

Haihao (Sean) Lu

CONTACT INFORMATION	Room 585, Sloan School of Management 100 Main St, Cambridge, MA 02142	Email: <a href="mailto:haihao@mit.edu">haihao@mit.edu</a> Phone: 857-998-3092
RESEARCH INTERESTS	Large-scale optimization, machine learning and data-driven decision making.	
APPOINTMENT	<b>Massachusetts Institute of Technology</b> , Cambridge, MA Sloan School of Management The Cecil and Ida Green Career Development Assistant Professor Assistant Professor of Operations Research/Statistics	July 2024-present
	<b>Google Research</b> , New York City, NY Research scientist (20% part time consulting)	December 2022-present
	<b>University of Chicago</b> , Chicago, IL Booth School of Business Assistant Professor of Operations Management	July 2020-June 2024
	<b>Google Research</b> , New York City, NY Visiting Faculty Researcher Large-Scale Optimization Team	July 2019-June 2020
EDUCATION	<b>Massachusetts Institute of Technology</b> , Cambridge, MA 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	2014-2019
	<b>Shanghai Jiao Tong University</b> , Shanghai, China B.S. in Applied Mathematics, Zhiyuan College, Graduation with distinction Advisor: Prof. David Cai	2010-2014
SELECTED HONORS	<b>COIN-OR Cup winner</b> ( <a href="#">link</a> ). 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.	2024
	<b>The Beale–Orchard-Hays Prize</b> ( <a href="#">link</a> ). This prize, awarded every three years by the Mathematical Optimization Society, recognizes excellence in computational mathematical programming.	2024
	<b>INFORMS Revenue Management and Pricing Section Prize</b> ( <a href="#">link</a> ). This prize is awarded for the best contribution to the science of pricing and revenue management published in English.	2023
	<b>Michael H. Rothkopf Junior Researcher Paper Prize (first place)</b> ( <a href="#">link</a> ). This prize recognizes the best paper by a young researcher in the area of Auctions and Market Design.	2022
	<b>Special invited session on “spotlight beyond WINE”</b> ( <a href="#">link</a> ). 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.	2022
	<b>INFORMS Optimization Society Young Researcher Prize</b> ( <a href="#">link</a> ).	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.

#### 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](#)).

#### JOURNAL PUBLICATIONS

"cuPDLP. jl: A GPU Implementation of Restarted Primal-Dual Hybrid Gradient for Linear Programming in Julia", with Jinwen Yang, to appear in *Operations Research*.

"On the Sparsity of Optimal Linear Decision Rules for a Class of Robust Optimization Problems with Box Uncertainty Sets", with Bradley Sturt, to appear in *Operations Research*.

"A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming", with Jinwen Yang, to appear in *Mathematical Programming*.

"Optimizing Scalable Targeted Marketing Policies with Constraints", with Duncan Simester and Yuting Zhu, to appear in *Marketing Science*.

"Regularized Online Allocation Problems: Fairness and Beyond", with Santiago Balseiro and Vahab Mirrokni, *Manufacturing & Service Operations Management*, 27.3 (2025):720-735.

"Auto-bidding and Auctions in Online Advertising: A Survey", with many people from Google, *ACM SIGecom Exchanges* 22.1 (2024): 159-183.

"On the Geometry and Refined Rate of Primal-Dual Hybrid Gradient for Linear Programming", with Jinwen Yang, to appear at *Mathematical Programming*.

"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.

"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.

"A J-Symmetric Quasi-Newton Method for Minimax Problems", with Azam Asl and Jinwen Yang, *Mathematical Programming* 204.1 (2024): 207-254.

"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.

“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.

“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.

“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.

“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.

“Randomized Gradient Boosting Machines”, Haihao Lu and Rahul Mazumder, *SIAM Journal on Optimization* 2020, 30(4), 2780-2808.

“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.

“ ‘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.

“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.

“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.

“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.

“Renormalized Dispersion Relations of  $\beta$ -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

“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.

PREPRINTS AND  
PAPERS UNDER  
REVIEWS

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

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

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

“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.

“MPAX: Mathematical Programming in JAX”, with Zedong Peng and Jinwen Yang.

“A New Crossover Algorithm for LP Inspired by the Spiral Dynamic of PDHG”, with Tianhao Liu.

“PDOT: a Practical Primal-Dual Algorithm and a GPU-Based Solver for Optimal Transport”, with Jinwen Yang.

“Restarted Halpern PDHG for Linear Programming”, with Jinwen Yang.

“Achieving Fairness and Accuracy in Regressive Property Taxation”, with Feiyu Han and Ozan Candogan.

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

“On the Infimal Sub-differential Size of Primal-Dual Hybrid Gradient Method”, with Jinwen Yang.

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

“Analysis of Dual-Based PID Controllers through Convolutional Mirror Descent”, with Santiago Balseiro, Vahab Mirrokni and Balasubramanian Sivan.

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

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

## TECHNICAL REPORTS

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

## TEACHING EXPERIENCE

### MIT:<sup>1</sup>

#### *Instructor*

- 15.071. The Analytics edge
  - 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:<sup>2</sup>

#### *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

<sup>1</sup>Two evaluation questions: “Overall rating” (out of 7), “Recommend Professor” (out of 5).

<sup>2</sup>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, 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.”

## PRESENTATIONS

### *GPU-Accelerated Linear Programming and Beyond*

- International Conference on Continuous Optimization (ICCOPT), July 2025
- MIP Workshop, June 2025
- Workshop on Strategies for Handling Applications with Nonconvexity, May 2025
- Workshop on Combinatorial Optimization and Data Science, May 2025
- Gurobi, January 2025
- NVIDIA, December 2024
- ORIE Department, Cornell University, November 2024
- Operations Research Center, MIT, October 2024
- School of Engineering and Applied Sciences, Harvard University, October 2024
- INFORMS annual meeting, October 2024
- AFOSR annual meeting, Virginia, September 2024
- Google Research, Cambridge, August 2024
- Workshop for Yinyu Ye's Retirement Celebration, Stanford University, July 2024
- Industrial Engineering and Management Sciences, Northwestern University, May 2024
- Industrial and Systems Engineering, University of Minnesota, May 2024

### *A Practical and Optimal First-Order Method for Large-Scale Convex Quadratic Programming*

- International Symposium on Mathematical Programming, Montreal, July 2024
- INOFMRS Conference on Optimization, Houston, March 2024

### *First-Order Methods for Constrained Continuous Optimization*

- Operations Research Center, MIT, December 2023
- DRO and IEOR joint seminar, Columbia University, November 2023
- Viterbi School of Engineering, University of Southern California, November 2023
- School of Data Science, Chinese University of Hong Kong, October 2023

### *First-Order Methods for Linear Programming: Theory, Computation and Applications*

- INFORMS Annual Meeting, October 2023
- SIAM Conference on Optimization, May 2023
- Shanghai University of Finance and Economics, March 2023
- INFORMS Annual Conference, October 2022
- Argonne National Lab, IL, September 2022
- Applied Mathematics Department, John Hopkins University, September 2022
- School of Operations Research and Information Engineering, Cornell University, August 2022
- Operations Research Center, MIT, IL, March 2022
- Institute for Interdisciplinary Information Sciences (Yao's class), Tsinghua University, February 2022
- Booth School of Business, University of Chicago, November 2021
- Tepper School of Business, Carnegie Mellon University, September 2021

### *The Best of Many Worlds: Dual Mirror Descent for Online Allocation Problems*

- INFORMS Annual Conference, IN, October 2022
- Simons Institute for the Theory of Computing, October 2022
- Institute for Data, Systems, and Society, MIT, December 2021
- Stern School of Business, New York University, October 2021
- Liautaud Graduate School of Business, University of Illinois at Chicago, March 2021
- Ebay Inc, August, 2020
- Google Research NYC, June 2020

### *Accelerated Gradient Boosting Machines*

- International Conference on Artificial Intelligence and Statistics, August 2020

*An  $O(s^r)$ -Resolution ODE Framework for Discrete-Time Optimization Algorithms and Applications to Minimax Problems*

- Chinese Academy of Sciences, December 2020
- AI+Math Colloquia, Shanghai Jiao Tong University, November 2020
- INFORMS Annual Meeting, November 2020
- Google Research NYC, June 2020
- SIAM Conference on Mathematics of Data Science, June 2020

*Ordered-SGD: A New Stochastic Optimization Framework for Empirical Risk Minimization*

- International Conference on Artificial Intelligence and Statistics, August 2020
- International Conference on Continuous Optimization, Berlin, August 2019
- Google Research, New York City, Sep 2019
- Rensselaer Polytechnic Institute, Mathematics Department, Nov 2019
- NYU, Courant Institute of Mathematical Science, Nov 2019

*Gradient Boosting Machines: New Insights, Algorithms, and Improved Complexity*

- National University of Singapore, ISEM, January 2019
- University of Toronto, Mathematics and Computer Science, January 2019
- University of Toronto, MIE, January 2019
- University of Chicago, Booth, January 2019
- University of Minnesota, Twin Cities, ISE, January 2019
- Duke University, Mathematics and Computer Science, January 2019
- Columbia University, IEOR, February 2019
- University of Wisconsin, Madison, ISE, February 2019
- University of Illinois, Urbana-Champaign, ISE, February 2019
- Google Research, Cambridge, February 2019
- Google Research, New York City, March 2019

*Randomized Gradient Boosting Machines*

- INFORMS Annual Meeting, Phoenix, November 2018

*Scalable Linear Programming via First-Order Methods*

- Princeton Optimization Day (Poster), Princeton, September 2018
- Google Research, New York City, August 2018

*Generalized Stochastic Frank-Wolfe Algorithm with Stochastic ‘Substitute’ Gradient for Structured Convex Optimization*

- Columbia University, Statistics Department, August 2018
- International Symposium on Mathematical Programming (ISMP), Bordeaux, July 2018

*Accelerating Greedy Coordinate Descent Methods*

- Google Research, New York City, July 2018
- International Conference on Machine Learning, Stockholm, July 2018
- NYAS Meeting on Machine Learning (Poster), New York City, March 2018

*Approximate Leave-One-Out for Fast Parameter Tuning in High Dimensions*

- International Conference on Machine Learning (Poster), Stockholm, July 2018
- NYAS Meeting on Machine Learning (Poster), New York City, March 2018

*“Relative-Continuity” for Non-Lipschitz Non-Smooth Convex Optimization using Stochastic (or Deterministic) Mirror Descent*

- INFORMS Meeting on Optimization, Denver, March 2018

- INFORMS Annual Meeting, Houston, Oct 2017

*Relatively-Smooth Convex Optimization by First-Order Methods, and Applications*

- SIAM Conference on Optimization, Vancouver, May 2017
- INFORMS Annual Meeting, Nashville, November 2016

*Extending the Scope of ‘Smooth’ and ‘Non-Smooth’ Convex Optimization via First-Order Methods*

- University of Edinburgh, Edinburgh, UK, April 2016

*Some New Results for Randomized Coordinate Gradient Descent*

- International Symposium on Mathematical Programming (ISMP), Pittsburg, July 2015

#### ACADEMIC SERVICE

##### **Community Service and Conference/Workshop Organizer**

- Vice chair of linear and conic optimization for INFORMS Optimization Society (2023-now)
- Program committee for INFORMS Optimization Society Conference (2024)
- Organizer of the workshop on modern continuous optimization at MIT (2023)
- Judge: George Nicholson student paper competition (2023, 2024)

##### **University Committee Service**

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

#### MISCELLANEOUS

**Computing:** Python, Julia, C++, R

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