

CURRICULUM VITAE

Name: RAHUL MAZUMDER

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

I Education

<i>Degree</i>	<i>School</i>	<i>Date</i>
PhD, Statistics	Stanford University, Stanford, CA	2012
M.Stat.	Indian Statistical Institute, Kolkata, India	2007
B.Stat.	Indian Statistical Institute, Kolkata, India	2005

II Title of Doctoral Thesis and Name of Thesis Advisor

Title: Topics in Sparse Multivariate Statistics

Advisor: Trevor Hastie, Stanford University

III Principal Fields of Interest

- Statistics, Machine Learning, Mathematical Optimization (Convex optimization, Mixed Integer optimization), Large Scale Optimization Algorithms
- High dimensional statistics and sparsity, combinatorial statistical modeling & computation, decision trees & ensembles, nonparametric function estimation (e.g., shape constrained inference). Conditional computing & sparsity in neural networks. Sparse mixture of experts. Post-training compression of foundation models.
- Applications of the above in recommender systems, biomedical sciences & healthcare, survey research, computational finance, etc.

IV Name and Rank of Other Sloan Faculty in Same Field

<i>Name</i>	<i>Rank</i>
Arnold Barnett	Professor
Cynthia Barnhart	Professor
Dimitris Bertsimas	Professor
David Gamarnik	Professor
Swati Gupta	Associate Professor (without tenure)
Alexandre Jacquillat	Associate Professor (without tenure)
Haihao Lu	Assistant Professor
Thomas Magnanti	Institute Professor
James 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>
Columbia University New York	Tenure Track Assistant Professor Department of Statistics	07/2013	06/2015
Yahoo! Research Santa Clara, CA	Research Intern Machine Learning & Statistics Group	07/2010	09/2010
Yahoo! Research Santa Clara, CA	Research Intern Machine Learning & Statistics Group	07/2009	09/2009

VI History of MIT Appointments

<i>School</i>	<i>Position</i>	<i>Start</i>	<i>End</i>
Laboratory for Information & Decision Systems	Affiliate	11/2024	Present
Sloan School of Management	Nanyang Technological University Associate Professor of Operations Research and Statistics	08/2023	Present
Sloan School of Management	Associate Professor with tenure (OR & Stats group)	07/2022	Present
MIT IBM Watson AI Lab	Principal Investigator	09/2019	Present
Center for Statistics & Data Science	Core Faculty	07/2017	Present
Operations Research Center	Affiliate	07/2015	Present
Sloan School of Management	Robert G. James Career Development Associate Professor	07/2019	06/2022
Sloan School of Management	Associate Professor, without tenure (OR & Stats group)	07/2019	06/2022
Sloan School of Management	Assistant Professor (OR & Stats group)	09/2015	06/2019
Operations Research Center	Postdoctoral Associate	10/2012	06/2013

VII MIT Activities

<i>Committee</i>	<i>Start</i>	<i>End</i>
MIT Sloan OR & Stats group head	2025	Present
MIT Sloan OR & Stats faculty hiring committee chair	2025	Present
Committee on Academic Performance (Undergraduate studies)	2024	Present
PhD General Exam Committee Member, MIT Operations Research Center	2016	Present
PhD Thesis Committee Member, MIT Operations Research Center	2016	Present
MIT Sloan Associate with Tenure Review Subcommittee (Dr. Negin Golrezaei)	2024	2025
MIT Sloan OR & Stats faculty hiring committee	2017	2025
MIT Sloan First Reappointment Committee Member (Dr. Johan Chu)	2024	2024
MIT Schwarzman College of Computing & Architecture shared faculty search committee	2023	2024
Dean's Social and Ethical Responsibilities of Computing (SERC) Action Group for Teaching	2021	2022
Admissions Committee, MIT Sloan Master of Business Analytics	2017	2022
MIT Sloan First Reappointment Committee Member (Dr. Jackson Lu)	2020	2020
Committee Member, Probability Qualifiers Examination, MIT Operations Research Center	2016	2020
Best Student Paper Award Committee, MIT Operations Research Center	2019	2019
MIT Operations Research Center Seminar Series Faculty Coordinator	2018	2018
Coordinator, Probability Qualifiers Examination, MIT Operations Research Center	2016	2018
Graduate Admissions Committee, MIT Operations Research Center	2016	2016
Chair, Best Student Paper Award Committee, MIT Operations Research Center	2016	2016

VIII Governmental Committees and Service

None

IX Consulting Activities

- Academic Scholar (part-time), LinkedIn (AI Foundations), USA, May 31, 2023 - current
- Sr. Staff Research Scientist (part-time) at LinkedIn (AI Foundations), USA, Nov 1, 2021 - April 30, 2023.

X Other Activities

None

XI Awards, Honors

SELECTED AWARDS & HONORS

- Leo Breiman Junior Award, 2024
- International Indian Statistical Association Early Career Award in Statistics and Data Sciences, 2024
- INFORMS Computing Society Prize (Honorable Mention), 2023
- Donald P. Gaver, Jr. Early Career Award for Excellence in Operations Research, 2021
- INFORMS Optimization Society Young Researchers Award, 2020
- Office of Naval Research Young Investigator Award (ONR YIP), 2018
- Annals of Statistics *Special Invited Session* Speaker at the Joint Statistical Meetings (JSM), 2017

PAPER AWARDS TO STUDENT CO-AUTHORS

- | | |
|---|------|
| • American Statistical Association (ASA) Best Student Paper Award Finalist in Nonparametric Statistics Section (as advisor) | 2026 |
| • American Statistical Association (ASA) Best Student Paper Award in Statistical Computing or Graphics (as advisor) | 2025 |
| • INFORMS Data Mining Best Student Paper Award (as advisor) | 2024 |
| • Conference of Artificial Intelligence and Statistics (AISTATS), Student Paper Highlight Award (as advisor) | 2024 |
| • ASA Section on Statistics in Genetics and Genomics Student Paper Award (as co-author) | 2024 |
| • 28th SIGKDD Conference on Knowledge Discovery and Data Mining, Best Student Paper Award (as advisor) | 2022 |
| • Mixed Integer Programming Workshop Student Poster Award (most popular poster, as advisor) | 2021 |
| • INFORMS Computing Society Student Paper Prize (honorable mention, as advisor) | 2020 |
| • MIT Operations Research Center Best Student Paper Award (as advisor) | 2020 |
| • Mixed Integer Programming Workshop Student Poster Award (honorable mention, as advisor) | 2019 |
| • INFORMS Optimization Society Student Paper Prize (as co-author) | 2015 |

XII Professional Membership and Activities

Editorial Boards of Book Series

<i>Role</i>	<i>Book Series</i>	<i>Start</i>	<i>End</i>
Series Editor (Algorithms)	Cambridge University Press (Institute of Mathematical Statistics, TextBooks and Monographs)	1/2024	Present

Editorial Boards of Journals

<i>Role</i>	<i>Journal</i>	<i>Start</i>	<i>End</i>
Action Editor	Journal of Machine Learning Research	1/2024	Present
Associate Editor	Operations Research (Optimization)	1/2024	Present
Associate Editor	Operations Research (Machine Learning/Data Science)	1/2024	Present
Editorial Board Member	Journal of Machine Learning Research	6/2020	Present
Associate Editor	Annals of Statistics	4/2020	Present
Associate Editor	Bernoulli	1/2022	12/2024

Services as a Reviewer and Panelist

- Reviewer for the following journals and conferences:

Artificial Intelligence and Statistics (AISTATS), Annals of Applied Statistics, Annals of Statistics, Biometrika, Computational Statistics and Data Analysis, Journal of the American Statistical Association, Journal of Machine Learning Research, Journal of the Royal Statistical Society (Series B), Journal of the Royal Statistical Society (Series C), Journal of Computational and Graphical Statistics, Journal of Optimization Theory and Applications, IEEE Transactions on Signal Processing, IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Information Theory, Information and Inference, INFORMS Journal on Computing, INFORMS Journal on Optimization, Knowledge Discovery and Data Mining Conference (KDD), Management Science, Mathematical Programming, Neural and Information Processing Systems (NeurIPS), Operations Research, Operations Research Letters, Optimization Methods and Software, SIAM Journal on Optimization, Statistical Science

- Panelist on the NSF panels:

<i>Division</i>	<i>Date</i>
NSF Operations Engineering, CMMI	2025
Division of Mathematical Sciences	2025
Division of Computing and Communication Foundations	2018
Division of Information and Intelligent Systems	2017

- Ad-hoc reviewer of grant proposals for:

<i>Agency</i>	<i>Date</i>
Office of Naval Research	2020, 2022, 2023
US Department of Energy	2020
Army Research Office	2018
NSF (Computing and Communication Foundations)	2018
Israel Science Foundation	2015, 2016
National Security Agency-American Mathematical Society	2015

Other professional activities

- Professional Activities as a Visiting Fellow/Scholar *Date*
 - Visiting Fellow at “Initiative on Data Science in Business and Economics”
at the University of Chicago, Booth School of Business 2018
 - Summer at Census Scholar, United States Census Bureau 2018
- Membership in Professional Societies
ASA (American Statistical Association), ACM (Association of Computing Machinery), IMS (Institute of Mathematical Statistics), IISA (International Indian Statistical Association), INFORMS (Institute for Operations Research and the Management Sciences), SIAM (Society for Industrial and Applied Mathematics)
- Conference, Tutorial, Workshop Organization and Committee Member *Date*
 - Co-organizer (with Behdin, Dai and other LinkedIn engineers) of the Web Conference (WWW ’25) Tutorial: “Efficient Algorithms for Leveraging LLMs for Generative and Predictive Recommender Systems” 2025
 - Founding Cluster Chair: ML Cluster within INFORMS Optimization Society 2025
 - Co-organizer (with Guillaume Pouliot, Alfred Galichon, Roger Koenker) of “Optimization-Conscious Econometrics Conference”-II and III at Becker Friedman Institute for Economics at the University of Chicago 2023–24
 - Co-organizer (with K. Selvaraj, A. Gupta, and other LinkedIn engineers) of KDD 2023 Tutorial: “Practical design of performant recommender systems using large-scale linear programming-based global inference” 2023
 - Committee Member of INFORMS Optimization Society Young Researchers Prize 2022
 - Program Committee Member (with Damek Davis): “Optimization for Data Science Track” at the INFORMS Optimization Society Conference 2021
 - Founder of Special Interest Group (SIG): ML SIG within INFORMS Optimization Society 2021
 - Co-organizer (with Roger Koenker and Guillaume Pouliot) of “Optimization-Conscious Econometrics Conference” at Becker Friedman Institute for Economics at the University of Chicago 2021
 - Co-organizer of “Mixed Integer Programming” Workshop at MIT, Cambridge 2019

- | | <i>Date</i> |
|--|-------------|
| • Session Chair and co-organizer | |
| – Session Chair/co-organizer INFORMS Annual Conference, USA | 2019-25 |
| – Session co-organizer Computational and Methodological Statistics, UK | 2017 |
| – Session Chair/co-organizer, IISA Annual Conference, USA | 2016 |
| – Session Chair/co-organizer at INFORMS Annual Conference, USA | 2016 |

XIII Subjects Taught

Teaching at MIT

<i>Number</i>	<i>Title</i>	<i>Date</i>
15.087	Engineering Statistics & Data Science (LGO Core)	Summer 2023, 2024, 2025
15.072	Advanced Analytics Edge (MBAn Core)	Fall 2024, 2025
15.086	Engineering Probability (LGO Core)	Summer 2024, 2025
15.060	Data Models and Decisions (MBA Core)	Fall 2016, 2017, 2019–2022
15.S15	Readings in Statistics (PhD course, co-taught with Colin Fogarty)	Fall 2016
15.075	Statistical Thinking and Data Analysis (Undergraduate Course)	Spring 2016, 2024
15.097	Statistical Learning via a Modern Optimization Lens (PhD course, co-taught with Dimitris Bertsimas)	Spring 2016

Teaching prior to MIT

- Teaching at Columbia University
 - W4240 “Data Mining”, Spring 2015 (MS Class)
 - W4105 “Introduction to Probability”, Fall 2014 (MS Class)
 - W4107 “Introduction to Statistical Inference”, Fall 2013 (MS Class)
 - G8325 “Topics in Modern Multivariate Statistics: Modeling, Computation and Theory” (Advanced Topics Course in Statistics, PhD Class)
 - S4204 Data Mining, Summer 2014 (MS Class)
- Teaching at Stanford University as a Teaching Assistant
 - STATS 60: Spring 2011, Undergraduate Course in Introductory Statistics (pre-calculus).
 - STATS 110: Spring 2011, Introductory Statistics course for Engineers (post-calculus).
 - STATS 306B: Spring 2009, PhD first year course in Applied Statistics

- STATS 300A: Fall 2008, PhD first year course in Theoretical Statistics
- STATS 217/218/219: Winter 2008, Spring 2008, Summer 2008, Fall 2009, Winter 2009, Summer 2011, Graduate course in Stochastic Processes
- STATS 305: Fall 2010, PhD first year course in Applied Statistics
- STATS 315A: Winter 2011, Modern Applied Statistics

XIV Theses and Student Supervision

1. DOCTORAL THESES SUPERVISED

- (i) Shibal Ibrahim, MIT EECS, 2019–2024
PhD in EECS
Title: *Nonparametric High-dimensional Models: Sparsity, Efficiency, Interpretability*
Next Position: Software Engineer, Google, CA
- (ii) Haoyue Wang, MIT Operations Research Center, 2019–2024
PhD in Operations Research
Title: *Large-scale Algorithms for Machine Learning: Efficiency, Estimation Errors, and Beyond*
Next Position: Susquehanna International Group, PA
- (iii) Kayhan Behdin, MIT Operations Research Center, 2019–2024
PhD in Operations Research
Title: *Statistical Learning with Discrete Structures: Statistical and Computational Perspectives*
Next Position: Machine Learning Engineer/Researcher, LinkedIn, CA
- (iv) Wenyu Chen, MIT Operations Research Center, 2018–2023
PhD in Operations Research
Title: *Optimization Methods for Machine Learning under Structural Constraints*
Next Position: Quantitative Researcher, Citadel, NY
- (v) Hussein Hazimeh, MIT Operations Research Center, 2016–2021
PhD in Operations Research
Title: *Sparse Learning using Discrete Optimization: Scalable Algorithms and Statistical Insights*
Next Position: Research Scientist, Google Research, NY.

2. MASTER'S THESES SUPERVISED

- (i) Riade Benbaki, MIT Operations Research Center, 2021–2023
MS in Operations Research
Title: *Topics in Sparsity and Compression: From High dimensional statistics to Overparametrized Neural Networks*
Next Position: Jump Trading, UK
- (ii) Antoine Dedieu, MIT Operations Research Center, 2016–2018
MS in Operations Research
Title: *Sparse learning: Statistical and optimization perspectives*
Next Position: Researcher at Vicarious AI, CA.
- (iii) Zachary Blanks, MIT Operations Research Center, 2017– 2019
MS in Operations Research (co-advisor Dr. Troy Lau, DRAPER)

Title: *A Generalized Hierarchical Approach for Data Labeling*

Next Position: PhD Candidate at UVA School of Data Science

(v) Suyash Gupta, Indian Statistical Institute, 2017

MStat in Statistics

Title: *Sentiment Analysis*

(Co-supervisor: Indranil Mukhopadhyay, Indian Statistical Institute)

Next Position: PhD candidate, Stanford University.

3. THESES IN PROGRESS

- Currently serving as the thesis advisor for the following students:

<i>Name</i>	<i>Program</i>	<i>Degree</i>	<i>Year</i>
Adam Deng	MIT Operations Research Center	PhD	2024–
Charlie Liu (Co-Advisor Haihao Lu)	MIT Operations Research Center	PhD	2024–
Ryan Lucas	MIT Operations Research Center	PhD	2023–
Mehdi Makni	MIT Operations Research Center	PhD	2023–
Peter Prastakos	MIT Operations Research Center	PhD	2023–
Gabriel Afriat	MIT Operations Research Center	PhD	2022–
Brian Liu	MIT Operations Research Center	PhD	2021–
Xiang Meng	MIT Operations Research Center	PhD	2021–

XV Publications

1. Theses

Phd: “Topics in Sparse Multivariate Statistics”, Stanford University, Dept. of Statistics, 2012

Advisor: Dr. Trevor J. Hastie

Committee: Emmanuel Candes, Jerome Friedman, Rob Tibshirani and Michael Saunders (external).

Masters: “Local scale-space contrasts via Gaussian mixture ensembles for speech signal segmentation”.

Indian Statistical Institute, 2007.

Advisor: Dr. Debapriya Sengupta.

(awarded the *TCS award for best Masters thesis*, Indian Statistical Institute, 2005-2007.)

2. Refereed Journal Articles

Note: Different author ordering conventions are used based on area-specific norms and expectations:

- (i) Papers where all authors are alphabetically ordered are indicated with the footnote^{††}
- (ii) Student co-authors are denoted by * (asterisk)
- (iii) When a subset of authors have equal contributions, the authors are indicated by the footnote[†]. In all other cases, the first author is the primary contributor
- (iv) In papers where all co-authors are students, I am the senior author. In other papers where I am a senior author, I indicate it by the superscript ^S next to my name.

Methodological Publications

- [J1] “*Multi-Task Learning for Sparsity Pattern Heterogeneity: Statistical and Computational Perspectives*”
Kayhan Behdin*, Gabriel Loewinger*, Kenneth T. Kishida, Giovanni Parmigiani, Rahul Mazumder^S, 2025
Journal of the Royal Statistical Society, Series B
- [J2] “*Sparse PCA: A New Scalable Estimator Based On Integer Programming*”
Kayhan Behdin* and Rahul Mazumder, 2025+
Annals of Statistics (forthcoming)
- [J3] “*Randomization Can Reduce Both Bias and Variance: A Case Study in Random Forests*”
Brian Liu* and Rahul Mazumder, 2025
Journal of Machine Learning Research
- [J4] “*Sparse Gaussian Graphical Models with Discrete Optimization: Computational and Statistical Perspectives*”
Kayhan Behdin*, Wenyu Chen*, and Rahul Mazumder, 2025+
Operations Research (accepted pending a minor revision)
- [J5] “*Nonparametric Finite Mixture Models with Possible Shape Constraints: A Cubic Newton Approach*”
Haoyue Wang*, Shibal Ibrahim*, and Rahul Mazumder, 2025
SIAM Journal on Mathematics of Data Science
- [J6] “*Predicting Census Survey Response Rates with Parsimonious Additive Models and Structured Interactions*”
Shibal Ibrahim*, Peter Radchenko, Emanuel Ben-David, and Rahul Mazumder^S, 2025
Annals of Applied Statistics
- [J7] “*PolyCD: Optimization via Cycling through the Vertices of a Polytope*”
^{††}Rahul Mazumder and Haoyue Wang*, 2024
SIAM Journal on Optimization
- [J8] “*Sparse NMF with Archetypal Regularization: Computational and Robustness Properties*”
Kayhan Behdin*, and Rahul Mazumder, 2024
Journal of Machine Learning Research
- [J9] “*Subgradient Regularized Multivariate Convex Regression at Scale*”
Wenyu Chen* and Rahul Mazumder, 2024

^{††}All authors in this paper are arranged alphabetically

[†]These authors contributed equally.

SIAM Journal on Optimization

- [J10] “*L0Learn: A Scalable Package for Sparse Learning using L0 Regularization*”

††Hussein Hazimeh*, Rahul Mazumder and Tim Nonet*, 2023

Journal of Machine Learning Research

- [J11] “*A New Computational Framework for Log-Concave Density Estimation*”

††Wenyu Chen*, Rahul Mazumder and Richard J. Samworth, 2023

Mathematical Programming Computation

- [J12] “*Grouped Variable Selection with Discrete Optimization: Computational and Statistical Perspectives*”

††Hussein Hazimeh*, Rahul Mazumder and Peter Radchenko, 2023

Annals of Statistics

- [J13] “*Subset Selection with Shrinkage: Sparse Linear Modeling when the SNR is low*”

Rahul Mazumder[†], Peter Radchenko[†] and Antoine Dedieu*, 2023

Operations Research

- [J14] “*Linear regression with partially mismatched data: local search with theoretical guarantees*”

††Rahul Mazumder and Haoyue Wang*, 2023

Mathematical Programming

An extended abstract of this paper appeared in IPCO

- [J15] “*Frank-Wolfe Methods with an Unbounded Feasible Region and Applications to Structured Learning*”

Haoyue Wang*, Haihao Lu and Rahul Mazumder^S, 2022

SIAM Journal on Optimization

- [J16] “*Solving L1-regularized SVMs and related linear programs: Revisiting the effectiveness of Column and Constraint Generation*”

††Antoine Dedieu*, Rahul Mazumder and Haoyue Wang*, 2022

Journal of Machine Learning Research

- [J17] “*Sparse Regression at Scale: Branch-and-Bound rooted in First-Order Optimization*”

††Hussein Hazimeh*, Rahul Mazumder and Ali Saab*, 2022

Mathematical Programming

► MIT Operations Research Center Best Student Paper Award, 2020. (Awardee: Hazimeh)

► INFORMS Computing Society Student Paper Prize (honorable mention), 2020. (Awardee: Hazimeh)

- [J18] “*Using ℓ_1 -relaxation and integer programming to obtain dual bounds for sparse PCA*”

††Santanu Dey, Rahul Mazumder and Guanyi Wang*, 2021.

Operations Research

- [J19] “*Learning Sparse Classifiers: Continuous and Mixed Integer Optimization Perspectives*”

††Antoine Dedieu*, Hussein Hazimeh* and Rahul Mazumder, 2021

Journal of Machine Learning Research

- [J20] “*Randomized Gradient Boosting Machine*”

††Haihao Lu* and Rahul Mazumder, 2020

SIAM Journal on Optimization

- [J21] “*Fast Best Subset Selection: Coordinate Descent and Local Combinatorial Optimization Algorithms*”

††Hussein Hazimeh* and Rahul Mazumder, 2020

Operations Research

► Awarded the INFORMS Optimization Society Young Researchers Prize, 2020. (Awardees: Hazimeh and Mazumder)

- [J22] “Computing the degrees of freedom of rank-regularized estimators and cousins”
^{††}Rahul Mazumder and Haolei Weng*, 2020
Electronic Journal of Statistics
- [J23] “Matrix completion with nonconvex regularization: spectral operators and scalable algorithms”
^{††}Rahul Mazumder, Diego Saldana* and Haolei Weng*, 2020
Statistics and Computing
- [J24] “Computation of the Maximum Likelihood estimator in low-rank Factor Analysis”
^{††}Koulik Khamaru* and Rahul Mazumder, 2019
Mathematical Programming
- [J25] “A Computational Framework for Multivariate Convex Regression and its Variants”
 Rahul Mazumder, Arkopal Choudhury, Garud Iyengar and Bodhisattva Sen, 2019
Journal of the American Statistical Association, Theory and Methods
- [J26] “Learning a Mixture of Gaussians via Mixed Integer Optimization”
^{††}Hari Bandi*, Dimitris Bertsimas and Rahul Mazumder, 2019
Inform Journal on Optimization
- [J27] “Flexible low-rank statistical modeling with missing data and side information”
^{††}William Fithian and Rahul Mazumder, 2018
Statistical Science
- [J28] “Certifiably Optimal Low Rank Factor Analysis”
^{††}Dimitris Bertsimas, Martin Copenhaver* and Rahul Mazumder, 2017
Journal of Machine Learning Research
- [J29] “The Discrete Dantzig Selector: Estimating Sparse Linear Models via Mixed Integer Linear Optimization”
^{††}Rahul Mazumder and Peter Radchenko, 2017
IEEE Transactions on Information Theory
- [J30] “An Extended Frank-Wolfe Method with “In-Face” Directions, and its Application to Low-Rank Matrix Completion”
^{††}Robert Freund, Paul Grigas* and Rahul Mazumder, 2017
SIAM Journal on Optimization
- [J31] “A New Perspective on Boosting in Linear Regression via Subgradient Optimization and Relatives”
^{††}Robert Freund, Paul Grigas* and Rahul Mazumder, 2017
Annals of Statistics
 ► Editors’ choice for one of the best four papers accepted to the Annals of Statistics in the previous two years. Special Invited Session presentation by Mazumder at the Joint Statistical Meetings, 2017.
 ► INFORMS Optimization Society Student Paper Award, 2015. (Awardee Grigas)
- [J32] “Best Subset Selection via a Modern Optimization Lens”
^{††}Dimitris Bertsimas, Angela King* and Rahul Mazumder, 2016
Annals of Statistics
- [J33] “Matrix Completion and Low-Rank SVD via Fast Alternating Least Squares”
 Trevor Hastie[†], Rahul Mazumder[†], Jason Lee and Reza Zadeh, 2015
Journal of Machine Learning Research
- [J34] “Least Quantile of Squares Regression via Modern Optimization”

^{††}Dimitris Bertsimas and Rahul Mazumder, 2014

Annals of Statistics

[J35] “*The Graphical Lasso: New Insights and Alternatives*”

Rahul Mazumder and Trevor Hastie, 2012

Electronic Journal of Statistics

[J36] “*Exact covariance thresholding into connected components for large-scale Graphical Lasso*”

Rahul Mazumder and Trevor Hastie, 2012

Journal of Machine Learning Research

[J37] “*SparseNet: Coordinate Descent with Non-Convex Penalties*”

Rahul Mazumder, Jerome Friedman and Trevor Hastie, 2011

Journal of American Statistical Association, Theory and Methods

[J38] “*Spectral Regularization Algorithms for Learning Large Incomplete Matrices*”

Rahul Mazumder, Trevor Hastie and Robert Tibshirani, 2010

Journal of Machine Learning Research

Invited Discussion Articles

[J39] “*Discussion of “Best Subset, Forward Stepwise or Lasso? Analysis and Recommendations Based on Extensive Comparisons”*”

Rahul Mazumder, 2020

Statistical Science

Application-oriented Journal Publications

[J40] “*Fast and scalable ensemble learning method for versatile polygenic risk prediction*”

Tony Chen*, Haoyu Zhang, Rahul Mazumder, Xihong Lin, 2024

Proceedings of the National Academy of Sciences

[J41] “*Optimal ensemble construction for multistudy prediction with applications to mortality estimation.*”

Gabriel Loewinger*, Rolando Acosta Nunez*, Rahul Mazumder, and Giovanni Parmigiani, 2024

Statistics in Medicine

[J42] “*Accurate and Efficient Estimation of Local Heritability using Summary Statistics and LD Matrix*”

Hui Li*, Rahul Mazumder, Xihong Lin, 2023

Nature Communications

[J43] “*Analysis of correlations between local geographic atrophy growth rates and local OCT angiography-measured choriocapillaris flow deficits*”

Eric M Moul, Yingying Shi, Qinqin Zhang, Liang Wang, Rahul Mazumder, Siyu Chen, Zhongdi Chu, William Feuer, Nadia K Waheed, Giovanni Gregori, Ruikang K Wang, Philip J Rosenfeld, James G Fujimoto, 2021

Biomedical Optics Express

[J44] “*Integration of Survival Data from Multiple Studies*”

Steffen Ventz, Rahul Mazumder and Lorenzo Trippa, 2021

Biometrics

- [J45] “*Mining Events with Declassified Diplomatic Documents*”
Yuanjun Gao*, Jack Goetz*, Matthew Connelly and Rahul Mazumder^S, 2020
Annals of Applied Statistics
- [J46] “*Assessing the significance of global and local correlations under spatial autocorrelation: a nonparametric approach*”
Julia Validomat, Rahul Mazumder, Alex McInturff, Douglas McCauley and Trevor Hastie, 2014
Biometrics
- [J47] “*Turbulence, suspension and downstream fining over a sand-gravel mixture bed*”
Koeli Ghoshal, Rahul Mazumder, Chandan Chakraborty and Bijoy Mazumder, 2013
International Journal of Sediment Research
- [J48] “*Modeling Item-Item Similarities for Personalized Recommendations on Yahoo! Front Page*”
Deepak Agarwal, Liang Zhang and Rahul Mazumder, 2011
Annals of Applied Statistics
- [J49] “*Fluid flow pattern analysis in a trough region: a nonparametric approach*”
Rahul Mazumder, 2008
Journal of Applied Statistics
- [J50] “*Clustering based on geometry and interactions of turbulence bursting rate processes in a trough region*”
Rahul Mazumder, 2007
Environmetrics
- [J51] “*Statistical characterization of circulation patterns and direction of turbulent flow over a waveform structure*”
Rahul Mazumder and Bijoy Mazumder, 2006
Environmetrics

3. Articles in Refereed Conference Proceedings (& Workshops)

[Note: NeurIPS, ICML, ICLR, KDD, closely followed by AISTATS, are leading conferences in machine learning, data mining, AI. The conference proceedings are very competitive with low acceptance rates 25-27% (NeurIPS, ICML, AISTATS), 30% (ICLR), and 15% (KDD).]

- [C1] “*TSENor: Highly-Efficient Algorithm for Finding Transposable N:M Sparse Masks*”
Xiang Meng, Mehdi Makni, Rahul Mazumder, 2025
The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS ’25)
- [C2] “*3BASiL: An Algorithmic Framework for Sparse plus Low-Rank Compression of LLMs*”
Mehdi Makni, Xiang Meng, Rahul Mazumder, 2025
The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS ’25)
- [C3] “*Differentially Private High-dimensional Variable Selection via Integer Programming*”
Petros Prastakos, Kayhan Behdin, Rahul Mazumder, 2025
The Thirty-ninth Annual Conference on Neural Information Processing Systems (NeurIPS ’25)
- [C4] “*An Optimization Framework for Differentially Private Sparse Fine-Tuning*”
Mehdi Makni*, Kayhan Behdin*, Gabriel Afriat*, Zheng Xu, Sergei Vassilvitskii, Natalia Ponomareva, Hussein Hazimeh, and Rahul Mazumder^S, 2025
SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’25)

- [C5] “*MOSS: Multi-Objective Optimization for Stable Rule Sets*”
Brian Liu* and Rahul Mazumder, 2025
SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’25)
► *Recipient of INFORMS Data Mining Best Student Paper Award, 2024 (first place)*
- [C6] “*A unified framework for Sparse plus Low-Rank Matrix Decomposition for LLMs*” Mehdi Makni*, Kayhan Behdin*, Zheng Xu, Natalia Ponomareva, and Rahul Mazumder^S, 2025
The Second Conference on Parsimony and Learning (CPAL ’25)
► *Selected for Oral Presentation*
- [C7] “*Preserving Deep Representations In One-Shot Pruning: A Hessian-Free Second-Order Optimization Framework*”
Ryan Lucas* and Rahul Mazumder, 2025
International Conference on Learning Representations (ICLR ’25)
- [C8] “*ALPS: Improved Optimization for Highly Sparse One-Shot Pruning for Large Language Models*”
Xiang Meng*, Kayhan Behdin*, Haoyue Wang*, and Rahul Mazumder, 2024
Proceedings of the 38th Conference on Neural Information Processing Systems (NeurIPS ’24)
- [C9] “*FAST: An Optimization Framework for Fast Additive Segmentation in Transparent ML*”
Brian Liu* and Rahul Mazumder 2024
30th SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD ’24)
► *2025 American Statistical Association Statistical Computing Student Paper Competition Award Winner*
- [C10] “*OSSCAR: One-Shot Structured Pruning in Vision and Language Models with Combinatorial Optimization*”
Xiang Meng*, Shibal Ibrahim*, Kayhan Behdin*, Hussein Hazimeh, Natalia Ponomareva, and Rahul Mazumder^S, 2024
International Conference on Machine Learning (ICML ’24)
- [C11] “*End-to-end Feature Selection Approach for Learning Skinny Trees*”
Shibal Ibrahim*, Kayhan Behdin*, and Rahul Mazumder, 2024
International Conference on Artificial Intelligence and Statistics (AISTATS ’24)
► *Recipient of Student Paper Highlight Award*
► *Selected for Oral Presentation*
- [C12] “*FALCON: FLOP-Aware Combinatorial Optimization for Neural Network Pruning*”
Xiang Meng*, Wenyu Chen*, Riade Benbaki* and Rahul Mazumder
International Conference on Artificial Intelligence and Statistics (AISTATS ’24)
- [C13] “*Dyn-GWN: Time-Series Forecasting using Time-varying Graphs with Applications to Finance and Traffic Prediction*”
Shibal Ibrahim*, Max Tell* and Rahul Mazumder, 2023
4th ACM International Conference on AI in Finance (ICAIF 23)
- [C14] “*Dynamic Covariance Estimation under Structural Assumptions via a Joint Optimization Approach*”
Riade Benbaki*, Wenyu Chen*, Yada Zhu and Rahul Mazumder^S, 2023
4th ACM International Conference on AI in Finance (ICAIF 23)
- [C15] “*GRAND-SLAMIN’ Interpretable Additive Modeling with Structural Constraints*”
Shibal Ibrahim*, Gabriel Isaac Afriat*, Kayhan Behdin* and Rahul Mazumder, 2023
Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 23)

- [C16] “*On the Convergence of CART under Sufficient Impurity Decrease Condition*”
^{††}Rahul Mazumder and Haoyue Wang*, 2023
Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 23)
- [C17] “*Optimizing for Member Value in an Edge Building Marketplace*”
 Ayan Acharya, Siyuan Gao, Ankan Saha, Borja Ocejó, Kinjal Basu, Keerthi Selvaraj, Rahul Mazumder, Aman Gupta and Parag Agrawal, 2023
Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM 23)
- [C18] “*COMET: Learning Cardinality Constrained Mixture of Experts with Trees and Local Search*”
 Shibal Ibrahim*, Wenyu Chen*, Hussein Hazimeh, Natalia Ponomareva, Zhe Zhao, Rahul Mazumder^S, 2023
29th SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 23)
- [C19] “*FIRE: An Optimization Framework for Fast Interpretable Rule Extraction*”
^{††}Brian Liu* and Rahul Mazumder, 2023
29th SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 23)
- [C20] “*ForestPrune: Compact Depth-Pruned Tree Ensembles*”
^{††} Brian Liu* and Rahul Mazumder, 2023
International Conference on Artificial Intelligence and Statistics (AISTATS 23)
- [C21] “*Fast as CHITA: Neural Network Pruning with Combinatorial Optimization*”
 Riade Benbaki*, Wenyu Chen*, Xiang Meng*, Hussein Hazimeh, Natalia Ponomareva, Zhe Zhao, Rahul Mazumder^S, 2023
Proceedings of the 40th International Conference on Machine Learning (ICML 23)
- [C22] “*Improved Deep Neural Network Generalization Using m -Sharpness-Aware Minimization*”
 Kayhan Behdin, Qingquan Song, Aman Gupta, David Durfee, Ayan Acharya, Sathiya Keerthi, Rahul Mazumder^S, 2022
NeurIPS Optimization for Machine Learning Workshop (OPTML 22)
- [C23] “*A Light-speed Linear Program Solver for Personalized Recommendation with Diversity Constraints*”
 Haoyue Wang, Miao Cheng, Kinjal Basu, Aman Gupta, Keerthi Selvaraj, Rahul Mazumder^S, 2022
NeurIPS Optimization for Machine Learning Workshop (OPTML 22)
- [C24] “*Network Pruning at Scale: A Discrete Optimization Approach*”
 Wenyu Chen*, Riade Benbaki*, Xiang Meng*, Rahul Mazumder, 2022
NeurIPS Optimization for Machine Learning Workshop (OPTML 22)
- [C25] “*Pushing the limits of fairness impossibility: Who’s the fairest of them all?*”
 Brian Hsu, Rahul Mazumder, Preetam Nandy, Kinjal Basu, 2022
Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS 22)
- [C26] “*Flexible Modeling and Multitask Learning using Differentiable Tree Ensembles*”
 Shibal Ibrahim*, Hussein Hazimeh, and Rahul Mazumder^S, 2022
28th SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 22)
► KDD Best Student Paper Award 2022 (Awardee: Ibrahim).
- [C27] “*Quant-BnB: A Scalable Branch-and-Bound Method for Optimal Decision Trees with Continuous Features*”
^{††}Rahul Mazumder, Xiang Meng*, Haoyue Wang*, 2022
Proceedings of the 39th International Conference on Machine Learning (ICML 22)

- [C28] “Newer is not always better: Rethinking transferability metrics, their peculiarities, stability and performance”
Shibal Ibrahim*, Natalia Ponomareva and Rahul Mazumder^S, 2022
ECML PKDD 2022
Earlier version NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications
- [C29] “Knowledge Graph Guided Simultaneous Forecasting and Network Learning for Multivariate Financial Time Series”
Shibal Ibrahim*, Wenyu Chen*, Yada Zhu, Pin-Yu Chen, Yang Zhang, Rahul Mazumder^S, 2022
3rd ACM International Conference on AI in Finance (ICAIF 22)
Earlier version at KDD Workshop on Machine Learning in Finance 21
- [C30] “DSelect-k: Differentiable Selection in the Mixture of Experts with Applications to Multi-Task Learning”
Hussein Hazimeh, Zhe Zhao, Aakanksha Chowdhery, Maheswaran Sathiamoorthy, Yihua Chen, Rahul Mazumder, Lichan Hong and Ed H. Chi, 2021
Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 21)
- [C31] “Linear Regression with Mismatched Data: a Provably Optimal Local Search Algorithm”
†† Rahul Mazumder and Haoyue Wang*, 2021
Proceedings of the 22nd International Conference on Integer Programming and Combinatorial Optimization (IPCO 21)
- [C32] “ECLIPSE: An Extreme-Scale Linear Program Solver for Web-Applications”
†† Kinjal Basu, Amol Ghoting, Rahul Mazumder and Yao Pan, 2020
Proceedings of the 37th International Conference of Machine Learning (ICML 20)
- [C33] “The Tree Ensemble Layer: Differentiability meets Conditional Computation”
Hussein Hazimeh*, Natalia Ponomareva, Petros Mol, Zhenyu Tan and Rahul Mazumder^S, 2020
Proceedings of the 37th International Conference on Machine Learning (ICML 20)
- [C34] “Learning Hierarchical Interactions at Scale: A Convex Optimization Approach”
Hussein Hazimeh* and Rahul Mazumder, 2020
Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AIS-TATS 20)
► *Mixed Integer Programming Workshop student poster award (honorable mention), 2019. (Awardee: Hazimeh).*
- [C35] “Hierarchical Modeling and Shrinkage for User Session Length Prediction in Media Streaming”
Antoine Dedieu, Rahul Mazumder, Zhen Zhu and Hossein Vahabi, 2018
Proceedings of the 27th ACM International Conference on Information and Knowledge Management (CIKM 18).
- [C36] “Non-Negative Matrix Completion for Bandwidth Extension: A Convex Optimization Approach”
Dennis Sun and Rahul Mazumder, 2013
IEEE Workshop on Machine Learning for Signal Processing (MLSP)
- [C37] “Projected likelihood contrasts for testing homogeneity in finite mixture models with nuisance parameters”
Debapriya Sengupta and Rahul Mazumder, 2008
Beyond Parametrics in Interdisciplinary Research: Festschrift in Honor of Professor Pranab K. Sen (Beachwood, Ohio, USA: Institute of Mathematical Statistics), 272-281.

4. Articles in Non-Refereed Conference Proceedings

None

5. Selected Articles in Progress or Under Review

Articles with recently decisioned revision requests are listed below

- [R1] “*Extracting Interpretable Models from Tree Ensembles: Computational and Statistical Perspectives*”
^{††} Brian Liu*, Rahul Mazumder^S and Peter Radchenko 2025+
Journal of the American Statistical Association (minor revision)
- [R2] “*Univariate convex regression: ℓ_q risk bounds under heavy-tailed noise*”
^{††} Rahul Mazumder and Haoyue Wang* 2025+
Mathematics of Operations Research (major revision)
- [R3] “*Modeling with Categorical Features via Exact Fusion and Sparsity Regularization*”
 Kayhan Behdin*, Riade Benbaki*, Peter Radchenko, Rahul Mazumder,^S, 2025+
Journal of the Royal Statistical Society, Series B (major revision submitted)
- [R4] “*Locally Transparent Rule Sets for Explainable Machine Learning*”
 Brian Liu* and Rahul Mazumder, 2025+
Operations Research (Reject & Resubmit)
- [R5] “*Improved heritability partitioning and enrichment analyses using summary statistics with graphREML*”
 Hui Li*, Tushar Kamath*, Rahul Mazumder, Xihong Lin, and Luke O’Connor, 2025+
Nature Genetics (Revision submitted)
- [R6] “*SPLendid incorporates continuous genetic ancestry in biobank-scale data to improve polygenic risk prediction across diverse populations*”
 Tony Chen, Haoyu Zhang, Rahul Mazumder, and Xihong Lin, 2025+
Nature Methods (Revision submitted)

6. Other Publications: Software Products

- [S1] My research group open-sources software accompanying our research publications. Most recent software hosted on Lab GitHub repository <https://github.com/mazumder-lab>
- [S2] Some software packages are hosted on CRAN and some are hosted on pypi (Python). Examples of some R packages arising from my methodological research that I’ve co-developed: `softImpute`, `sparsenet`, `FACTMLE`, `LOLearn`.
- [S3] Download stats for three software packages I’ve co-developed: matrix completion algorithm `softImpute` (~162,000 downloads from RStudio), primal heuristic `LOLearn` for sparse regression (~90,000 downloads across Python/pypi, RStudio), branch-and-bound based algorithm `LOBnB` for sparse regression (~37,000 downloads from Python/pypi).
- [S4] My research methods have been implemented as software hosted on Gurobi’s website, Google’s official

github repository, and LinkedIn's official github repository* and are widely used.

*Gurobi: https://gurobi.github.io/modeling-examples/linear_regression/, Google: https://github.com/google-research/google-research/tree/master/tf_trees and https://github.com/google-research/google-research/tree/master/dselect_k_moe; and LinkedIn <https://github.com/linkedin/DuaLip>