talwar. Resolving the mixing time of the Langevin Algorithm to its stationary distribution for log-concave sampling. *SIAM Journal on Mathematics of Data Science, 2025.* (Preliminary conference version in C4.)

- J8 Jason Altschuler, Pablo Parrilo. Acceleration by StepSize Hedging: Silver StepSize Schedule for Smooth Convex Optimization. *Mathematical Programming*, 2024. **INFORMS Computing Society Prize 1/2**
- J9 Jason Altschuler, Pablo Parrilo. Acceleration by StepSize Hedging: Multi-Step Descent and the Silver StepSize Schedule. *Journal of the ACM*, 2024. **INFORMS Computing Society Prize 2/2**
- J10 Jason Altschuler, Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Journal of the ACM*, 2024. (Preliminary conference version in C3.)
- J11 Jason Altschuler, Jinho Bok, Kunal Talwar. On the privacy of Noisy Stochastic Gradient Descent for convex optimization. *SIAM Journal on Computing*, 2024. (Preliminary conference version in C5.)
- J12 Jason Altschuler, Pablo Parrilo. Kernel approximation on algebraic varieties. *SIAM Journal on Applied Algebra and Geometry*, 2023.
- J13 Jason Altschuler, Enric Boix-Adserà. Polynomial-time algorithms for Multimarginal Optimal Transport problems with structure. *Mathematical Programming*, 2023.
- J14 Jason Altschuler, Jonathan Niles-Weed, Austin Stromme. Asymptotics for semi-discrete entropic optimal transport. *SIAM Journal on Mathematical Analysis*, 2022.
- J15 Jason Altschuler, Pablo Parrilo. Near-linear convergence of the Random Osborne algorithm for Matrix Balancing. *Mathematical Programming*, 2023.
- J16 Jason Altschuler, Pablo Parrilo. Approximating Min-Mean-Cycle for low-diameter graphs in near-optimal time and memory. *SIAM Journal on Optimization*, 2022.
- J17 Jason Altschuler, Enric Boix-Adserà. Wasserstein barycenters are NP-hard to compute. *SIAM Journal on Mathematics of Data Science*, 2022.
- J18 Jason Altschuler, Enric Boix-Adserà. Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research*, 2021.
- J19 Jason Altschuler, Enric Boix-Adserà. Hardness results for Multimarginal Optimal Transport problems. *Discrete Optimization*, 2021.
- J20 Jason Altschuler, Pablo Parrilo. Lyapunov exponent of rank one matrices: ergodic formula and inapproximability of the optimal distribution. *SIAM Journal on Control and Optimization*, 2020. (Preliminary conference version in C8.)
- J21 Jason Altschuler, Kunal Talwar. Online learning over a finite action set with limited switching. *Mathematics of Operations Research*, 2021. (Preliminary conference version in C9.)
- J22 Jason Altschuler, Victor-Emmanuel Brunel, Alan Malek. Best arm identification for contaminated bandits. *Journal of Machine Learning Research*, 2019.
- J23 Jason Altschuler, Elizabeth Yang. Inclusion of forbidden minors in random representable matroids. *Discrete Mathematics*, 2017.
- J24 Benjamin Pavie, Satwik Rajaram, Austin Ouyang, Jason Altschuler, Robert Steininger, Lani Wu, Steven Altschuler. Rapid analysis and exploration of fluorescence microscopy images. *Journal of Visualized Experiments*, 2014.

### Conference proceedings

- C1 Jinho Bok, Jason Altschuler. Accelerating proximal gradient descent via silver stepsizes. *Conference on Learning Theory (COLT)*, 2025.
- C2 Jinho Bok, Weijie Su, Jason Altschuler. Shifted interpolation for differential privacy. *International Conference on Machine Learning (ICML)*, 2024.
- C3 Jason Altschuler, Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Foundations of Computer Science (FOCS)*, 2023.
- C4 Jason Altschuler, Kunal Talwar. Resolving the mixing time of the Langevin Algorithm to its stationary distribution for log-concave sampling. *Conference on Learning Theory (COLT)*, 2023.

- C5 Jason Altschuler, Kunal Talwar. Privacy of Noisy Stochastic Gradient Descent: more iterations without more privacy loss. *Neural Information Processing Systems Conference (NeurIPS)*, 2022. **selected for oral presentation**
- C6 Jason Altschuler, Sinho Chewi, Patrik Gerber, Austin Stromme. Averaging on the Bures-Wasserstein manifold: dimension-free convergence of gradient descent. *Neural Information Processing Systems Conference (NeurIPS)*, 2021. **selected for spotlight presentation**
- C7 Jason Altschuler, Francis Bach, Alessandro Rudi, Jonathan Weed. Massively scalable Sinkhorn distances via the Nyström method. *Neural Information Processing Systems Conference (NeurIPS)*, 2019.
- C8 Jason Altschuler, Pablo Parrilo. Lyapunov exponent of rank one matrices: ergodic formula and inapproximability of the optimal distribution. *Conference on Decision and Control (CDC)*, 2019.
- C9 Jason Altschuler, Kunal Talwar. Online learning over a finite action set with limited switching. *Conference on Learning Theory (COLT)*, 2018.
- C10 Jason Altschuler, Jonathan Weed, Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via Sinkhorn iteration. *Neural Information Processing Systems Conference (NeurIPS)*, 2017. **selected for spotlight presentation**
- C11 Jason Altschuler, Aditya Bhaskara, Gang Fu, Vahab Mirrokni, Afshin Rostamizadeh, Morteza Zadimoghaddam. Greedy column subset selection: new bounds and distributed algorithms. *International Conference on Machine Learning (ICML)*, 2016.

### Expository writing

- E1 Jason Altschuler. Flows, scaling, and entropy revisited: a unified perspective via optimizing joint distributions. *SIAM Group on Optimization's Views and News, invited article*, 2022.

## TEACHING

---

### UPenn

- STAT 9912: Acceleration and Hedging in Optimization (**New course**; grad course) *Fall 2025*
- STAT 4300: Probability (**Undergraduate Teaching Excellence Award**) *Spring 2025*
- STAT 9910: Stability in Optimization and Statistics (**New course**; grad course) *Fall 2024*
- STAT 4300: Probability (**Undergraduate Teaching Excellence Award**) *Spring 2024*

### NYU

- DS-GA 1014: Optimization and Computational Linear Algebra (grad course) *Fall 2022*

### MIT

- 6.255/15.093/IDS.200: Optimization Methods (grad course), TA *Fall 2019*

### Princeton

- COS 511: Theoretical Machine Learning (grad course), TA *Spring 2016*
- MAT 340: Applied Algebra, TA *Fall 2015*
- MAT 216: Accelerated Real Analysis, TA *Fall 2014*
- MAT 217: Honors Linear Algebra, TA *Spring 2014*
- MAT 215: Honors Real Analysis, TA *Fall 2013*

Chess.com — chess coach. Over 40,000 hits on coaching profile. *2007 - 2014*

## ADVISING AND MENTORING

---

### PhD students

Jinho Bok (UPenn, Statistics and Data Science) *2023-current*

Henry Shugart (UPenn, Statistics and Data Science) *2023-current*

### Postdocs

Chris Criscitiello (UPenn, Statistics and Data Science departmental postdoc) *2025-current*

### Undergraduate students

Nikhil Kumar (UPenn, EE & Economics → next: PhD at Princeton) *2025*

Shuyi Wang (UPenn, Mathematics → next: PhD at Yale) *2024*

Vincent Huang (MIT, Mathematics) *2020-2021*

## DOCTORAL / QUALIFYING COMMITTEES

---

Kyurae Kim (UPenn, Computer Science), committee chair, advised by Jacob Gardner *2025*

Salma Tarmoun (UPenn, Applied Mathematics), advised by Rene Vidal *2025*

Zexuan Wang (UPenn, Applied Mathematics), advised by Li Shen *2025*

Thomas Zhang (UPenn, Electrical Engineering), advised by Nikolai Matni *2025*

Orlando Romero (UPenn, Electrical Engineering), advised by George Pappas *2024*

Patrick Chao (UPenn, Statistics and Data Science), advised by Edgar Dobriban *2024*

## SERVICE AND LEADERSHIP

---

### Organizing seminars and workshops

Initiator & organizer, **UPenn FOLDS Seminar** *2025+*

- University-wide seminar series on the **Foundations of Optimization, Learning, and Data Science**
- Primary seminar on foundational research for UPenn's new data science center
- Evolved from the UPenn Optimization Seminar that I created in 2023
- Website: <https://jasonaltschuler.github.io/folds-seminar>

Initiator & organizer, **UPenn Optimization Seminar** *2023+*

- University-wide seminar series to unite siloed communities across the UPenn campus studying optimization
- 50+ seminars organized
- Past talks: <https://jasonaltschuler.github.io/opt-seminar-fall-2023>  
<https://jasonaltschuler.github.io/opt-seminar-spring-2024>  
<https://jasonaltschuler.github.io/opt-seminar-fall-2024>  
<https://jasonaltschuler.github.io/opt-seminar-spring-2025>

Organizer, ICCOPT session on “Optimization for data science” *2025*

Co-organizer, NYU Math and Data Science seminar *2022-2023*

Co-organizer, NeurIPS workshop on “Optimal Transport and Machine Learning” *2021*

Co-organizer, SIAM Optimization mini-symposium on “Optimal Transport: Theory and Algorithms” *2021*

### Internal

Co-chair, postdoc recruitment committee for UPenn Statistics and Data Science department *2024-2025*

Committee member, Brown lectures and student workshop *2025*

Committee member, PhD admissions for UPenn Statistics and Data Science department *2022-2024*

Mentor underrepresented minorities in MIT's Graduate Application Assistance Program (GAAP)	2020-2022
Reader for MIT EECS PhD applications	2019-2020

### External

Mathematical Optimization Society, Tucker Prize Committee	2027
Institute of Mathematical Statistics (IMS), Committee on Awards	2024-2027
Conference on Learning Theory (COLT), Senior Program Committee	2025

### Reviewing

Foundations of Computational Mathematics (FOCM), Mathematical Programming (MAPR), SIAM Journal on Optimization (SIOPT), SIAM Journal on Mathematics of Data Science (SIMODS), Journal of Machine Learning Research (JMLR), Annals of Statistics (AoS), Annals of Probability (AoP), Annals of Applied Probability (AAP), Probability Theory and Related Fields (PTRF), Operations Research (OR), INFORMS Journal on Optimization (IJO), Transactions on Automatic Control (TACON), Information and Inference, Conference on Decision and Control (CDC), Neural Information Processing Systems (NeurIPS), International Conference on Machine Learning (ICML), Conference on Learning Theory (COLT), Symposium on Discrete Algorithms (SODA), Symposium on Theory of Computing (STOC), Symposium on Foundations of Computer Science (FOCS), American Mathematical Society book chapter, etc.

### SELECTED RECENT TALKS

---

Swiss Optimization Symposium ( <i>keynote talk</i> )	August 2026
FOCM, Workshop on Continuous Optimization	July 2026
SIAM Conference on Optimization, Lecture for Early Career Award ( <i>award talk</i> )	June 2026
SIAM Conference on Optimization, Session on analysis of iterative algorithms	June 2026
Oberwolfach, Workshop on Optimization for Computer-Assisted Proofs	May 2026
École Polytechnique Seminar	May 2026
ELLIT, Symposium on Optimization for Learning	May 2026
John Hopkins, Data Science Seminar	April 2026
UPenn, Econometrics Seminar	April 2026
MIT, EECS Seminar	April 2026
Princeton, ORFE Colloquium	February 2026
TwoSigma, Seminar	December 2025
CMU, Tepper Operations Research Seminar	November 2025
Pitt, Workshop on Mathematical Analysis and Machine Learning	November 2025
UPenn, Applied Math and Computational Sciences Seminar	November 2025
UT Austin, Applied and Computational Mathematics Seminar	November 2025
Rice, Computational Mathematics and Operations Research Colloquium	October 2025
INFORMS, Special session for ICS Prize ( <i>award talk</i> )	October 2025
Harvard CMSA, Workshop on Mathematical Foundations of AI	October 2025
LMU Munich, MIP Seminar	October 2025
Georgia Tech, Algorithms and Randomness Center (ARC) Colloquium	September 2025
University of Florida, Statistics Seminar	September 2025
International Conference on Continuous Optimization	July 2025

Portuguese American Optimization Workshop ( <i>invited lecture</i> )	June 2025
Stanford, ISL Colloquium	May 2025
Institute for Pure & Applied Mathematics, Workshop on Optimal Transport	May 2025
UPenn, Optimization Seminar	April 2025
Apple, Workshop on Privacy Preserving ML	March 2025
Les Houches, Workshop on Physics of AI Algorithms	January 2025
NeurIPS, OPT + ML workshop ( <i>plenary talk</i> )	December 2024
UW Madison, SILO seminar	November 2024
Rutgers, MSIS seminar	November 2024
Columbia, Applied Probability Seminar	November 2024
MIT, Sloan OR/STAT seminar	September 2024
Harvard, Statistics Colloquium	September 2024
Oaxaca BIRS, Workshop on Computational Harmonic Analysis in Data Science and ML	September 2024
Princeton ML Theory Summer School ( <i>principal speaker</i> )	August 2024
ISMP, Special session for Tucker Prize ( <i>award talk</i> )	July 2024
ISMP, Session on semidefinite programming	July 2024
AMS Sectional Meeting, Workshop on Recent Advances in Optimal Transport	April 2024
CISS, Special session on ML Theory	March 2024
Apple, Workshop on Privacy Preserving ML	March 2024
UPenn, CS Theory Seminar	February 2024
Singapore Institute for Mathematical Sciences, Mathematics of Data Program	January 2024
MIT, ORC Seminar	December 2023
Berkeley, BLISS Seminar	December 2023
Princeton, Machine Learning Theory Seminar	November 2023
UPenn, Electrical Engineering colloquium	October 2023
Kantorovich Initiative Seminar	October 2023
Columbia, Workshop on Robust Statistics & Privacy	October 2023
Duke, Decision Sciences seminar	October 2023
MIT, LIDS seminar	September 2023
INSEAD, Workshop on Information and Learning	July 2023
FOCM, Foundations of Data Science and Machine Learning workshop	June 2023
SIAM Conference on Optimization, Algebraic and Geometric Optimization session	May 2023
Princeton CCR colloquium	May 2023
Columbia, Statistics seminar	April 2023
MIT, LIDS	April 2023
Harvard, differential privacy seminar	April 2023
École Normale Supérieure, interteam seminar	January 2023
NYU, Data Science lunch seminar	October 2022
Yale, Statistics and Data Science seminar	September 2022

NYU, Mathematics and Data Science seminar	<i>September 2022</i>
Microsoft Research, ML Foundations seminar	<i>July 2022</i>
ICCOPT, Optimization for ML session	<i>July 2022</i>
Singapore UTD, Engineering Systems and Design seminar	<i>May 2022</i>
Berkeley, EECS seminar	<i>February 2022</i>
UWaterloo, Computer Science seminar	<i>February 2022</i>
CMU, Operations Research seminar	<i>February 2022</i>
UPenn, Statistics & Data Science seminar	<i>February 2022</i>
Princeton, ORFE colloquium	<i>February 2022</i>
Stanford, OIT seminar	<i>February 2022</i>
TTIC, Theory seminar	<i>January 2022</i>
Cornell, ORIE special seminar	<i>January 2022</i>
UWaterloo, Combinatorics and Optimization colloquium	<i>January 2022</i>
Caltech, Frontiers in Computing and Mathematical Sciences seminar	<i>January 2022</i>
INFORMS Conference, Optimal Transport session	<i>October 2021</i>
Caltech, Computational Mathematics special seminar	<i>October 2021</i>
NYU Courant, Numerical Analysis and Scientific Computing seminar	<i>October 2021</i>
Simons Institute, graduate student seminar	<i>September 2021</i>
SIAM Conference on Optimization, Optimal Transport session	<i>July 2021</i>
MIT, LIDS & Stats Tea	<i>April 2021</i>
INRIA + CNRS + Université Paris-Dauphine, joint MokaMeeting seminar	<i>March 2021</i>
MIT, LIDS Student Conference ( <i>best talk award</i> )	<i>January 2021</i>
MIT, LIDS & Stats Tea	<i>October 2020</i>
MIT, LIDS & Stats Tea	<i>April 2020</i>
MIT, Applied Mathematics graduate student seminar	<i>March 2020</i>
Harvard CMSA, workshop on noncommutative analysis and computational complexity	<i>October 2019</i>
Google Brain, invited seminar	<i>August 2019</i>

## INDUSTRY EXPERIENCE

---

Apple Research — Research Intern	<i>Summer 2021</i>
DE Shaw — Quant Intern	<i>Summer 2016</i>
Google Research — Research Intern	<i>Summer 2015</i>
Google — Software Engineering Intern	<i>Summer 2014</i>
Tower Research Capital — Quant Intern	<i>Summer 2013</i>

## OTHER ACTIVITIES

---

### Chess

· US National Scholastic Champion	<i>4 times</i>
· “International Master” norm at Blau Escacs Open (Spain)	<i>2015</i>
· “National Master” title at age 15	<i>2010</i>
· “Expert” title at age 12	<i>2007</i>

- Ranked top 5 in world for speed-chess (bullet and blitz) on chess.com

*2013*

- Won 2014 US Amateur Team Championship with Princeton team  
(featured in New York Times and US Chess news articles)

*2014*

### **Swimming**

- Varsity swim team at Oxford University

*Spring 2015*