Haihao (Sean) Lu

Contact Information

Room 585, Sloan School of Management
100 Main St, Cambridge, MA 02142

Email: haihao@mit.edu
Phone: 857-998-3092

Research Interests

I develop optimization theory, algorithms, and software for large-scale decision-making. My research connects mathematical foundations with practical questions: how to exploit modern computing architectures, explain and improve algorithmic performance, allocate scarce resources under uncertainty, and make better data-driven decisions.

Appointment

Massachusetts Institute of Technology, Cambridge, MA — September 2024-present
Sloan School of Management
The Cecil and Ida Green Career Development Assistant Professor (since July 2025)
Assistant Professor of Operations Research and Statistics

Google Research, New York City, NY — December 2022-present
Research scientist (20% part time consulting)

University of Chicago, Chicago, IL — July 2020-June 2024
Booth School of Business
Assistant Professor of Operations Management
Paternity leave in 2020-2021

Google Research, New York City, NY
Visiting Faculty Researcher — July 2019-June 2020
Large-Scale Optimization Team

Education

Massachusetts Institute of Technology, Cambridge, MA — 2014-2019
Ph.D. in Mathematics & Operations Research,
Thesis Title: Large-Scale Optimization Methods for Data-Science Applications
Advisor: Prof. Robert M. Freund
Thesis Committee Members: Robert M. Freund, Rahul Mazumder, Ankur Moitra, Jonathan Kelner

Shanghai Jiao Tong University, Shanghai, China — 2010-2014
B.S. in Applied Mathematics, Zhiyuan College, Graduation with distinction
Advisor: Prof. David Cai

Selected Honors

INFORMS Computing Society Prize. 2026
This prize is awarded to the best paper or group of related papers that advance the state of the art in the interface of operations research and computer science.

Sloan Research Fellowship (Mathematics). 2026
The two-year $75,000 fellowships are awarded annually to early-career researchers whose creativity, innovation, and research accomplishments make them stand out as the next generation of leaders.

COIN-OR Cup winner. 2024
The COIN-OR Cup recognizes and celebrates the best contributions to open-source operations research software development and use associated with the COIN-OR software collection.

The Beale–Orchard-Hays Prize. 2024

This prize, awarded every three years by the Mathematical Optimization Society, recognizes excellence in computational mathematical programming.

INFORMS Revenue Management and Pricing Section Prize. 2023

This prize is awarded for the best contribution to the science of pricing and revenue management published in English.

Michael H. Rothkopf Junior Researcher Paper Prize (first place). 2022

This prize recognizes the best paper by a young researcher in the area of Auctions and Market Design.

Special invited session on "spotlight beyond WINE". 2022

This session at WINE highlights some of the best works in algorithmic game theory that appeared in conferences and journals other than WINE, or mature working papers.

INFORMS Optimization Society Young Researcher Prize. 2021

This prize serves as an esteemed recognition of promising colleagues in the field of optimization who are at the beginning of their academic or industrial career.

Research Grants

GE Vernova 2026-2027

"GPU Solver Engine for Grid Orchestration at Scale." My share: $150,000.

Air Force Office of Scientific Research 2024-2027

"Scalable First-Order Methods for Large Scale Optimization: Theory and Computation." Award: $400,000.

Office of Naval Research 2024-2027

"Collaborative Proposal: Scaling up Linear Programming with First-Order Methods and MPI." My share: $456,409.

MIT Social and Ethical Responsibilities of Computing 2025

"Addressing Regressive Property Taxation with Fairness-Constrained Differentiable Machine Learning." Award: $75,000.

MIT Research Support Committee 2025

Award: $90,000.

MIT Sloan Junior Faculty Research Assistance Program 2025

Award: $30,000.

MIT Sloan Junior Faculty Research Assistance Program 2024

Award: $30,000.

eBay Research Gift 2022-2024

"Efficient Algorithms for Solving Large Scale Choice-Based Model and Applications to Recommendation System." Award: $100,000.

Journal Publications[1]

1. "PDOT: a Practical Primal-Dual Algorithm and a GPU-Based Solver for Optimal Transport", with Jinwen Yang, to appear in *SIAM Journal on Optimization*.
2. "PDLP: A Practical First-Order Method for Large-Scale Linear Programming", with David Applegate, Mateo Díaz, Oliver Hinder, Miles Lubin, Brendan O'Donoghue, and Warren Schudy, *Mathematical Programming Computation* (2026).

[1] Publications denoted as *with coauthors* follow alphabetical author order per convention in the field; publications with full author lists follow non-alphabetical order reflecting contributions.

3. "Analysis of Dual-Based PID Controllers through Convolutional Mirror Descent", with Santiago R. Balseiro, Vahab Mirrokni and Balasubramanian Sivan, to appear in *Operations Research*.
4. "A New Crossover Algorithm for LP Inspired by the Spiral Dynamic of PDHG", with Tianhao Liu, forthcoming in *INFORMS Journal on Computing*.
5. "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.
   - *INFORMS Computing Society Prize (2026).*
   - *COIN-OR Cup winner (2024).*
6. "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.
7. "A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming", with Jinwen Yang, *Mathematical Programming*, 215.1 (2026): 771-808.
   - *INFORMS Computing Society Prize (2026).*
8. "Optimizing Scalable Targeted Marketing Policies with Constraints", with Duncan Simester and Yuting Zhu, *Marketing Science*, 44.5 (2025): 1082-1103.
9. "Regularized Online Allocation Problems: Fairness and Beyond", with Santiago Balseiro and Vahab Mirrokni, *Manufacturing & Service Operations Management*, 27.3 (2025): 720-735.
10. "Auto-bidding and Auctions in Online Advertising: A Survey", with many people from Google, *ACM SIGecom Exchanges* 22.1 (2024): 159-183.
11. "On the Geometry and Refined Rate of Primal-Dual Hybrid Gradient for Linear Programming", with Jinwen Yang, *Mathematical Programming*, 212 (2025): 349-387.
    - *INFORMS Computing Society Prize (2026).*
12. "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.
    - *Beale–Orchard-Hays Prize (2024).*
13. "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.
14. "A J-Symmetric Quasi-Newton Method for Minimax Problems", with Azam Asl and Jinwen Yang, *Mathematical Programming* 204.1 (2024): 207-254.
15. "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.
16. "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.
    - *Beale–Orchard-Hays Prize (2024).*
17. "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.

18. "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.
    - *INFORMS Revenue Management and Pricing Section Prize (2023).*
    - *Michael H. Rothkopf Junior Researcher Paper Prize (first place, 2022).*
19. "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.
    - *INFORMS Optimization Society Young Researcher Prize (2021).*
20. "Randomized Gradient Boosting Machines", Haihao Lu and Rahul Mazumder, *SIAM Journal on Optimization* 2020, 30(4), 2780-2808.
21. "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.
22. " '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.
23. "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.
24. "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.
25. "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.
26. "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.

Refereed Conference Publications

1. "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*.
2. "Online Ad Procurement in Non-stationary Autobidding Worlds", with Jason Liang and Baoyu Zhou, *NeurIPS 2023*.
3. "Limiting Behaviors of Nonconvex-Nonconcave Minimax Optimization via Continuous-Time Systems", Benjamin Grimmer, Haihao Lu, Pratik Worah and Vahab Mirrokni, *ALT 2022*.
4. "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*.
    - *Beale–Orchard-Hays Prize (2024; conference version).*
5. "Regularized Online Allocation Problems: Fairness and Beyond", with Santiago Balseiro and Vahab Mirrokni, *ICML* 2021.

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

7. "Dual Mirror Descent for Online Allocation Problems", with Santiago Balseiro and Vahab Mirrokni, *ICML*, 2020.

8. "Accelerating Gradient Boosting Machines", Haihao Lu, Sai Praneeth Karimireddy, Natalia Ponomareva and Vahab Mirrokni, *AISTATS*, 2020.

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

10. "Accelerating Greedy Coordinate Descent Methods", Haihao Lu, Robert M. Freund and Vahab Mirrokni, *ICML*, 2018.

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

Refereed Workshop Publications

1. "MPAX: Mathematical Programming in JAX", with Zedong Peng and Jinwen Yang, *NeurIPS Workshop on Scalable Optimization*, 2025.

Preprints and Papers under Review

1. "Reformulation Techniques for PDLP with Applications to QUBO and Joint Replenishment", with Jacob Dentes.

2. "PDFP: GPU-Accelerated Feasibility Pump via PDHG", with Ruofeng Liu and Rahul Mazumder.

3. "cuPDQPx: A GPU-Accelerated Restarted Halpern PDHG for Large-Scale Convex Quadratic Programming", with Zedong Peng and Jinwen Yang.

4. "H-PDLP: GPU-Accelerated Linear Programming via Generalized PDHG and Homogeneous Self-Dual Embedding", with Bo Tang.

5. "PDNQP: A GPU-based Factorization-free Method for Large-scale Nonconvex Quadratic Programming", with Zixi Chen.

6. "On the Local Linear Convergence of Operator Splitting Methods for Conic Programming", with Lijun Ding and Jinwen Yang.

7. "FlashFolio: A GPU-Accelerated Solver for Portfolio Optimization", with Yilun Jiang, Zedong Peng and Jinwen Yang.

8. "End-to-End Learning of Correlated Operating Reserve Requirements in Security-Constrained Economic Dispatch", Owen Shen, Hung-po Chao, Haihao Lu and Patrick Jaillet.

9. "Active Set Identification and Rapid Convergence for Degenerate Primal-Dual Problems", with Mateo Díaz, Pedro Izquierdo Lehmann and Jinwen Yang.

10. "cuPDLPx: A Further Enhanced GPU-Based First-Order Solver for Linear Programming", with Zedong Peng and Jinwen Yang.

11. "Enhanced PDHG for Linear Programming with Online Preconditioning", with Wanyu Zhang.

12. "An Overview of GPU-based First-Order Methods for Linear Programming and Extensions", with Jinwen Yang.

13. "New Understandings and Computation on Augmented Lagrangian Methods for Low-Rank Semidefinite Programming", with Lijun Ding and Jinwen Yang.

14. "MPAX: Mathematical Programming in JAX", with Zedong Peng and Jinwen Yang.

15. "Restarted Halpern PDHG for Linear Programming", with Jinwen Yang.

16. "The Power of Linear Programming in Sponsored Listings Ranking: Evidence from a Large-Scale Field Experiment", with Luyang Zhang and Yuting Zhu.

17. "cuPDLP-C: A Strengthened Implementation of cuPDLP for Linear Programming by C Language", Haihao Lu, Jinwen Yang, Haodong Hu, Qi Huangfu, Jinsong Liu, Tianhao Liu, Yinyu Ye, Chuwen Zhang and Dongdong Ge.

18. "Achieving Fairness and Accuracy in Regressive Property Taxation", with Ozan Candogan and Feiyu Han.

19. "On the Convergence of L-shaped Algorithms for Two-Stage Stochastic Programming", with John R. Birge and Baoyu Zhou.

20. "On a Unified and Simplified Proof for the Ergodic Convergence Rates of PPM, PDHG and ADMM", with Jinwen Yang.

Technical Reports

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

2. "Nearly Optimal Linear Convergence of Stochastic Primal-Dual Methods for Linear Programming", with Jinwen Yang.

3. "Approximate Leave-One-Out for High-Dimensional Non-Differentiable Learning Problems", Shuaiwen Wang, Wenda Zhou, Arian Maleki, Haihao Lu and Vahab Mirrokni.

Teaching Experience

MIT:[2]

Instructor

- 15.071. The Analytics Edge
  - Two sections (MBA): 6.0/7, 4.3/5. Spring 2026
  - One section (MBA): 6.5/7, 4.4/5. Spring 2025
- 15.C57/15.C571/6.C57/6.C571/IDS.C57. Optimization
  - Two sections (Undergraduate and Master): 5.1/7, 4.0/5. Fall 2024

Guest Instructor

- 6.251/15.081. Introduction to Mathematical Programming
  - Two lectures in Fall 2022
  - Two lectures in Fall 2024

Booth School of Business, University of Chicago:[3]

[2]Two evaluation questions: "Overall rating" (out of 7), "Recommend Professor" (out of 5).

[3]Five evaluation questions (out of 5): "Did the instructor convey the course material clearly," "Did the instructor convey the material in an interesting way," "Did you take away useful tools,

Instructor

- BUSN 32100 Data Analysis with R and Python (MBA)
  - Section 32100-01 (Full-Time MBA): 4.6, 4.6, 4.7, 4.6, 4.6. 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
- BUSN 36919 Modern Large-Scale Optimization: Theory and Computation (PhD)
  - Section 36919-50: 4.8, 4.5, 5.0, 4.5, 4.5. Spring 2023

Students

Current Members

Ph.D. Students (Theses in Progress)

- Charlie Liu (MIT ORC), co-advised with Rahul Mazumder, expected 2029
- Jacob Dentes (MIT ORC), expected 2030

Postdoctoral Researchers

- Zedong Peng (MIT)
- Bo Tang (MIT)
- Vinit Ranjan (MIT)
- Yuchen Lou (MIT)

Alumni

Ph.D. Students

- Jinwen Yang (UChicago Statistics)
  Next position: Assistant Professor, Columbia University

Postdoctoral Researchers

- Azam Asl (UChicago Booth)
  Next position: consulting firm
- Baoyu Zhou (UChicago Booth)
  Next position: Assistant Professor, Arizona State University

Master's Theses Supervised

- Feiyu Han (UChicago Statistics), M.Sc. 2023
  Next position: Ph.D. program, UChicago Booth
- Luyang Zhang (UChicago Applied Mathematics), M.Sc. 2022
  Next position: Ph.D. program, CMU Heinz
- Hanyang Jiang (UChicago Applied Mathematics), M.Sc. 2021
  Next position: Ph.D. program, Georgia Tech ISE

Undergraduates

- Luke Fitzgerald (MIT Mathematics)
- Yang Meng (UChicago CS)
  Next position: Ph.D. program, UChicago CS

concepts, and/or insights from this course," "How much did you get out of this course," and "Would you recommend this course to other students."

- Wanyu Zhang (Shanghai University of Finance and Economics)
  Next position: Ph.D. program, Stanford MS&E
- Hongpei Li (Shanghai University of Finance and Economics)
  Next position: Ph.D. program, Northwestern IEMS

Additional Advising

MIT MBAn Capstone Advising

- Yi Wang and Joyce Zhang 2026
- Felicia Ann Nappi and Gabe Miller 2025

Service on Thesis or Dissertation Committees

- Zikai Xiong (MIT ORC), Ph.D. 2025, advised by Robert M. Freund
- Xiang Meng (MIT ORC), Ph.D. expected 2026, advised by Rahul Mazumder
- Junhui Zhang (MIT EECS), Ph.D. expected 2026, advised by Patrick Jaillet

Academic Service

Editorial Service

- Associate Editor,
  Mathematics of Operations Research 2026-present
  Computational Optimization and Applications 2026-present
- Editorial Board of Reviewers, *Journal of Machine Learning Research* 2020-present

Professional Society Leadership

- Vice Chair (Linear and Conic Optimization), INFORMS Optimization Society 2024-2026

Conference and Workshop Organization

- Stream Organizer, International Symposium on Mathematical Programming 2027
- Program Committee, INFORMS Optimization Society Conference 2026, 2024
- Cluster Chair, INFORMS Annual Meeting 2025, 2024
- Organizer, Workshop on Modern Continuous Optimization at MIT 2023

Award Committees

- Member, COIN-OR Cup Committee 2026
- Member, Land-Doig MIP Competition Committee 2026
- Chair, COIN-OR Cup Committee 2025
- Judge, INFORMS George Nicholson Student Paper Competition 2024, 2023

University Committee Service

- MIT Social and Ethical Responsibilities of Computing (SERC) Seed Grant Review Committee 2026
- MIT Sloan OR/Statistics Group Hiring Committee 2025
- MIT ORC Master's Program Admissions Committee 2025
- MIT Sloan MBAn Program Admissions Reader 2024, 2025
- Chicago Booth OM Faculty Recruiting Committee Member 2022-2023
- Chicago Booth OM Workshop Co-Organizer 2021-2022

Reviewer for Journals and Conferences:

Mathematical Programming, *SIAM Journal on Optimization*, *SIAM Journal on Control and Optimization*, *Mathematics of Operations Research*, *Mathematical Programming*

Computation, *INFORMS Journal on Computing*, *Foundations of Computational Mathematics*, *Operations Research*, *Management Science*, *Manufacturing & Service Operations Management*, *Production and Operations Management*, *Journal of Machine Learning Research*, *IEEE Transactions on Image Processing*, *Computational Optimization and Applications*, *INFORMS Journal on Optimization*, *Machine Learning*, *Journal of Optimization Theory and Applications*, *NeurIPS*, *WebConf*, *ICML*.

Professional Memberships:

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

Selected Presentations

2026 *(scheduled)* Duke University; Fuqua School of Business
(scheduled) Keynote Talk at Workshop on Optimisation, Metric Bounds, Approximation and Transversality
(scheduled) Keynote Talk at Australian Mathematical Society Conference 2026
(scheduled) Curtin University; Mathematics
(scheduled) RMIT University; Mathematics
University of Chicago; Booth School of Business
Stanford University; Information Systems Laboratory Colloquium
MIT; CSE Seminar:
INFORMS Annual Meeting
Invited speaker at AFOSR Mathematical Optimization Workshop
Zuse Institute in Berlin
Invited speaker at workshop on Foundations of Computational Mathematics
Invited speaker at Mostly OM workshop
Invited speaker at workshop on Foundation Models of the Electric Grid
Amazon; STOC team

2025 Plenary speaker at Cargese Workshop on Combinatorial Optimization
INFORMS Annual Meeting
Carnegie Mellon University; Tepper School of Business
NeurIPS Workshop on ScaleOPT: GPU-Accelerated and Scalable Optimization
Invited speaker at AFOSR Mathematical Optimization Workshop
International Conference on Continuous Optimization (ICCOPT)
MIP Workshop
Invited speaker at workshop on Strategies for Handling Applications with Nonconvexity
Invited speaker at workshop on Combinatorial Optimization and Data Science
Gurobi Inc

2024 NVIDIA; cuOpt team
Cornell University; Operations Research and Information Engineering
MIT; Operations Research Center
Harvard University; School of Engineering and Applied Sciences
INFORMS Annual Meeting
Invited speaker at AFOSR Mathematical Optimization Workshop
Google Research; Cambridge
Stanford University; Workshop for Yinyu Ye's Retirement Celebration
Northwestern University; Industrial Engineering and Management Sciences
University of Minnesota; Industrial and Systems Engineering
International Symposium on Mathematical Programming (ISMP)
INFORMS Conference on Optimization

2023 MIT; Operations Research Center
Columbia University; DRO and IEOR joint seminar
University of Southern California; Viterbi School of Engineering

	Chinese University of Hong Kong; School of Data Science
	INFORMS Annual Meeting
	SIAM Conference on Optimization
	Shanghai University of Finance and Economics
2022	INFORMS Annual Meeting
	Argonne National Laboratory
	Johns Hopkins University; Applied Mathematics Department
	Cornell University; Operations Research and Information Engineering
	MIT; Operations Research Center
	Tsinghua University; Institute for Interdisciplinary Information Sciences
	Simons Institute for the Theory of Computing
2021	University of Chicago; Booth School of Business
	Carnegie Mellon University; Tepper School of Business
	MIT; Institute for Data, Systems, and Society
	New York University; Stern School of Business
	University of Illinois at Chicago; Liautaud Graduate School of Business
2020	eBay
	Google Research; New York City
	International Conference on Artificial Intelligence and Statistics
	Chinese Academy of Sciences
	Shanghai Jiao Tong University; AI+Math Colloquia
	INFORMS Annual Meeting
	SIAM Conference on Mathematics of Data Science
2019	International Conference on Continuous Optimization (ICCOPT)
	Google Research; New York City
	Rensselaer Polytechnic Institute; Mathematics Department
	New York University; Courant Institute of Mathematical Sciences
	National University of Singapore; ISEM
	University of Toronto; Mathematics and Computer Science
	University of Toronto; MIE
	University of Chicago; Booth School of Business
	University of Minnesota, Twin Cities; ISE
	Duke University; Mathematics and Computer Science
	Columbia University; IEOR
	University of Wisconsin–Madison; ISE
	University of Illinois Urbana-Champaign; ISE
	Google Research; Cambridge
2018	INFORMS Annual Meeting
	Princeton University; Princeton Optimization Day
	Google Research; New York City
	Columbia University; Statistics Department
	International Symposium on Mathematical Programming (ISMP)
	International Conference on Machine Learning
	NYAS Meeting on Machine Learning
	INFORMS Conference on Optimization
2017	INFORMS Annual Meeting
	SIAM Conference on Optimization
2016	INFORMS Annual Meeting
	University of Edinburgh
2015	International Symposium on Mathematical Programming (ISMP)

Media and Blogs

Google research blog: Scaling up linear programming with PDLP (link).

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

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

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

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 (link).

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

Miscellaneous

Computing: Python, Julia, C++, R

Hobbies: Food/Cuisine, Kayaking, Sailing, Hiking, Skiing