

Curriculum Vitae

Name: Haihao Lu

Sloan Department (Group):
Operations Research (OR) and Statistics

I. Education

<i>Degree, Subject</i>	<i>School</i>	<i>Date</i>
Ph.D., Mathematics & Operations Research	MIT	2014-2019
B.S., Applied Mathematics	Shanghai Jiao Tong University	2010-2014

II. Title of Doctoral Thesis

“Large-Scale Optimization Methods for Data-Science Applications” Robert Freund,
Thesis Advisor

III. Principal Fields of Interest

Large-scale optimization, machine learning and data-driven decision making.

IV. Name and Rank of Other Sloan Faculty in Same Field

Arnold Barnett	Professor
Cynthia Barnhart	Professor
Dimitris Bertsimas	Professor
David Gamarnik	Professor
Swati Gupta	Associate Professor (without tenure)
Alexandre Jacquillat	Associate Professor (without tenure)
Thomas Magnanti	Institute Professor
Rahul Mazumder	Associate Professor (with tenure)
James B. Orlin	Professor
Georgia Perakis	Professor
Chara Podimata	Assistant Professor
Andy Sun	Associate Professor (with tenure)
Roy Welsch	Professor

V. Non-MIT Employment

<i>Employer</i>	<i>Position</i>	<i>Start</i>	<i>End</i>
University of Chicago	Assistant Professor of Operations Management	2020	2024
Google Research	Visiting Researcher	2019	2020

VI. History of MIT Appointments

<i>Rank</i>	<i>Start</i>	<i>End</i>
The Cecil and Ida Green Career Development Assistant Professor	07/2025	Present
Assistant Professor, Operations Research and Statistics	09/2024	Present

VII. MIT Activities

<i>Committee</i>	<i>Start</i>	<i>End</i>
MIT Social and Ethical Responsibilities of Computing (SERC) Seed Grant Review Committee	2026	2026
Master of Business Analytics (MBAn) Capstone advisor Yi Wang and Joyce Zhang	2026	2026
Felicia Ann Nappi and Gabe Miller	2025	2025
Sloan OR/Statistics Group Hiring Committee	2025	2025
MIT Operations Research Center (ORC) Master's Program Admissions Committee	2025	2025
MIT Sloan MBAn Program Admissions Reader	2024, 2025	

VIII. Governmental Committees and Service

IX. Consulting Activities

<i>Name</i>	<i>Location</i>	<i>Start</i>	<i>End</i>
Google Research, Research Scientist (20% part-time consulting)	New York City, NY	2022	Present

X. Other Activities

XI. Awards

	<i>Date</i>
Sloan Research Fellowship (in Mathematics)	2026
COIN-OR (Computational Infrastructure for Operations Research) Cup winner	2024
Mathematical Programming Society Beale-Orchard-Hays Prize	2024
INFORMS Revenue Management and Pricing Section Prize	2023
INFORMS Michael H. Rothkopf Junior Researcher Paper Prize (First Place)	2022
Special invited session on "spotlight beyond WINE" (conference on Web and Internet Economics)	2022
INFORMS Optimization Society Young Researcher Prize	2021

<i>Grants</i>	<i>Date</i>
• "Scalable First-Order Methods for Large Scale Optimization: Theory and Computation", Air Force Office of Scientific Research (\$400,000)	2024-2027
• "Collaborative Proposal: Scaling up Linear Programming with First-Order Methods and MPI", Office of Naval Research (my part \$456,409)	2024-2027

- “Addressing Regressive Property Taxation with Fairness-Constrained Differentiable Machine Learning,” MIT Social and Ethical Responsibilities of Computing Seed Grant (\$75,000) 2025
- MIT Research Support Committee award (\$90,000) 2025
- MIT Sloan Junior Faculty Research Assistance Program (\$30,000) 2025
- MIT Sloan Junior Faculty Research Assistance Program (\$30,000) 2024

XII. Professional Membership and Activities

	<i>Start</i>	<i>End</i>
<i>Professional Society Leadership</i>		
Vice Chair (Linear and Conic Optimization), INFORMS Optimization Society	2024	2026
<i>Conference and Seminar Organizer</i>		
Stream Organizer, International Symposium on Mathematical Programming (scheduled)	2027	2027
Program committee, INFORMS Optimization Society Conference	2026	2026
Cluster chair, INFORMS Annual Meeting	2025	2025
Cluster chair, INFORMS Annual Meeting	2024	2024
Program committee, INFORMS Optimization Society Conference	2024	2024
Organizer, workshop on modern continuous optimization at MIT	2023	2023
<i>Award Committees</i>		
Member, Land-Doig MIP (Mixed Integer Programming) Competition committee	2026	2026
Chair, COIN-OR Cup committee	2025	2025
Judge, INFORMS George Nicholson student paper competition	2023, 2024	
<i>Review Boards</i>		
Associate Editor, Mathematics of Operations Research	2026	Present
Editorial Board of Reviewers, Journal of Machine Learning Research	2020	Present
<i>Reviewer for Journals and Conferences</i>		
Computational Optimization and Applications		
Foundations of Computational Mathematics		
ICML (International Conference on Machine Learning)		
IEEE Transactions on Image Processing		
INFORMS Journal on Computing		
INFORMS Journal on Optimization		
Journal of Optimization Theory and Applications		
Machine Learning		
Management Science		
Manufacturing & Service Operations Management		
Mathematical Programming		
Mathematical Programming Computation		
Mathematics of Operations Research		
NeurIPS (Conference on Neural Information Processing Systems)		

Operations Research
 Production and Operations Management
 SIAM Journal on Control and Optimization
 SIAM Journal on Optimization
 WebConf (Web Conference)

Memberships

The Institute for Operations Research and the Management Sciences (INFORMS)
 Mathematical Optimization Society (MOS)
 Society for Industrial and Applied Mathematics (SIAM)

Prior Professional Activities at Chicago Booth	Start	End
Chicago Booth Operations Management (OM) Faculty Recruiting Committee Member	2022	2023
Chicago Booth OM Workshop Co-Organizer	2021	2022

XIII. Subjects Taught

MIT

Instructor

<i>Number</i>	<i>Title</i>	<i>Date</i>
15.071	The Analytics edge One section (MBA) 6.5/7 Overall Teaching; 4.4/5 Recommend Professor	Spring 2025

15.C57/15.C571/ 6.C57/6.C571/ IDS.C57	Optimization Two sections (Undergraduate and Master's) 5.1/7 Overall Teaching; 4.0/5 Recommend Professor	Fall 2024
---	--	-----------

Guest Instructor

6.251/15.081	Introduction to Mathematical Programming – Two lectures in Fall 2024; two in Fall 2022	Fall 2024, 2022
--------------	---	-----------------

Booth School of Business, University of Chicago

Instructor

BUSN 32100	Data Analysis with R and Python (MBA)	2021-2024
– Section 32100-01 (Full-Time MBA): 4.6, 4.6, 4.7, 4.6, 4.6 (of 5)		Winter 2024
– Section 32100-02 (Full-Time MBA): 4.6, 4.4, 4.8, 4.5, 4.5		Winter 2024
– Section 32100-81 (Evening MBA): 4.2, 3.8, 4.4, 4.4, 3.9		Winter 2024
– Section 32100-01 (Full-Time MBA): 4.8, 4.7, 4.7, 4.5, 4.7		Spring 2023
– Section 32100-02 (Full-Time MBA): 4.9, 4.7, 4.8, 4.7, 4.7		Spring 2023
– Section 32100-01 (Full-Time MBA): 4.7, 4.5, 4.7, 4.3, 4.5		Spring 2022
– Section 32100-81 (Evening MBA): 4.8, 4.7, 4.9, 4.6, 4.8		Spring 2022
– Section 32100-02 (Full-Time MBA): 5.0, 4.3, 5.0, 4.8, 5.0		Spring 2021

XIV. Thesis Supervision

1. Doctoral Theses Supervised

Thesis or Dissertation Committees

<i>Name</i>	<i>Program</i>	<i>Degree</i>	<i>Year</i>
Zikai Xiong (advisor Freund)	MIT ORC	PhD	2025

2. Master's Theses Supervised

<i>Name</i>	<i>Program</i>	<i>Degree</i>	<i>Year</i>
Feiyu Han	Univ. of Chicago Statistics	MSc	2023
Luyang Zhang	Univ. of Chicago Statistics	MSc	2022
Hanyang Jiang	Univ. of Chicago Statistics	MSc	2021

3. Bachelor's Theses Supervised

4. Theses in Progress

<i>Name</i>	<i>Program</i>	<i>Degree</i>	<i>Year</i>
Luke Fitzgerald	MIT Math	BS	2026
Jinwen Yang	Univ. of Chicago Statistics	PhD	2026
Charlie Liu	MIT ORC	PhD	2029
Nicolas Villena	MIT ORC	PhD	2029
Jacob Dentes	MIT ORC	PhD	2030

Thesis or Dissertation Committees

Xiang Meng (advisor Mazumder)	MIT ORC	PhD	2026
Junhui Zhang (advisor Jaillet)	MIT EECS	PhD	2026

XV. Publications (*including order of co-authors, if any*)¹

1. Theses

“Large-Scale Optimization Methods for Data-Science Applications.

2. Refereed Journal Articles

1. “A New Crossover Algorithm for LP Inspired by the Spiral Dynamic of PDHG”, with Tianhao Liu, forthcoming in *INFORMS Journal on Computing*.

¹ Publications denoted as “with coauthors” follow alphabetical author order per convention in field, whereas publications with full author lists follow non-alphabetical order reflecting contributions.

2. "cuPDLP. jl: A GPU Implementation of Restarted Primal-Dual Hybrid Gradient for Linear Programming in Julia", with Jinwen Yang , Operations Research, 73.6 (2025): 3440-3452.
3. "On the Sparsity of Optimal Linear Decision Rules for a Class of Robust Optimization Problems with Box Uncertainty Sets", with Bradley Sturt, Operations Research, 74.1 (2026): 500-516.
4. "A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming", with Jinwen Yang, Mathematical Programming, 215.1 (2026): 771-808.
5. "Optimizing Scalable Targeted Marketing Policies with Constraints", with Duncan Simester and Yuting Zhu, Marketing Science, 44.5 (2025): 1082-1103.
6. "Regularized Online Allocation Problems: Fairness and Beyond", with Santiago Balseiro and Vahab Mirrokni, Manufacturing & Service Operations Management, 27.3 (2025):720-735.
7. "Auto-bidding and Auctions in Online Advertising: A Survey", with many people from Google, ACM SIGecom Exchanges 22.1 (2024): 159-183.
8. "On the Geometry and Refined Rate of Primal-Dual Hybrid Gradient for Linear Programming", with Jinwen Yang, Mathematical Programming 212 (2025): 349-387.
9. "Infeasibility Detection with Primal-Dual Hybrid Gradient for Large-Scale Linear Programming", with David Applegate, Mateo Diaz and Miles Lubin, SIAM Journal on Optimization 34.1 (2024): 459-484.
10. "On the Linear Convergence of Extra-Gradient Methods for Nonconvex-Nonconcave Minimax Problems", Saeed Hajizadeh, Haihao Lu and Benjamin Grimmer, INFORMS Journal on Optimization 6.1 (2024): 19-31.
11. "A J-Symmetric Quasi-Newton Method for Minimax Problems", with Azam Asl and Jinwen Yang, Mathematical Programming 204.1 (2024): 207-254.
12. "The landscape of the proximal point method for nonconvex-nonconcave minimax optimization", Benjamin Grimmer, Haihao Lu, Pratik Worah and Vahab Mirrokni, Mathematical Programming 201.1-2 (2023): 373-407.
13. "Faster First-Order Primal-Dual Methods for Linear Programming using Restarts and Sharpness", with David Applegate, Oliver Hinder and Miles Lubin, Mathematical Programming 201.1-2 (2023): 133-184.

14. “Frank-Wolfe Methods with an Unbounded Feasible Region and Applications to Structured Learning”, Haoyue Wang, Haihao Lu and Rahul Mazumder, *SIAM Journal on Optimization* 32.4 (2022): 2938-2968.
15. “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems”, with Santiago Balseiro and Vahab Mirrokni, *Operations Research* 71.1 (2022): 101-119.
16. “An $O(s^r)$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to the Linear Convergence of Minimax Problems”, Haihao Lu, *Mathematical Programming* 194 (2022): 1061–1112.
17. “Randomized Gradient Boosting Machines”, Haihao Lu and Rahul Mazumder, *SIAM Journal on Optimization* 2020, 30(4), 2780-2808.
18. “Generalized Stochastic Frank-Wolfe Algorithm with Stochastic ‘Substitute’ Gradient for Structured Convex Optimization”, Haihao Lu and Robert M. Freund, *Mathematical Programming* 187.1 (2021): 317-349.
19. “ ‘Relative-Continuity’ for Non-Lipschitz Non-Smooth Convex Optimization using Stochastic (or Deterministic) Mirror Descent”, Haihao Lu, *INFORMS Journal on Optimization*, 2019, 1(4): 288-303.
20. “Relatively-Smooth Convex Optimization by First-Order Methods, and Applications”, Haihao Lu, Robert M. Freund and Yurii Nesterov, *SIAM Journal on Optimization*, 2018, 28(1): 333-354.
21. “New Computational Guarantees for Solving Convex Optimization Problems with First Order Methods, via a Function Growth Condition Measure”, Robert M. Freund, Haihao Lu, *Mathematical Programming* 2018, Vol.170, No.2: 445-477.
22. “Stochastic linearization of turbulent dynamics of dispersive waves in equilibrium and nonequilibrium state”, Shi-xiao W. Jiang, Haihao Lu, Douglas Zhou, and David Cai, *New Journal of Physics*, 2016, 18(8): 083028.
23. “Renormalized Dispersion Relations of β -Fermi-Pasta-Ulam Chains in Equilibrium and Nonequilibrium states”, Shi-xiao W. Jiang, Haihao Lu, Douglas Zhou, and David Cai. *Physical Review E*, 2014, 90(3): 032925.

3. Articles in Refereed Conference Proceedings

- “A Field Guide for Pacing Budget and ROS Constraints”, with Santiago Balseiro, Kshipra Bhawalkar, Zhe Feng, Vahab Mirrokni, Balasubramanian Sivan and Di Wang, ICML 2024.
- “Online Ad Procurement in Non-stationary Autobidding Worlds”, with Jason Liang and Baoyu Zhou, NeurIPS 2023.

“Limiting Behaviors of Nonconvex-Nonconcave Minimax Optimization via Continuous-Time Systems”, Benjamin Grimmer, Haihao Lu, Pratik Worah and Vahab Mirrokni, ALT 2022.

“Practical Large-Scale Linear Programming using Primal-Dual Hybrid Gradient”, with David Applegate, Oliver Hinder, Mateo Diaz, Miles Lubin, Brendan O’Donoghue, and Warren Schudy, NeurIPS 2021.

“Regularized Online Allocation Problems: Fairness and Beyond”, with Santiago Balseiro and Vahab Mirrokni, ICML 2021.

“Contextual Reserve Price Optimization in Auctions via Mixed Integer Programming”, Joey Huchette, Haihao Lu, Hossein Esfandiari and Vahab Mirrokni, NeurIPS, 2020.

“Dual Mirror Descent for Online Allocation Problems”, with Santiago Balseiro and Vahab Mirrokni, ICML, 2020.

“Accelerating Gradient Boosting Machines”, Haihao Lu, Sai Praneeth Karimireddy, Natalia Ponomareva and Vahab Mirrokni, AISTATS, 2020.

“Ordered SGD: A Stochastic First-Order Method for Ordered Empirical Risk Minimization”, Kenji Kawaguchi and Haihao Lu, AISTATS, 2020.

“Accelerating Greedy Coordinate Descent Methods”, Haihao Lu, Robert M. Freund and Vahab Mirrokni, ICML, 2018.

“Approximate Leave-One-Out for Fast Parameter Tuning in High Dimensions”, Shuaiwen Wang, Wenda Zhou, Haihao Lu, Arian Maleki, Vahab Mirrokni, ICML, 2018.

3b. Articles in Refereed Workshop

“MPAX: Mathematical Programming in JAX”, with Zedong Peng and Jinwen Yang, NeurIPS workshop on scalable optimization, 2025.

4. Articles in Non-Refereed Conference Proceedings

5. Papers/Articles in Progress or Under Review

24. “cuPDLP+: A Further Enhanced GPU-Based First-Order Solver for Linear Programming”, with Zedong Peng and Jinwen Yang.

25. “An Overview of GPU-based First-Order Methods for Linear Programming and Extensions”, with Jinwen Yang.

26. “New Understandings and Computation on Augmented Lagrangian Methods for Low-Rank Semidefinite Programming”, with Lijun Ding and Jinwen Yang.
27. “PDLP: A Practical First-Order Method for Large-Scale Linear Programming”, with David Applegate, Mateo D’iaz, Oliver Hinder, Miles Lubin, Brendan O’Donoghue, and Warren Schudy.
28. “MPAX: Mathematical Programming in JAX”, with Zedong Peng and Jinwen Yang.
29. “PDOT: a Practical Primal-Dual Algorithm and a GPU-Based Solver for Optimal Transport”, with Jinwen Yang.
30. “Restarted Halpern PDHG for Linear Programming”, with Jinwen Yang.
31. “Achieving Fairness and Accuracy in Regressive Property Taxation”, with Feiyu Han and Ozan Candogan.
32. “On the Convergence of L-shaped Algorithms for Two-Stage Stochastic Programming”, with John Birge and Baoyu Zhou.
33. “On the Infimal Sub-differential Size of Primal-Dual Hybrid Gradient Method”, with Jinwen Yang.
34. “On a Unified and Simplified Proof for the Ergodic Convergence Rates of PPM, PDHG and ADMM”, with Jinwen Yang.
35. “Analysis of Dual-Based PID Controllers through Convolutional Mirror Descent”, with Santiago Balseiro, Vahab Mirrokni and Balasubramanian Sivan.
36. “Nearly Optimal Linear Convergence of Stochastic Primal-Dual Methods for Linear Programming”, with Jinwen Yang.
37. “Approximate Leave-One-Out for High-Dimensional Non-Differentiable Learning Problems”, Shuaiwen Wang, Wenda Zhou, Arian Maleki, Haihao Lu and Vahab Mirrokni.

6. Other Publications: Media Coverage and Blogs

Google research blog: Scaling up linear programming with PDLP.

SIAG on Optimization Views and News: First-Order Methods for Linear Programming.

NVIDIA technical blog: Accelerate Large Linear Programming Problems with NVIDIA cuOpt.

Quantum Zeitgeist: Google's PDLP Solver Revolutionizes Linear Programming Problem Solving Efficiency.

Marktechpost: PDLP (Primal-Dual Hybrid Gradient Enhanced for LP): A New FOM-based Linear Programming LP Solver that Significantly Scales Up Linear Programming LP Solving Capabilities

TMT Post: Cardinal Optimizer is ushering in a new "Age of Exploration" for operations research with GPU chips.

7. Technical Reports

"Depth Creates No Bad Local Minima", Haihao Lu and Kenji Kawaguchi, Technical Report.

XVI. Invited Oral Presentations

2026:

- "An Overview of GPU-Based First-Order Methods for Linear Programming and Beyond", Amazon Inc, STOC team

2025:

- "An Overview of GPU-Based First-Order Methods for Linear Programming and Beyond", Tepper Business School, Carnegie Mellon University
- "An Overview of GPU-Based First-Order Methods for Linear Programming and Beyond" NeurIPS Workshop on ScaleOPT: GPU-Accelerated and Scalable Optimization
- "An Overview of GPU-Based First-Order Methods for Linear Programming and Beyond" INFORMS Annual Meeting
- "GPU-Accelerated Linear Programming and Beyond" AFOSR Annual Meeting
- "An Overview of GPU-Based First-Order Methods for Linear Programming and Beyond" Cargese Workshop on Combinatorial Optimization (a 4-hours tutorial)
- "GPU-Accelerated Linear Programming and Beyond" International Conference on Continuous Optimization (ICCOPT)
- "GPU-Accelerated Linear Programming and Beyond" MIP Workshop
- "GPU-Accelerated Linear Programming and Beyond" Workshop on Strategies for Handling Applications with Nonconvexity
- "GPU-Accelerated Linear Programming and Beyond" Workshop on Combinatorial Optimization and Data Science
- "GPU-Accelerated Linear Programming and Beyond" Gurobi Inc

2024:

- “GPU-Accelerated Linear Programming and Beyond” NVIDIA
- “GPU-Accelerated Linear Programming and Beyond” ORIE Department, Cornell University
- “GPU-Accelerated Linear Programming and Beyond” Operations Research Center, MIT
- “GPU-Accelerated Linear Programming and Beyond” School of Engineering and Applied Sciences, Harvard University
- “GPU-Accelerated Linear Programming and Beyond” INFORMS annual meeting
- “GPU-Accelerated Linear Programming and Beyond” AFOSR annual meeting
- “GPU-Accelerated Linear Programming and Beyond” Google Research
- “GPU-Accelerated Linear Programming and Beyond” Workshop for Yinyu Ye’s Retirement Celebration, Stanford University
- “A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming” International Symposium on Mathematical Programming
- “GPU-Accelerated Linear Programming and Beyond” Industrial Engineering and Management Sciences, Northwestern University
- “GPU-Accelerated Linear Programming and Beyond” Industrial and Systems Engineering, University of Minnesota
- “A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming” INFORMS Conference on Optimization

2023:

- “First-Order Methods for Constrained Continuous Optimization” Operations Research Center, MIT
- “First-Order Methods for Constrained Continuous Optimization” DRO and IEOR joint seminar, Columbia University
- “First-Order Methods for Constrained Continuous Optimization” Viterbi School of Engineering, University of Southern California
- “First-Order Methods for Constrained Continuous Optimization” School of Data Science, Chinese University of Hong Kong
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” INFORMS Annual Meeting
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” SIAM Conference on Optimization
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Shanghai University of Finance and Economics

2022:

- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” INFORMS Annual Conference
- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Simons Institute for the Theory of Computing

- “First-Order Methods for Linear Programming: Theory, Computation and Applications” INFORMS Annual Conference
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Argonne National Lab
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Applied Mathematics Department, John Hopkins University
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” School of Operations Research and Information Engineering, Cornell University
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Operations Research Center, MIT
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Institute for Interdisciplinary Information Sciences (Yao’s class), Tsinghua University

2021:

- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Institute for Data, Systems, and Society, MIT
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Booth School of Business, University of Chicago
- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Stern School of Business, New York University
- “First-Order Methods for Linear Programming: Theory, Computation and Applications” Tepper School of Business, Carnegie Mellon University
- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Stern School of Business, New York University
- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Lioutaud Graduate School of Business, University of Illinois at Chicago

2020:

- “An $O(\text{sr})$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems” Chinese Academy of Sciences
- “An $O(\text{sr})$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems” AI+Math Colloquia, Shanghai Jiao Tong University
- “An $O(\text{sr})$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems” INFORMS Annual Meeting
- “Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization” International Conference on Artificial Intelligence and Statistics
- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Ebay Inc
- “Accelerated Gradient Boosting Machines” International Conference on Artificial Intelligence and Statistics

- “The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems” Google Research NYC
- “An $O(sr)$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems” Google Research NYC
- “An $O(sr)$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems” SIAM Conference on Mathematics of Data Science

2019:

- “Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization” Rensselaer Polytechnic Institute, Mathematics Department
- “Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization” NYU, Courant Institute of Mathematical Science
- “Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization” Google Research
- “Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization” International Conference on Continuous Optimization
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” Google Research
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” Columbia University, IEOR
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Wisconsin, Madison, ISE
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Illinois, Urbana-Champaign, ISE
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” Google Research, Cambridge
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” National University of Singapore, ISEM
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Toronto, Mathematics and Computer Science
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Toronto, MIE
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Chicago
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” University of Minnesota, Twin Cities, ISE
- “Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity” Duke University, Mathematics and Computer Science

2018:

- “Randomized Gradient Boosting Machines” INFORMS Annual Meeting
- “Scalable Linear Programming via First-Order Methods” Princeton Optimization Day (Poster)

- “Scalable Linear Programming via First-Order Methods” Google Research, New York City
- “Generalized Stochastic Frank-Wolfe Algorithm with Stochastic ‘Substitute’ Gradient for Structured Convex Optimization” Columbia University, Statistics Department
- “Generalized Stochastic Frank-Wolfe Algorithm with Stochastic ‘Substitute’ Gradient for Structured Convex Optimization” International Symposium on Mathematical Programming (ISMP)
- “Accelerating Greedy Coordinate Descent Methods” Google Research, New York City
- “Accelerating Greedy Coordinate Descent Methods” International Conference on Machine Learning
- “Approximate Leave-One-Out for Fast Parameter Tuning in High Dimensions” International Conference on Machine Learning (Poster)
- “Accelerating Greedy Coordinate Descent Methods” NYAS Meeting on Machine Learning (Poster)
- “Approximate Leave-One-Out for Fast Parameter Tuning in High Dimensions” NYAS Meeting on Machine Learning (Poster)
- “‘Relative-Continuity’ for Non-Lipschitz Non-Smooth Convex Optimization using Stochastic (or Deterministic) Mirror Descent” INFORMS Meeting on Optimization

2017:

- “‘Relative-Continuity’ for Non-Lipschitz Non-Smooth Convex Optimization using Stochastic (or Deterministic) Mirror Descent” INFORMS Annual Meeting
- “Relatively-Smooth Convex Optimization by First-Order Methods, and Applications” SIAM Conference on Optimization

2016:

- “Relatively-Smooth Convex Optimization by First-Order Methods, and Applications” INFORMS Annual Meeting
- “Extending the Scope of ‘Smooth’ and ‘Non-Smooth’ Convex Optimization via First-Order Methods” University of Edinburgh

2015:

- “Some New Results for Randomized Coordinate Gradient Descent” International Symposium on Mathematical Programming (ISMP)