

JASON ALTSCHULER

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

ACADEMIC POSITIONS

UPenn	<i>2023-current</i>
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	
NYU	<i>2022-2023</i>
Faculty Fellow	

EDUCATION

Massachusetts Institute of Technology (PhD)	<i>2018-2022</i>
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 & Decision Making (MIT)</i>	
Massachusetts Institute of Technology (MS)	<i>2016-2018</i>
Electrical Engineering and Computer Science	
Advisor: Pablo Parrilo	
Thesis: Greed, hedging, and acceleration in convex optimization.	
<i>GPA: 5.0/5.0</i>	
Princeton University (BSE)	<i>2012-2016</i>
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>	
Oxford University (undergrad semester abroad)	<i>Spring 2015</i>
Math and Computer Science	
<i>GPA 4.0/4.0 (converted to US)</i>	

SELECTED ACADEMIC AWARDS

Wharton Teaching Excellence Award for undergraduate teaching	<i>2024</i>
Seed Grant Award from Apple	<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)	<i>2022-2024</i>
TwoSigma Ph.D. Fellowship	<i>2020-2022</i>
Best talk, MIT LIDS Student Conference	<i>2021</i>
NeurIPS Top Reviewers Award	<i>2019</i>
ICML Top Reviewers Award	<i>2019</i>
National Science Foundation Graduate Fellowship	<i>2016-2021</i>

Columbia University Presidential Fellowship (declined)	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 Jason Altschuler, Sinho Chewi. Shifted Composition II: Shift Harnack Inequalities and Curvature Upper Bounds. *Preprint at arXiv:2401.00071, 2024.*
- P2 Jason Altschuler, Pablo Parrilo. Acceleration by StepSize Hedging I: Multi-Step Descent and the Silver StepSize Schedule. *Preprint at arXiv:2309.07879, 2023.*
- P3 Jason Altschuler, Kunal Talwar. Resolving the mixing time of the Langevin Algorithm to its stationary distribution for log-concave sampling. *Preprint, 2022.* (Preliminary conference version in C3.)

Journal articles

- J1 Jason Altschuler, Pablo Parrilo. Acceleration by StepSize Hedging II: Silver StepSize Schedule for Smooth Convex Optimization. *Mathematical Programming, 2024, to appear.*
- J2 Jason Altschuler, Sinho Chewi. Shifted Composition I: Harnack and Reverse Transport Inequalities. *IEEE Transactions on Information Theory, 2024, to appear.*
- J3 Jason Altschuler, Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Journal of the ACM, 2024.* (Preliminary conference version in C2.)
- J4 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 C4.)
- J5 Jason Altschuler, Pablo Parrilo. Kernel approximation on algebraic varieties. *SIAM Journal on Applied Algebra and Geometry, 2023.*
- J6 Jason Altschuler, Enric Boix-Adserà. Polynomial-time algorithms for Multimarginal Optimal Transport problems with structure. *Mathematical Programming, 2023.*
- J7 Jason Altschuler, Jonathan Niles-Weed, Austin Stromme. Asymptotics for semi-discrete entropic optimal transport. *SIAM Journal on Mathematical Analysis, 2022.*
- J8 Jason Altschuler, Pablo Parrilo. Near-linear convergence of the Random Osborne algorithm for Matrix Balancing. *Mathematical Programming, 2023.*
- J9 Jason Altschuler, Pablo Parrilo. Approximating Min-Mean-Cycle for low-diameter graphs in near-optimal time and memory. *SIAM Journal on Optimization, 2022.*
- J10 Jason Altschuler, Enric Boix-Adserà. Wasserstein barycenters are NP-hard to compute. *SIAM Journal on Mathematics of Data Science, 2022.*
- J11 Jason Altschuler, Enric Boix-Adserà. Wasserstein barycenters can be computed in polynomial time in fixed dimension. *Journal of Machine Learning Research, 2021.*
- J12 Jason Altschuler, Enric Boix-Adserà. Hardness results for Multimarginal Optimal Transport problems. *Discrete Optimization, 2021.*

- J13 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 C7.)
- J14 Jason Altschuler, Kunal Talwar. Online learning over a finite action set with limited switching. *Mathematics of Operations Research*, 2021. (Preliminary conference version in C8.)
- J15 Jason Altschuler, Victor-Emmanuel Brunel, Alan Malek. Best arm identification for contaminated bandits. *Journal of Machine Learning Research*, 2019.
- J16 Jason Altschuler, Elizabeth Yang. Inclusion of forbidden minors in random representable matroids. *Discrete Mathematics*, 2017.
- J17 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, Weijie Su, Jason Altschuler. Shifted Interpolation for Differential Privacy. *International Conference on Machine Learning (ICML)*, 2024, to appear.
- C2 Jason Altschuler, Sinho Chewi. Faster high-accuracy log-concave sampling via algorithmic warm starts. *Foundations of Computer Science (FOCS)*, 2023.
- C3 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.
- C4 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**
- C5 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**
- C6 Jason Altschuler, Francis Bach, Alessandro Rudi, Jonathan Weed. Massively scalable Sinkhorn distances via the Nyström method. *Neural Information Processing Systems Conference (NeurIPS)*, 2019.
- C7 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.
- C8 Jason Altschuler, Kunal Talwar. Online learning over a finite action set with limited switching. *Conference on Learning Theory (COLT)*, 2018.
- C9 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**
- C10 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 EXPERIENCES

UPenn

- STAT 9910: Stability in Optimization and Statistics (grad course) *Fall 2024*
- STAT 4300: Probability **Wharton Teaching Excellence Award** *Spring 2024*
- STAT 9910: Advanced Topics in Optimization (grad course) *Fall 2023*

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*

SUMMER INTERNSHIP EXPERIENCES

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>
UT Southwestern Medical Center — Research Intern	<i>Summer 2012</i>

SERVICE AND LEADERSHIP

Organizing seminars and workshops

Initiator & organizer, UPenn Optimization seminar [[link]]	<i>2023-current</i>
Organizer, ICCOPT session on “Optimization for data science”	<i>2025</i>
Co-organizer, NYU Math and Data Science seminar	<i>2022-2023</i>
Co-organizer, NeurIPS workshop on “Optimal Transport and Machine Learning”	<i>2021</i>
Co-organizer, SIAM Optimization mini-symposium on “Optimal Transport: Theory and Algorithms”	<i>2021</i>

Internal

PhD admissions committee, UPenn Statistics and Data Science department	<i>2022-2024</i>
Mentor underrepresented minorities in MIT’s Graduate Application Assistance Program (GAAP)	<i>2020-2022</i>
Reader for MIT EECS PhD applications	<i>2019-2020</i>

Reviewing

SIAM Journal on Mathematics of Data Science (SIMODS), Journal of Machine Learning Research (JMLR), Foundations of Computational Mathematics (FOCM), Mathematical Programming (MAPR), Operations Research (OR), Annals of Statistics (AoS), Information and Inference, Transactions on Automatic Control (TACON), 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

Portuguese American Optimization Workshop (invited lecture)	<i>June 2025</i>
Institute for Pure & Applied Mathematics, Workshop on Optimal Transport	<i>May 2025</i>
Les Houches, Workshop on Physics of AI Algorithms	<i>January 2025</i>
NeurIPS, OPT + ML workshop (plenary speaker)	<i>December 2024</i>
UW Madison, SILO seminar	<i>November 2024</i>
Rutgers, MSIS seminar	<i>November 2024</i>
Columbia, Applied Probability Seminar	<i>November 2024</i>
MIT, Sloan OR/STAT seminar	<i>September 2024</i>
Harvard, Statistics Colloquium	<i>September 2024</i>
Oaxaca BIRS, Workshop on Computational Harmonic Analysis in Data Science and ML	<i>September 2024</i>
Princeton ML Theory Summer School (principal speaker)	<i>August 2024</i>
ISMP, Special session for Tucker Prize recipients	<i>July 2024</i>
ISMP, Session on semidefinite programming	<i>July 2024</i>
AMS Sectional Meeting, Workshop on Recent Advances in Optimal Transport	<i>April 2024</i>
CISS, Special session on ML Theory	<i>March 2024</i>
Apple, Workshop on Privacy Preserving ML	<i>March 2024</i>
UPenn, CS Theory Seminar	<i>February 2024</i>
Singapore Institute for Mathematical Sciences, Mathematics of Data Program	<i>January 2024</i>
MIT, ORC Seminar	<i>December 2023</i>
Berkeley, BLISS Seminar	<i>December 2023</i>
Princeton, Machine Learning Theory Seminar	<i>November 2023</i>
UPenn, Electrical Engineering colloquium	<i>October 2023</i>
Kantorovich Initiative Seminar	<i>October 2023</i>
Columbia, Workshop on Robust Statistics & Privacy	<i>October 2023</i>
Duke, Decision Sciences seminar	<i>October 2023</i>
MIT, LIDS seminar	<i>September 2023</i>
INSEAD, Workshop on Information and Learning	<i>July 2023</i>
FOCM, Foundations of Data Science and Machine Learning workshop	<i>June 2023</i>
SIOP, Algebraic and Geometric Optimization session	<i>May 2023</i>
Princeton CCR colloquium	<i>May 2023</i>
Columbia, Statistics seminar	<i>April 2023</i>
MIT, LIDS	<i>April 2023</i>
Harvard, differential privacy seminar	<i>April 2023</i>
École Normale Supérieure, interteam seminar	<i>January 2023</i>
NYU, Data Science lunch seminar	<i>October 2022</i>
Yale, Statistics and Data Science seminar	<i>September 2022</i>
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	<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 best talk award	<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, workshop on noncommutative analysis and computational complexity	<i>October 2019</i>
Google Brain, invited seminar	<i>August 2019</i>

OTHER ACTIVITIES

Chess

- US National Scholastic Champion *4 times*
- “International Master” norm at Blau Escacs Open (Spain) *2015*
- “National Master” title at age 15 *2010*
- “Expert” title at age 12 *2007*
- Ranked top 5 in world for speed-chess (bullet and blitz) on chess.com *2013*
- Won 2014 US Amateur Team Championship with Princeton team *2014*
(featured in New York Times and US Chess news articles)

Swimming

- Varsity swim team at Oxford University *Spring 2015*