

Madeleine Udell

Assistant Professor, Management Science & Engineering
475 Via Ortega, Huang Engineering Center, Stanford, CA 94305
udell@stanford.edu · 415-515-7872 · <https://web.stanford.edu/~udell>

Academic Employment

| | |
|--|--|
| Stanford University Assistant Professor Management Science and Engineering and, by courtesy, Electrical Engineering Gabilan Fellow Affiliated with the Institute of Computational and Mathematical Engineering, Stanford Data Science, Human Centered AI, Center for AI Safety, and BioX | Stanford, CA <i>July 2022 –</i> |
| Cornell University Associate Professor with tenure (on leave) Assistant Professor Department of Operations Research and Information Engineering Richard and Sybil Smith Sesquicentennial Fellow Graduate field member in Operations Research, Computer Science, Applied Mathematics, Data Science, Statistics, Electrical and Computer Engineering, and Systems Engineering | Ithaca, NY <i>July 2022 – June 2024</i> <i>July 2016 – June 2022</i> |
| California Institute of Technology Postdoctoral Fellow, Center for the Mathematics of Information | Pasadena, CA <i>June 2015 – June 2016</i> |

Education

| | |
|--|-----------------------------------|
| Stanford University Ph.D. in Computational and Mathematical Engineering GPA: 4.0. Thesis: <i>Generalized Low Rank Models</i> . Advisor: Stephen P. Boyd. | Stanford, CA <i>June 2015</i> |
| Yale University B.S. in Mathematics and Physics <i>Summa cum laude</i> , with honors in mathematics and honors in physics. GPA: 3.95. Thesis: <i>Local Parametrizations via Laplacian Eigenfunctions</i> . Advisor: Peter W. Jones. | New Haven, CT <i>June 2009</i> |

Papers

In the pipeline

- 13.
12. W. Gao, C. He, and M. Udell. Small gradient norm regret for online convex optimization. 2026, 2601.13519
11. S. Fazliani and M. Udell. Pde-sharp: Pde solver hybrids through analysis and refinement passes. 2025, 2511.00183
10. C. He, W. Gao, B. Jiang, M. Udell, and S. Zhang. New results on the polyak stepsize: Tight convergence analysis and universal function classes. 2025, 2512.06231

9. C. Lawless, Y. Li, A. Wikum, M. Udell, and E. Vitercik. LLMs for cold-start cutting plane separator configuration. 2025, 2412.12038
8. C. Lawless, J. Schöffer, and M. Udell. “It was a magical box”: Understanding practitioner workflows and needs in optimization. 2025, 2509.16402
7. Y.-C. Chu, W. Gao, Y. Ye, and M. Udell. Gradient methods with online scaling part 2. practical aspects. 2025, 2509.11007
6. W. Gao, Y.-C. Chu, Y. Ye, and M. Udell. Gradient methods with online scaling part i. theoretical foundations. 2025, 2505.23081
5. J. Sun, Z. Frangella, and M. Udell. Sapphire: Preconditioned stochastic variance reduction for faster large-scale statistical learning. *Major revision at SIAM Journal on Mathematics of Data Science (SIMODS)*, 2025, 2501.15941
4. L. Lessard and M. Udell. Algebraic characterization of equivalence between optimization algorithms. *Minor revision at Math Programming*, 2025, 2501.04972
3. W. Gao, H. Liu, Y. Ye, and M. Udell. When does primal interior point method beat primal-dual in linear optimization? *Submitted*, 2024, 2411.16015
2. P. Rathore, Z. Frangella, J. Yang, M. Dereziński, and M. Udell. Have ASkotch: A neat solution for large-scale kernel ridge regression. *Submitted*, 2025, 2407.10070
1. A. AhmadiTeshnizi, W. Gao, H. Brunborg, S. Talaei, C. Lawless, and M. Udell. OptiMUS-0.3: Using large language models to model and solve optimization problems at scale. *Submitted*, 2024, 2407.19633

Refereed Journal Articles

27. W. Gao, Z. Qu, M. Udell, and Y. Ye. Scalable approximate optimal diagonal preconditioning. *Computational Optimization and Applications*, 2026, 2312.15594
26. Y.-C. Chu, L.-R. Santos, and M. Udell. Randomized nyström preconditioned interior point-proximal method of multipliers. *SIAM Journal on Scientific Computing (SISC)*, 2026, 2404.14524
25. Z. Frangella, S. Zhao, T. Diamandis, B. Stellato, and M. Udell. On the (linear) convergence of generalized newton inexact ADMM. *Transactions on Machine Learning Research (TMLR)*, 2026, 2302.03863
24. T. Diamandis, Z. Frangella, S. Zhao, B. Stellato, and M. Udell. GeNIOS: an (almost) second-order operator-splitting solver for large-scale convex optimization. *Math Programming Computation*, 2026, 2310.08333
23. S. Fazliani, Z. Frangella, and M. Udell. Enhancing physics-informed neural networks through feature engineering. *Transactions on Machine Learning Research (TMLR)*, 2025, 2502.07209
22. P. Rathore, Z. Frangella, and M. Udell. PROMISE: Preconditioned stochastic optimization methods by incorporating scalable curvature estimates. *JMLR*, 2024, 2309.02014

21. Z. Frangella, P. Rathore, S. Zhao, and M. Udell. SketchySGD: Reliable stochastic optimization via robust curvature estimates. *SIAM Journal on Mathematics of Data Science (SIMODS)*, 2024, 2211.08597
20. Y. Zhao and M. Udell. gcompute: A package for missing data imputation. *Journal of Statistical Software*, 2024, 2203.05089
19. D. Kouri, Z. Hua, and M. Udell. A greedy galerkin method to efficiently select sensors for linear dynamical systems. *Linear Algebra and its Applications*, 2023
18. J. Fan, L. Ding, C. Yang, and M. Udell. Low-rank tensor recovery with Euclidean-norm-induced Schatten-p quasi-norm regularization. *Transactions on Machine Learning Research*, 2023, 2012.03436
17. Z. Frangella, J. A. Tropp, and M. Udell. Randomized Nyström preconditioning. *SIAM Journal on Matrix Analysis and Applications*, 2022, 2110.02820
16. L. Ding and M. Udell. A strict complementarity approach to error bound and sensitivity of solution of conic programs. *Optimization Letters*, 2022, 2012.00183
15. N. Sengupta, M. Udell, N. Srebro, and J. Evans. Sparse data reconstruction, missing value and multiple imputation through matrix factorization. *Sociological Methodology*, 2022
14. L. Ding and M. Udell. On the simplicity and conditioning of low rank semidefinite programs. *SIAM Journal on Optimization (SIOPT)*, 2021, 2002.10673
13. L. Ding, A. Yurtsever, V. Cevher, J. A. Tropp, and M. Udell. An optimal-storage approach to semidefinite programming using approximate complementarity. *SIAM Journal on Optimization (SIOPT)*, 2021, 1902.03373
12. J. Fan, C. Yang, and M. Udell. Robust non-linear matrix factorization for dictionary learning, denoising, and clustering. *IEEE Trans. Signal Processing (TSP)*, 69:1755–1770, 2021, 2005.01317
11. R. Muthukumar, D. P. Kouri, and M. Udell. Randomized sketching algorithms for low-memory dynamic optimization. *SIAM Journal on Optimization (SIOPT)*, 31(2):1242–1275, 2021
10. A. Yurtsever, J. A. Tropp, O. Fercoq, M. Udell, and V. Cevher. Scalable semidefinite programming. *SIAM Journal on Mathematics of Data Science (SIMODS)*, 3(1):171–200, 2021, 1912.02949
9. Y. Sun, Y. Guo, C. Luo, J. A. Tropp, and M. Udell. Low-rank Tucker approximation of a tensor from streaming data. *SIAM Journal on Mathematics of Data Science (SIMODS)*, 2(4):1123–1150, 2020, 1904.10951
8. J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Streaming low-rank matrix approximation with an application to scientific simulation. *SIAM Scientific Computing (SISC)*, 41(4):A2430–A2463, 2019, 1902.08651
7. M. Udell and O. Toole. Optimal design of efficient rooftop photovoltaic arrays. *INFORMS Journal on Applied Analytics (Interfaces)*, 49(4):281–294, 2019

6. M. Udell and A. Townsend. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science (SIMODS)*, 1(1):144–160, 2019, 1705.07474
5. N. Kallus and M. Udell. Dynamic assortment personalization in high dimensions. *Operations Research*, 2019, 1610.05604
4. J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Practical sketching algorithms for low-rank matrix approximation. *SIAM Journal of Matrix Analysis and Applications (SIMAX)*, 38(4):1454–1485, 2017, 1609.00048
3. M. Udell, C. Horn, R. Zadeh, and S. Boyd. Generalized low rank models. *Foundations and Trends in Machine Learning*, 9(1), 2016, 1410.0342
2. M. Udell and S. Boyd. Bounding duality gap for separable problems with linear constraints. *Computational Optimization and Applications*, 64(2):355–378, 2016, 1410.4158
1. E. Birch, M. Udell, and M. Covert. Incorporation of flexible objectives and time-linked simulation with flux balance analysis. *Journal of Theoretical Biology*, 345:12–21, 2014

Refereed Conference Proceedings

43. C. Lawless, T.-W. Weng, B. Ustun, and M. Udell. Understanding fixed predictions via confined regions. In *International Conference on Machine Learning (ICML)*, 2025, 2502.16380
42. P. Rathore, Z. Frangella, S. Garg, S. Fazliani, M. Dereziński, and M. Udell. Turbocharging gaussian process inference with approximate sketch-and-project. In *NeurIPS*, 2025, 2505.13723
41. Y.-C. Chu, W. Gao, Y. Ye, and M. Udell. Provable and practical online learning rate adaptation with hypergradient descent. In *International Conference on Machine Learning (ICML)*, 2025, 2502.11229
40. W. Gao, Y.-C. Chu, Y. Ye, and M. Udell. Gradient methods with online scaling. In *Conference on Learning Theory (COLT)*, 2025, 2411.01803
39. C. Lawless, Y. Li, A. Wikum, M. Udell, and E. Vitercik. LLMs for cold-start cutting plane separator configuration. In *International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR)*, 2025, 2412.12038
38. M. V. Ness and M. Udell. Interpretable prediction and feature selection for survival analysis. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2025, 2404.14689
37. M. Van Ness, B. Block, and M. Udell. DNAMite: interpretable calibrated survival analysis with discretized additive models. In *Machine Learning for Health (ML4H)*, 2024, 2411.05923
36. A. AhmadiTeshnizi, W. Gao, and M. Udell. OptiMUS: Scalable optimization modeling using MIP solvers and large language models. In *International Conference on Machine Learning (ICML)*, 2024, 2402.10172
35. P. Rathore, W. Lei, Z. Frangella, L. Lu, and M. Udell. Challenges in training PINNs: A loss landscape perspective. In *International Conference on Machine Learning (ICML)*, 2024, 2402.01868

34. M. Van Ness, T. Bosschieter, N. Din, A. Ambrosy, A. S. Singh, and M. Udell. Interpretable survival analysis for heart failure risk prediction. In *Machine Learning for Health (ML4H)*, 2023, 2310.15472
33. M. Van Ness, T. Bosschieter, R. Halpin-Gregorio, and M. Udell. The missing indicator method: From low to high dimensions. In *29th SIGKDD Conference on Knowledge Discovery and Data Mining - Applied Data Science Track*, 2023, 2211.09259
32. C.-H. Chang, J. Yoon, S. Arik, M. Udell, and T. Pfister. Data-efficient and interpretable tabular anomaly detection. In *SIGKDD Conference on Knowledge Discovery and Data Mining - Applied Data Science Track*, 2023
31. I. Drori, S. J. Zhang, R. S. Shuttleworth, S. Zhang, K. Tyser, Z. Chin, P. Lantigua, S. Surbehera, G. Hunter, D. Austin, L. Tang, Y. Hicke, S. Simhon, S. Karnik, D. Granberry, and M. Udell. From human days to machine seconds: Automatically answering and generating machine learning final exams. In *SIGKDD Conference on Knowledge Discovery and Data Mining - Applied Data Science Track*, 2023, 2206.05442
30. Y. Zhao, A. Townsend, and M. Udell. Probabilistic missing value imputation for mixed categorical and ordered data. In *NeurIPS*, 2022, 2210.06673
29. C. Yang, G. Bender, H. Liu, P.-J. Kindermans, M. Udell, Y. Lu, Q. Le, and D. Huang. Resource-constrained neural architecture search on tabular datasets. In *NeurIPS*, 2022, 2204.07615
28. S. Zhao, Z. Frangella, and M. Udell. NysADMM: faster composite convex optimization via low-rank approximation. In *International Conference on Machine Learning (ICML)*, 2022, 2202.11599
27. C. Yang, Z. Wu, J. Chee, C. D. Sa, and M. Udell. How low can we go: Trading memory for error in low-precision training. In *International Conference on Learning Representations (ICLR)*, 2022, 2106.09686
26. Y. Zhao, E. Landgrebe, E. Shekhtman, and M. Udell. Online missing value imputation and correlation change detection for mixed-type data via gaussian copula. In *AAAI*, 2021, 2009.12326
25. W. T. Stephenson, Z. Frangella, M. Udell, and T. Broderick. Can we globally optimize cross-validation loss? Quasiconvexity in ridge regression. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021, 2107.09194
24. B. Liu, M. Xie, and M. Udell. ControlBurn: Feature selection by sparse forests. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2021, 2107.00219
23. C. Yang, L. Ding, Z. Wu, and M. Udell. TenIPS: Inverse propensity sampling for tensor completion. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021, 2101.00323
22. I. Drori, B. Kates, W. Sickinger, A. Kharkar, B. Dietrich, A. Shporer, and M. Udell. Galaxy TSP: A new billion-node benchmark for TSP. In *NeurIPS Workshop on Learning Meets Combinatorial Algorithms*, 2020
21. I. Drori, A. Kharkar, W. R. Sickinger, B. Kates, Q. Ma, S. Ge, E. Dolev, B. Dietrich, D. P. Williamson, and M. Udell. Learning to solve combinatorial optimization problems on real-world graphs in linear time. In *IEEE International Conference on Machine Learning and Applications (IEEE ICMLA)*, 2020, 2006.03750

20. W. Stephenson, M. Udell, and T. Broderick. Approximate cross-validation with low-rank data in high dimensions. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020, 2008.10547
19. Y. Zhao and M. Udell. Matrix completion with quantified uncertainty through low rank gaussian copula. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020, 2006.10829
18. C. Yang, J. Fan, Z. Wu, and M. Udell. AutoML pipeline selection: Efficiently navigating the combinatorial space. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2020, 2006.04216
17. Y. Zhao and M. Udell. Missing value imputation for mixed data through gaussian copula. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2020, 1910.12845
16. J. Fan, Y. Zhang, and M. Udell. Polynomial matrix completion for missing data imputation and transductive learning. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 3842–3849, 2020, 1912.06989
15. J. Fan, L. Ding, Y. Chen, and M. Udell. Factor group-sparse regularization for efficient low-rank matrix recovery. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, pages 5105–5115, 2019, 1911.05774
14. C. Yang, Y. Akimoto, D. W. Kim, and M. Udell. OBOE: Collaborative filtering for AutoML model selection. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, volume 25, pages 1173–1183. ACM, 2019, 1808.03233
13. J. Fan and M. Udell. Online high-rank matrix completion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8690–8698, 2019, 2002.08934
12. J. Chen, N. Kallus, X. Mao, G. Svacha, and M. Udell. Fairness under unawareness: Assessing disparity when protected class is unobserved. In *FAT*: Conference on Fairness, Accountability, and Transparency*, pages 339–348, 2019, 1811.11154
11. S. Zhou, S. Gupta, and M. Udell. Limited memory Kelley’s method converges for composite convex and submodular objectives. In *Advances in Neural Information Processing Systems*, 2018, 1807.07531
10. N. Kallus, X. Mao, and M. Udell. Causal inference with noisy and missing covariates via matrix factorization. In *Advances in Neural Information Processing Systems*, 2018, 1806.00811
9. J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. Fixed-rank approximation of a positive-semidefinite matrix from streaming data. In *Advances in Neural Information Processing Systems*, 2017, 1706.05736
8. A. Yurtsever, M. Udell, J. A. Tropp, and V. Cevher. Sketchy decisions: Convex low-rank matrix optimization with optimal storage. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1188–1196, 2017, 1702.06838
7. X. Shen, S. Diamond, M. Udell, Y. Gu, and S. Boyd. Disciplined multi-convex programming. In *Chinese Control and Decision Conference (CCDC)*, 2017, 1609.03285

6. D. Davis, B. Edmunds, and M. Udell. The sound of APALM clapping: Faster nonsmooth nonconvex optimization with stochastic asynchronous PALM. In *Advances in Neural Information Processing Systems*, 2016, 1606.02338
5. A. Schuler, V. Liu, J. Wan, A. Callahan, M. Udell, D. Stark, and N. Shah. Discovering patient phenotypes using generalized low rank models. In *Pacific Symposium on Biocomputing (PSB)*, 2016
4. N. Kallus and M. Udell. Revealed preference at scale: Learning personalized preferences from assortment choices. In *The 2016 ACM Conference on Economics and Computation*, New York, NY, USA, 2016. ACM
3. H. Mehmood, M. Udell, and J. Cioffi. Revenue maximization for broadband service providers using revenue capacity. In *IEEE Global Communications Conference*, 2015
2. E. Lee, M. Udell, and S. Wong. Factorization for analog-to-digital matrix multiplication. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015
1. P. LePendou, Y. Liu, S. Iyer, M. Udell, and N. Shah. Analyzing patterns of drug use in clinical notes for patient safety. *Proceedings of the AMIA Summits on Translational Science*, 2012:63, 2012

Refereed Workshops

15. S. Fazliani, Z. Frangella, and M. Udell. Enhancing physics-informed neural networks through feature engineering. In *Theory of AI for Scientific Computing Workshop at COLT*, 2025, 2502.07209
14. V. S. Lokhande, K. Sohn, J. Yoon, M. Udell, C. Lee, and T. Pfister. Towards group robustness in the presence of partial group labels. In *ICML 2022: Workshop on Spurious Correlations, Invariance and Stability*, volume abs/2201.03668, 2022, 2201.03668
13. M. Van Ness and M. Udell. CDF normalization for controlling the distribution of hidden layer activations. In *I (Still) Can't Believe It's Not Better! NeurIPS 2021 Workshop*, 2021
12. I. Drori, L. Liu, Q. Ma, B. Kates, and M. Udell. Zero-shot AutoML. In New York Academy of Sciences, editor, *Annual Machine Learning Symposium*, 2020
11. C. Yang, L. Ding, Z. Wu, and M. Udell. TenIPS: Inverse propensity sampling for tensor completion (workshop). In *OPT2020: 12th Annual Workshop on Optimization for Machine Learning*, 2020
10. E. Landgrebe, Y. Zhao, and M. Udell. Online mixed missing value imputation using gaussian copula. In *ICML Workshop on the Art of Learning with Missing Values (Artemiss)*, 2020
9. I. Drori, L. Liu, S. Koorathota, N. Yi, J. Li, A. Moretti, J. Freire, and M. Udell. AutoML using metadata language embeddings. In *NeurIPS Workshop on Meta-Learning*, 2019, 1910.03698
8. Y. Zhang, K. Song, Y. Sun, S. Tan, and M. Udell. "Why should you trust my explanation?" understanding uncertainty in LIME explanations. In *ICML Workshop AI for Social Good*, 2019, 1904.12991
7. Y. Sun, Y. Guo, J. A. Tropp, and M. Udell. Tensor random projection for low memory dimension reduction. In *NeurIPS Workshop on Relational Representation Learning*, 2018

6. C. Yang, Y. Akimoto, D. W. Kim, and M. Udell. OBOE: Collaborative filtering for AutoML initialization (workshop version). *NeurIPS Workshop on Automated Machine Learning*, 2018, 1808.03233
5. M. Paradkar and M. Udell. Graph-regularized generalized low rank models. In *CVPR Workshop on Tensor Methods in Computer Vision*, 2017
4. N. Kallus and M. Udell. Learning preferences from assortment choices in a heterogeneous population. In *ICML Workshop on Computational Frameworks for Personalization*, 2016, 1509.05113
3. M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd. Convex optimization in Julia. In *SC14 Workshop on High Performance Technical Computing in Dynamic Languages*, 2014, 1410.4821
2. M. Udell, C. Horn, R. Zadeh, and S. Boyd. Generalized low rank models. *NeurIPS Workshop on Distributed Machine Learning and Matrix Computations*, 2014
1. M. Udell and R. Takapoui. Linear bandits, matrix completion, and recommendation systems. *NeurIPS Workshop on Large Scale Matrix Analysis and Inference*, 2013

Miscellaneous: theses, technical reports, book chapters, newsletters, etc.

22. M. V. Ness and M. Udell. dnamite: A python package for neural additive models. *Submitted*, 2025, 2503.07642
21. I. Drori, G. Longhitano, M. Mao, S. Hyun, Y. Zhang, S. Park, Z. Meeks, X.-Y. Zhang, B. Segev, H. Yong, N. Verma, A. Shporer, A. Amit, and M. Udell. Diverse inference and verification for advanced reasoning. 2025, 2502.09955
20. K. Tyser, B. Segev, G. Longhitano, X.-Y. Zhang, Z. Meeks, J. Lee, U. Garg, N. Belsten, A. Shporer, M. Udell, D. Te'eni, and I. Drori. AI-driven review systems: Evaluating LLMs in scalable and bias-aware academic reviews. 2024, 2408.10365
19. K. Tyser, J. Lee, A. Shporer, M. Udell, D. Te'eni, and I. Drori. OpenReviewer: Mitigating challenges in LLM reviewing. 2023
18. A. AhmadiTeshnizi, W. Gao, and M. Udell. OptiMUS: Optimization modeling using MIP solvers and large language models. 2023, 2310.06116
17. M. Udell and Z. Frangella. Randomized numerical linear algebra for optimization. *SIAG/OPT Views and News*, 2022
16. B. Liu, M. Xie, H. Yang, and M. Udell. ControlBurn: Nonlinear feature selection with sparse tree ensembles, 2022
15. S. Zhao, L. Lessard, and M. Udell. An automatic system to detect equivalence between iterative algorithms. 2021, 2105.04684
14. N. Singh, B. Kates, J. Mentch, A. Kharkar, M. Udell, and I. Drori. Privileged zero-shot AutoML, 2021, 2106.13743

13. C. Yang, J. Fan, Z. Wu, and M. Udell. Efficient AutoML pipeline search with matrix and tensor factorization, 2020, 2006.04216
12. I. Drori, L. Liu, Q. Ma, J. Deykin, B. Kates, and M. Udell. Real-time AutoML, 2020
11. B. Liu and M. Udell. Impact of accuracy on model interpretations. 2020, 2011.09903
10. L. Ding, J. Fan, and M. Udell. k FW: A Frank-Wolfe style algorithm with stronger subproblem oracles. *In revision at Computational Optimization and Applications*, 2020, 2006.16142
9. E. A. Ricci, M. Udell, and R. A. Knepper. An information-theoretic approach to persistent environment monitoring through low rank model based planning and prediction. 2020, 2009.01168
8. M. Udell. Big data is low rank. *SIAG/OPT Views and News*, 2019
7. A. Ratner, D. Alistarh, G. Alonso, D. G. Andersen, P. Bailis, S. Bird, N. Carlini, B. Catanzaro, E. Chung, B. Dally, J. Dean, I. S. Dhillon, A. G. Dimakis, P. Dubey, C. Elkan, G. Fursin, G. R. Ganger, L. Getoor, P. B. Gibbons, G. A. Gibson, J. E. Gonzalez, J. Gottschlich, S. Han, K. M. Hazelwood, F. Huang, M. Jaggi, K. G. Jamieson, M. I. Jordan, G. Joshi, R. Khalaf, J. Knight, J. Konecný, T. Kraska, A. Kumar, A. Kyrillidis, J. Li, S. Madden, H. B. McMahan, E. Meijer, I. Mitliagkas, R. Monga, D. G. Murray, D. S. Papailiopoulos, G. Pekhimenko, T. Rekatsinas, A. Rostamizadeh, C. Ré, C. D. Sa, H. Sedghi, S. Sen, V. Smith, A. Smola, D. Song, E. R. Sparks, I. Stoica, V. Sze, M. Udell, J. Vanschoren, S. Venkataraman, R. Vinayak, M. Weimer, A. G. Wilson, E. P. Xing, M. Zaharia, C. Zhang, and A. Talwalkar. SysML: The new frontier of machine learning systems. *CoRR*, abs/1904.03257, 2019, 1904.03257
6. J. A. Tropp, A. Yurtsever, M. Udell, and V. Cevher. More practical sketching algorithms for low-rank matrix approximation. Technical Report 2018-01, California Institute of Technology, Pasadena, California, 2018
5. L. Ding and M. Udell. Frank-Wolfe style algorithms for large scale optimization. In *Large-Scale and Distributed Optimization*. Springer, 2018
4. M. Udell. *Generalized Low Rank Models*. PhD thesis, Stanford University, 2015
3. M. Udell and S. Boyd. PCA on a data frame. 2015
2. M. Udell and S. Boyd. Beyond principal component analysis (PCA). *Biomedical Computation Review*, 2014
1. M. Udell and S. Boyd. Maximizing a sum of sigmoids. 2013

Patents

Madeleine Udell and Oliver Toole. Optimal Design of Residential Photovoltaic Arrays.
Application No. 62/400,542, filed on 9/27/2017, published 3/29/2018.

Jinsung Yoon, Sercan Omer Arik, Madeleine Richards Udell, Chun-Hao Chang. Interpretable Anomaly Detection By Generalized Additive Models With Neural Decision Trees.
Application No. 18/113,267, filed on 2/23/2023, published 8/31/2023.

Da Huang, Chengrun Yang, Pieter-Jan Kindermans, Hanxiao Liu, Quoc V Le, Madeleine Richards Udell, Yifeng Lu, Gabriel Mintzer Bender. Neural Architecture Search with Improved Computational Efficiency.

Application No. 18/839,379, filed on 12/27/2025, published 5/15/2025.

Grants

Current

Stanford Human-Centered AI Seed Grant (PI): *Scientist-in-the-Loop Code Generation for PDE solvers*, \$75,000, April 2026 to March 2027.

AFOSR (PI): *Online Scaled Gradient Method for Stochastic and Distributed Optimization*, \$519,711, December 2025 to November 2028.

Stanford Human-Centered AI – IBM Research award (PI): *Democratizing Decision Optimization for Organizational Decision Makers via LLMs*, \$85,000, September 2024.

Stanford Woods Institute for the Environmental Venture Projects (co-I): *Sustainable and resilient performance-based design solutions using an inverse analysis approach employing multi-objective optimization*, \$250,000, June 2024 to November 2026.

ONR (PI): *Randomized Numerical Linear Algebra for Optimization*, \$523,238, May 2024 to November 2027.

Past

ONR (PI): *Fast Re-routing using Machine Learning*, \$350,000, July 2022 to June 2025.

Alfred P. Sloan Foundation Fellowship:, \$75,000, September 2021 to August 2023.

NSF IIS-1943131 (PI): *CAREER: Accelerating Machine Learning with Low Dimensional Structure*, \$550,000, October 2020 to September 2025.

ONR (PI): *Young Investigator Award: DREAMI: Dimension Reduction for Efficient Automated Machine Intelligence*, \$528,890, June 2020 to August 2024.

Canadian Institutes of Health Research (Senior Personnel): *Using data to guide population health management: A comprehensive evaluation of analytic approaches for population segmentation.* PI: Laura Rosella. \$291,082, March 2020 to March 2023. (Provides some student funding.)

NSF CCF-1740822 (Senior Personnel): *TRIPODS: Data Science for Improved Decision-Making: Learning in the Context of Uncertainty, Causality, Privacy, and Network Structures*, \$1,496,655, October 1, 2017 to September 30, 2020. (.09 calendar [ie, nominal level] effort.) This grant establishes a major data science center at Cornell.

Cornell: Digital Agriculture (co-PI): *Development of a High-Resolution Weather Forecast Database for Digital Agricultural Research and Outreach Applications.* Co-PI: Art Degaetano. \$214,000, 2017–2020.

DARPA FA8750-17-2-0101 (PI): *Composable Robust Structured Data Inference*, \$1,411,602, March 27, 2017 to February 24, 2020. (2 summer months, 1 academic month effort.) The goal of this project is to develop robust tools for imputing missing data in large scale, heterogeneous data tables by using side information about which rows or columns are similar.

Capital One (co-PI). Co-PI: Nathan Kallus. \$60,000, August 1 2018.

Cornell Tech Faculty Exchange Grant (co-PI). Co-PI: Nathan Kallus. \$3,000, 2016–2017.

Awards

| | |
|--|------|
| US National Academy of Sciences Kavli Fellow, | 2023 |
| Microsoft Research Faculty Fellowship Finalist, | 2021 |
| Alfred P. Sloan Foundation Research Fellowship, | 2021 |

| | |
|--|------|
| ONR Young Investigator Award (YIP), | 2020 |
| NSF CAREER Award, | 2020 |
| Cornell Engineering Research Excellence Award, Cornell University | 2020 |
| INFORMS Optimization Society Student Research Prize, INFORMS <i>First place for student Lijun Ding's paper on Storage-Optimal Semidefinite Programming</i> | 2019 |
| INFORMS Undergraduate Operations Research Prize, INFORMS <i>Honorable mention for student Song Zhou's paper on Limited Memory Kelley's Method</i> | 2018 |
| Douglas Whitney '61 Engineering Teaching Excellence Award, Cornell University | 2018 |
| Doing Good with Good OR Student Paper Competition, INFORMS <i>Second place, for Optimal design of efficient rooftop photovoltaic arrays.</i> | 2017 |
| Center for the Mathematics of Information Postdoctoral Fellowship California Institute of Technology | 2015 |
| Gerald J. Lieberman Fellowship, Stanford University <i>Awarded to doctoral students demonstrating the potential to become academic leaders. (12 Lieberman Fellows are selected among all doctoral candidates at Stanford each year.)</i> | 2014 |
| Best Force Multiplier, DARPA PlanX | 2013 |
| Graduate Research Fellowship, National Science Foundation | 2010 |
| Gabilan Graduate Fellowship, Stanford University | 2009 |
| Phi Beta Kappa, Yale University | 2009 |
| Henry Edwards Ellsworth Prize, Yale University <i>Awarded for the best senior thesis research paper in the sciences.</i> | 2009 |
| US Physics Olympics Team Member | 2005 |

Advising

Below I list my Stanford PhD advisees. See my webpage for details.

| | |
|---|--|
| 8. Wanyu Zhang, MS&E | <i>September 2025 – June 2030 (expected)</i> |
| 7. Yinjun Wang, MS&E | <i>September 2024 – June 2029 (expected)</i> |
| 6. Wenzhi Gao, ICME | <i>September 2023 – June 2028 (expected)</i> |
| 5. Shaghayegh Fazliani, Math | <i>April 2024 – June 2027 (expected)</i> |
| 4. Ya-Chi Chu, Math | <i>Jan 2023 – June 2026 (expected)</i> |
| 3. Pratik Rathore, EE | <i>July 2022 – June 2026 (expected)</i> |
| 2. Mike Van Ness, MS&E <i>Interpretable machine learning with applications in healthcare</i> | <i>Sept 2020 – June 2025</i> |
| 1. Zachary Frangella, MS&E <i>Randomized numerical linear algebra for large-scale optimization.</i> | <i>Jan 2020 – June 2025</i> |

PhD committee member

| | |
|-------------------------------|--------------------------------------|
| 16. Max Schaller (EE) | <i>expected graduation June 2026</i> |
| 15. Akshay Rao (CEE) | <i>expected graduation June 2026</i> |
| 14. Sam Liu (MS&E) | <i>expected graduation June 2026</i> |

- | | |
|---|----------------|
| 13. Sergio Gomez (GSB OIT) | June 2025 |
| 12. Tianyi Zhang (CS) | June 2025 |
| 11. Chunlin Sun (ICME) | June 2025 |
| 10. Zhuoyang Liu (GSB OIT) | December 2025 |
| 9. Izabel Aguilar (MS&E) | June 2024 |
| 8. Robert Dyro (AA) | June 2024 |
| 7. Jiahong Ouyang (EE) | June 2024 |
| 6. Robin Alexandra Brown (ICME) | June 2024 |
| 5. Ethan Steinberg (BDS) | September 2023 |
| 4. Yujia Jin (MS&E) | September 2023 |
| 3. Yiping Lu (MS&E) | September 2023 |
| 2. Aldo Carranzo (ICME) | June 2023 |
| 1. Gradey Wang (Mechanical Engineering) | March 2023 |

Postdoctoral advising

- | | |
|-------------------|----------------------------------|
| 1. Connor Lawless | July 2024 – June 2026 (expected) |
|-------------------|----------------------------------|

Cornell advising At Cornell, I advised 3 postdoctoral researchers; 6 PhDs (5 graduated as of June 2022); 16 undergraduates; and 26 Masters students. I served on several other PhD committees and was a member of seven graduate fields: ORIE, CS, CAM, Statistics, ECE, Systems, and Data Science.

Below I list my Cornell PhD advisees. See my webpage for details.

- | | |
|---|--------------|
| 6. Shipu Zhao, Systems Engineering | May 2023 |
| <i>New perspectives in continuous optimization: theory and methodology</i> | |
| 5. Yuxuan Zhao, Statistics | May 2022 |
| <i>Gaussian copula for mixed data with missing values: model estimation and imputation</i> | |
| 4. Chengrun Yang, Electrical and Computer Engineering | May 2022 |
| <i>Automated machine learning under resource constraints</i> | |
| 3. Lijun Ding, ORIE (co-advised by Yudong Chen) | August 2021 |
| <i>Large scale semidefinite programming: simplicity, conditioning, and an efficient algorithm.</i> | |
| 2. Xiaojie Mao, Statistics (co-advised by Nathan Kallus) | May 2021 |
| <i>Machine Learning Methods for Data-driven Decision Making: Contextual Optimization, Causal Inference, and Algorithmic Fairness.</i> | |
| 1. Yiming Sun, Statistics (co-advised by Sumanta Basu) | October 2019 |
| <i>High Dimensional Data Analysis with Dependency and Under Limited Memory.</i> | |

Cornell PhD committee member

- | | |
|---|------------------------|
| 3. Tianyi Shi (Applied Math) | graduated May 2022 |
| 2. Zhengze Zhou (Statistics) | graduated May 2021 |
| 1. Yingjie Bi (Electrical and Computer Engineering) | graduated January 2020 |

Teaching

| | |
|---|--|
| MS&E 318: Safe and Constrained AI <i>Instructor</i> PhD level elective course; 11 students (2026). | Stanford University <i>Winter 2026</i> |
| CME 307 / MS&E 311: Optimization <i>Instructor</i> PhD level required course; 46 students (2023), 28 students (W 2024), 44 students (F 2024). | Stanford University <i>Spring 2023, Winter 2024, Fall 2024, Fall 2025</i> |
| MS&E 478: ASES Breakthrough <i>Instructor</i> Student-led course on venture finance | Stanford University <i>Spring 2024</i> |
| CME 218 / MS&E 218: Applied Data Science <i>Instructor</i> masters level project-based elective; 24 students. | Stanford University <i>Spring 2023</i> |
| MS&E 125: Intro to Applied Statistics <i>Instructor</i> sophomore level required course; 71 students. | Stanford University <i>Spring 2023</i> |
| ORIE 7391: Faster: Algorithmic Ideas for Speeding Up Optimization <i>Instructor</i> PhD level elective course; 13 students. | Cornell University <i>Spring 2022</i> |
| ORIE 4741: Learning with Big Messy Data <i>Instructor</i> Undergraduate level elective course in data analysis; 2016: 85 students. 2017: 117 students. 2019: 147 students. 2020: 117 students. 2021: 155 students. | Cornell University <i>Fall 2016, 2017, 2019, 2020, 2021</i> |
| ORIE 3120: Practical Tools for Operations Research, Data Science, and Machine Learning <i>Instructor</i> Undergraduate level required course; 222 students | Cornell University <i>Spring 2020</i> |
| ORIE 7191: Optimization for Machine Learning <i>Instructor</i> PhD level elective course; 15 students. | Cornell University <i>Spring 2019</i> |
| CS+ORIE+STSCI 1380: Data Science for All <i>Co-instructor</i> Undergraduate level elective course in data analysis; 30 students. | Cornell University <i>Spring 2018</i> |
| ORIE 6326: Convex Optimization <i>Instructor</i> PhD level elective course; 47 students. | Cornell University <i>Spring 2017</i> |
| EE 364b: Convex Optimization II <i>Teaching assistant</i> | Stanford University <i>Spring 2014</i> |
| CVX 101: Convex Optimization <i>Head teaching assistant</i> Taught 10,000 students worldwide. | EdX Stanford <i>Winter 2014</i> |

| | |
|--|--|
| EE 364a: Convex Optimization I <i>Instructor</i> | Stanford University <i>Summer 2013</i> |
| EE 364a: Convex Optimization I <i>Teaching Assistant</i> | Stanford University <i>Winter 2012</i> |
| CME Refresher Course: Discrete Math and Algorithms <i>Instructor</i> | Stanford University <i>September 2011, September 2012</i> |
| CME 305: Discrete Mathematics and Algorithms <i>Teaching Assistant</i> | Stanford University <i>Winter 2011</i> |

**Academic
Service**

| | |
|---|------------------------------------|
| INFORMS Annual Meeting Organizing Committee <i>Invited Sessions Committee Member</i> | INFORMS <i>2026</i> |
| Conference on AI for Operations and Operations for AI: Charting the Future University <i>Co-organizer (with Mohsen Bayati, Ramesh Johari, Arvind Karunakaran, Vasilis Syrgkanis, Ben Van Roy, Gabriel Weintraub)</i> | Stanford <i>2023</i> |
| Generative AI Public Working Group <i>Committee Member</i> | NIST <i>2023–2024</i> |
| SciML Scientific Machine Learning Advisory Committee <i>Committee Member</i> | NumFocus <i>2020–</i> |
| JuliaOpt <i>Co-owner</i> The JuliaOpt organization curates high quality optimization software in the Julia language. | GitHub <i>Fall 2014 –</i> |
| INFORMS AI Strategy Advisory Committee <i>Committee Member</i> <i>Advise the INFORMS board to develop synergies and opportunities in artificial intelligence</i> | INFORMS <i>2019</i> |
| INFORMS Session: Large Scale Semidefinite Programming <i>Co-organizer (with Lijun Ding)</i> | Seattle <i>2019</i> |
| BoydFest: Stephen Boyd’s 60th Birthday Conference <i>Co-organizer (with Maryam Fazel and Mung Chiang)</i> | Stanford University <i>2018</i> |
| ICDM Workshop: Data Driven Discovery of Models <i>Co-organizer (with Christophe Giraud-Carrier and Ishanu Chattopadhyay)</i> | New Orleans <i>2017</i> |
| SIAM Annual: mini-symposium on Robust Low-Rank Models and Applications <i>Co-organizer (with Tamara Kolda)</i> | New Orleans <i>2016</i> |
| MOPTA: session on Large-scale Distributed Convex Optimization <i>Organizer</i> | Lehigh <i>2015</i> |

**University
Service**

| | |
|--|--------------------------------------|
| MS&E MS Program Review Committee <i>Committee Member</i> | Stanford University <i>2025–6</i> |
|--|--------------------------------------|

| | |
|---|--|
| Information Systems Lab Colloquium <i>Faculty Organizer</i> | Stanford University 2024– |
| MS&E Advisory Committee <i>Committee Member</i> | Stanford University 2023– |
| MS&E Teaching Committee <i>Committee Member</i> | Stanford University 2022–2025 |
| ICME Xpo <i>Co-organized workshop for affiliates</i> | Stanford University 2022–2023 |
| ORIE Curriculum Committee <i>Committee Member</i> | Cornell University 2020–2021 |
| CAM Colloquium Committee <i>Committee Member</i> | Cornell University 2018–2019 |
| CAM PhD Admissions Committee <i>Committee Member</i> | Cornell University 2018, 2021 |
| ORIE Curriculum Review Committee <i>Committee Member</i> | Cornell University 2017–2018 |
| ORIE PhD Admissions Committee <i>Committee Member</i> | Cornell University 2016, 2017 |
| Committee on the Future of the School of Engineering <i>Committee Member</i> | Stanford University Fall 2014 – Spring 2015 |
| Represented all engineering doctoral students on faculty committee. Collaborated on proposal addressing faculty hiring and development, research themes and centers, space and facilities, education and outreach, and interdisciplinary research. | |
| C² Computational Consulting <i>Consultant</i> | Stanford University Fall 2011 – Spring 2015 |
| Helped researchers across the university (in physics, computer science, neuroscience, law, immunology, ...) formulate and solve numerical problems. | |
| EE Faculty Search Committee <i>Committee Member</i> | Stanford University Fall 2014 – Spring 2015 |
| Student member on Electrical Engineering broad area search committee. | |
| Information Systems Laboratory Colloquium <i>Coordinator</i> | Stanford University Winter 2012 – Spring 2013 |
| Invited and hosted academic speakers for weekly seminar series. | |
| Committee on Graduate Studies <i>Committee Member</i> | Stanford University Fall 2011 – Spring 2013 |
| Debated and decided policies for all graduate students at Stanford. Approved and reauthorized interdisciplinary graduate programs. | |
| Graduate Student Housing <i>Community Associate</i> | Stanford University Winter 2011 – Spring 2013 |
| Planned and led events for 800 graduate students. | |

Editorial

Associate Editor Foundations and Trends in Machine Learning
2024–

Associate Editor Operations Research (OR)
2024–

Associate Editor SIAM Journal on Mathematics of Data Science (SIMODS)
2023–

Technical Editor Mathematical Programming Computation (MPC)
2021–

2025: NeurIPS Area Chair; NeurIPS Datasets and Benchmarks Senior Area Chair; ICML Senior Area Chair; ...

2024: NeurIPS Area Chair; ICML Senior Area Chair; AFRL; Journal of Machine Learning Research; Math Programming Computation (2); SIAM Journal on Matrix Analysis; Statistics and Computation ...

2023: NeurIPS Area Chair; NeurIPS Datasets and Benchmarks Senior Area Chair; Journal of the Royal Statistical Society; Journal of Machine Learning Research; Math Programming Computation (2); Math of OR; DeepMath; International Journal of Computer Vision ...

2022: Math of OR (2); Operations Research; IEEE transactions on signal processing; IEEE transactions on automatic control; NeurIPS Area Chair; technical editor for three papers at MPC.

2021: Linear Algebra and Applications (LAA); Proceedings of the National Academy of Sciences (PNAS); Mathematics of Operations Research (MOR); SIAM Journal on the Mathematics of Data Science (SIMODS) (2); NeurIPS area chair; Statistics and Computing; SIAM Journal on Matrix Analysis (SIMAX).

2020: SIAM Journal on Optimization (SIOPT); Journal of the American Statistical Society (JASA); Annals of Applied Statistics (AOAS); Linear Algebra and Applications (LAA); SIAM Journal on the Mathematics of Data Science (SIMODS); Science Advances; TPAMI Special Issue on AutoML (2); Proceedings of the National Academy of Sciences (PNAS). Conferences: Learning for Dynamics and Control (6); NeurIPS (6). Grants: NSF RI, ONR.

2019: Journal of the American Statistical Society (JASA); Annals of Applied Statistics (AOAS); SIAM Journal on the Mathematics of Data Science (SIMODS) (2); Journal of Statistical Software (JSS); Optimization Letters (OPTL); SIAM Journal on Matrix Analysis and Applications (SIMAX) (2); SIAM Journal on Optimization (SIOPT). Conferences: Learning for Dynamics and Control (5); ICML (3); NeurIPS (6); AAAI (3). Grants: NSF RI.

2018: SIAM Review; SIAM Journal on Scientific Computing; Linear Algebra and Applications; Stochastic Systems; NeurIPS (6); AAAI meta-reviewer (23); SysML (now MLSys).

2017: Mathematical Programming; NIPS (6); SIAM Journal on Matrix Analysis; AAAI metareviewer (26); ICDM Workshop on Data Driven Discovery of Models (2); Journal of Statistical Software; Automatica; Springer book chapters (2);

Industry Experience

Technical Consultant New York, NY
Coherence Alpha 2025 –
Designed fast optimization solvers for tax-aware portfolio allocation.

Technical Consultant Ithaca, NY
BMD Analytics 2018 –
Provide technical expertise on modeling, learning, and optimization, via structured presentations or hands-on consulting, to firms including BlackRock, Goldman Sachs, Schonfeld Quantitative, Two Sigma, and H2O.ai.

Technical Advisor Oakland, CA
Orchestrated Intelligence 2022 –
Advised on data-driven and robust solutions for supply chain modeling and optimization.

Technical Consultant Stanford, CA
Two Sigma 2021 –
Developed and advised on algorithms for risk management and portfolio optimization.

Visiting Researcher Ithaca, NY and Mountain View, CA
Google Cloud AI Research 2021
Research ideas to accelerate and improve deep learning for tabular and time-series data to solve high-impact business problems in finance, manufacturing, retail, and beyond.

Technical Advisor Santa Monica, CA
Retina AI 2017 – 2022
Advised on technical solutions for problems in e-commerce, including rapid assessment of long term value, retention analysis, and customer segmentation.

Technical Advisor Palo Alto, CA
Aurora Solar 2014 – 2019
Designed optimization algorithms tailored for problems in the solar industry, including design of efficient rooftop photovoltaic array configurations. Compared to designs produced by solar installation experts, the resulting optimized designs deliver the same energy output at lower cost for more than 70% of homes.

Senior Research Scientist San Francisco, CA
Qadium (renamed Expanse, acquired by Palo Alto Networks) 2012 – 2015
Won grants exceeding \$6.5M from DARPA for research in data analytics and cybersecurity.

Lead Data Scientist Arlington, VA
DARPA (via Data Tactics) 2012 – 2013
Wrote 3 white papers to define mission for \$100M DARPA cybersecurity program.

Data Scientist Chicago, IL
Obama for America Fall 2011
Analyzed graph of 70M Facebook users to identify potential donors and target voter registration campaign.

Research Scientist San Mateo, CA
Apixio (acquired by Centene) Summer 2011
Developed a tool to extract structured information about diseases from the unstructured text of doctors' notes.

Sales and Trading Strategist New York, NY
Goldman Sachs Summer 2009
Corrected model of commodities derivatives risk using multiple parameter estimation.

Market Risk Management Analyst New York, NY
Goldman Sachs Summer 2008

Designed and automated a system to evaluate and graph mutual fund risk.

Presentations (Plenary)

| | |
|--|------|
| International Conference on Preconditioning Techniques for Scientific and Industrial Applications , Edinburgh | 2026 |
| <i>Plenary: Learning a Preconditioner</i> | |
| Mathematics of Operations Research , Soesterland, Netherlands | 2025 |
| <i>Keynote: Big Data is Low Rank</i> | |
| <i>Keynote: Hunting the Hessian</i> | |
| <i>Plenary: AI and the Future of Optimization</i> | |
| TILOS workshop on optimization for deep learning , UCSD | 2025 |
| <i>Hunting the Hessian</i> | |
| AlgoPerf Workshop , Meta, Menlo Park | 2025 |
| <i>Hunting the Hessian</i> | |
| MOPTA , Lehigh | 2024 |
| <i>Plenary: Automated Optimization Modeling</i> | |
| AutoML Conference , Berlin | 2023 |
| <i>Plenary: Recommending Learners</i> | |
| East Coast Optimization Meeting , (online) | 2021 |
| <i>Keynote: Scalable Semidefinite Programming</i> | |
| JuliaCon , Baltimore | 2019 |
| <i>Keynote: Big Data is Low Rank using LowRankModels</i> | |
| Women in Data Science , Stanford | 2019 |
| <i>Plenary: Big Data is Low Rank</i> | |
| 100,000 conference attendees worldwide! | |
| JuliaCon , Berkeley | 2017 |
| <i>Plenary: Julia: the Type of Language for Mathematical Programming</i> | |
| MIIS (Tutorial and Invited Talk) , Chinese University of Hong Kong, Shenzhen | 2016 |

Presentations (Invited)

| | |
|---|------|
| Two Sigma Academic Seminar , New York | 2026 |
| <i>AI and the Future of Optimization</i> | |
| INFORMS , Atlanta | 2025 |
| <i>Interpretable Prediction and Feature Selection for Survival Analysis</i> | |
| KDD , Toronto | 2025 |
| <i>Interpretable Prediction and Feature Selection for Survival Analysis</i> | |
| International Conference on Continuous Optimization (ICCOPT) , USC | 2025 |
| <i>Online Scaled Gradient Method</i> | |
| Strategies for Handling Applications with Nonconvexity , Banff | 2025 |
| <i>Online Scaled Gradient Method</i> | |
| INFORMS , Seattle | 2024 |
| <i>AI and the Future of Optimization Modeling</i> | |
| Discrete Optimization Talks (DOTs) , https://talks.discreteopt.com/ | 2024 |
| <i>Automated Optimization Modeling</i> | |
| ORAI , China (online) | 2024 |

| | |
|--|------|
| <i>Automated Optimization Modeling</i> | |
| AAAI Workshop on AI in OR , Vancouver <i>AI and the Future of Optimization Modeling</i> | 2024 |
| Information Theory and Applications (ITA) , San Diego <i>AI and the Future of Optimization Modeling</i> | 2024 |
| Knowledge Discovery and Data Mining (KDD) , Long Beach <i>Panel on AI in Education</i> | 2023 |
| International Conference on Stochastic Programming , UC Davis <i>Low rank approximation for faster convex optimization</i> | 2023 |
| Kavli frontiers of science symposium , Irvine <i>Low rank approximation for faster convex optimization</i> | 2023 |
| One World MINDS seminar , online <i>Low rank approximation for faster convex optimization</i> | 2023 |
| SIAM Math of Data Science , San Diego (online) <i>Low rank approximation for faster convex optimization</i> | 2022 |
| Aspects of Logic and Machine Learning , ASL, Cornell University, <i>Automating Machine Learning</i> | 2022 |
| ICME XPO , Stanford University <i>Automating Machine Learning</i> | 2022 |
| Workshop on Automated Data Science , ECML (online), <i>Structured Models for Automated Machine Learning</i> | 2021 |
| Workshop on Low Rank models and Applications , Fields Institute (online) <i>Imputing Missing Data with the Gaussian Copula</i> | 2021 |
| Complexity of Matrix Computation Seminar , (online) <i>Panelist: What does it mean to compute a low rank approximation of a matrix?</i> | 2021 |
| E-NLA Numerical Linear Algebra Seminar , (online) <i>Scalable Semidefinite Programming</i> | 2021 |
| Mathematical Foundations and Algorithms for Tensor Computations IPAM, UCLA (online) 2021 <i>Low Rank Tucker Approximation of a Tensor from Streaming Data</i> | |
| Women in Data Science Workshop , WiDS Global Conference (online) <i>Automating Machine Learning</i> | 2021 |
| OPTML++ Seminar , MIT (online) <i>Big Data is Low Rank</i> | 2020 |
| QMnet Seminar , Melbourne (online) <i>Missing Data Imputation with Low Rank Models</i> | 2020 |
| The Art of Learning with Missing Values (ARTEMISS) Workshop , ICML (online) <i>Imputing Missing Data with the Gaussian Copula</i> | 2020 |
| SIAM Mathematics of Data Science , (online) <i>Imputing Missing Data with the Gaussian Copula</i> | 2020 |
| Science on Tap , Ithaca, NY <i>Filling in Missing Data: Elections, _____, Healthcare.</i> | 2020 |
| Low-rank models winter school , Villars-Sur-Ollon, Switzerland | 2020 |

| | |
|---|------|
| <i>Low Rank Models for Missing Data and Optimization</i> | |
| Statistics and Computation , Alan Turing Institute, London <i>Big Data is Low Rank</i> | 2020 |
| Reunion Conference on Foundations of Data Science , Simons Institute <i>Missing Value Imputation for Mixed Data Through Gaussian Copula</i> | 2019 |
| INFORMS , Seattle <i>Low Rank Tucker Approximation of a Tensor from Streaming Data</i> | 2019 |
| Knowledge Discovery and Data Mining (KDD) , Anchorage <i>Oboe: Collaborative Filtering for AutoML Initialization</i> | 2019 |
| Machine Learning for Health (ML4H) , Vector Institute, Toronto <i>Representation Learning, Patient Similarity, and Subtyping</i> | 2019 |
| Low Rank Optimization Workshop , Leipzig MPI for Mathematics in the Sciences <i>Low Rank Tucker Approximation of a Tensor from Streaming Data</i> | 2019 |
| Optimization and Statistical Learning , Les Houches <i>Optimal-Storage Semidefinite Programming using Approximate Complementarity</i> | 2019 |
| Women and Mathematics (WAM) Ambassador Program , Cornell University <i>Filling in Missing Data: Elections, _____, Healthcare.</i> | 2019 |
| NeurIPS workshop on AI in financial services , Montreal <i>Moderated Industry Panel</i> | 2018 |
| NeurIPS workshop on AI in financial services , Montreal <i>Fairness under Unawareness</i> | 2018 |
| NeurIPS spotlight talk , Montreal <i>Limited Memory Kelley's Method Converges for Composite Convex and Submodular Optimization</i> | 2018 |
| ISMP , Bordeaux <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2018 |
| DARPA D3M Workshop , Arlington <i>Composable Robust Structured Data Inference: AutoML, Causal Inference, Big Data is Low Rank</i> | 2018 |
| AI in advancement , Cornell <i>Panel Discussion</i> | 2018 |
| INFORMS , Houston <i>Optimal Design of Rooftop Photovoltaic Arrays</i> | 2017 |
| SIMONS Institute , Berkeley <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| Schonfeld Quantitative Conference , New York <i>Convex Optimization Modeling</i> | 2017 |
| STRATA , New York <i>Generalized Low Rank Models</i> | 2017 |
| LCCC workshop on Distributed Optimization (Invited) , Lund <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| SIOPT , Vancouver <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| DARPA D3M Kickoff , Arlington <i>Composable Robust Structured Data Inference</i> | 2017 |

| | |
|--|------|
| Optimization Under Uncertainty Workshop , Duke <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| Yale Alumni in Science and Engineering Talk , New York <i>Filling in Missing Data: Elections, _____, Healthcare.</i> | 2017 |
| INFORMS , Nashville <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2016 |
| ICCOPT , Tokyo <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2016 |
| SIAM Annual Meeting , Boston <i>Generalized Low Rank Models</i> | 2016 |
| DARPA ISAT Workshop on the Future of Storage , New York <i>Low Rank Models for Compression</i> | 2016 |
| TDA 2016 , Leuven <i>Generalized Low Rank Models</i> | 2016 |
| DARPA SIMPLEX program meeting , Stanford University <i>Generalized Low Rank Models</i> | 2015 |
| H2O World , Santa Clara <i>Generalized Low Rank Models</i> | 2015 |
| INFORMS , Philadelphia <i>Bounding Duality Gap for Separable Problems with Linear Constraints</i> | 2015 |
| ISMP , Pittsburgh <i>Convex optimization in Julia</i> | 2015 |
| Optimization in Julia , JuliaCon, Cambridge <i>Convex optimization in Julia</i> | 2015 |
| Bay Area Julia Users Meetup , San Francisco <i>Convex optimization in Julia</i> | 2014 |
| HPTCDL Workshop , SC14, New Orleans <i>Convex optimization in Julia</i> | 2014 |
| INFORMS , San Francisco <i>Sigmoidal Programming</i> | 2014 |
| IPAM Workshop on Optimization , UCLA <i>Sigmoidal Programming</i> | 2013 |
| Presentations (Contributed) | |
| SIAM optimization conference , Seattle <i>Detecting equivalence between iterative algorithms for optimization</i> | 2023 |
| Department Seminars | |
| IEOR-DRO Seminar , Columbia <i>Hunting the Hessian</i> | 2025 |
| Quantitative Sciences Unit Seminar , Stanford University <i>Interpretable Survival Analysis</i> | 2025 |
| Wharton Optimization Seminar , Philadelphia <i>AI and the Future of Optimization Modeling</i> | 2024 |

| | |
|--|------|
| ISE Seminar , USC <i>Low rank approximation for faster convex optimization</i> | 2024 |
| Two Sigma Tech Talk , London <i>Optimization for deep learning: a progress update</i> | 2024 |
| Data Science Seminar , Oxford <i>Low rank approximation for faster convex optimization</i> | 2024 |
| SIAM Conference on Linear Algebra , Paris <i>Low rank approximation for faster convex optimization</i> | 2024 |
| Keller Distinguished Lecture , Caltech <i>Low rank approximation for faster convex optimization</i> | 2024 |
| EE faculty meeting , Stanford University <i>AI and the Future of Optimization Modeling</i> | 2024 |
| BCG Tech Talk , online <i>AI and the Future of Optimization Modeling</i> | 2024 |
| MS&E Alumni Weekend , Stanford University(online) <i>AI and the Future of Optimization Modeling</i> | 2024 |
| Biostatistics Seminar , University of Tennessee Health Sciences (online) <i>Big Data is Low Rank</i> | 2024 |
| Applied Math Seminar , UC Santa Cruz <i>Low rank approximation for faster convex optimization</i> | 2023 |
| Applied Math Seminar , UC Boulder <i>Detecting equivalence between iterative algorithms for optimization</i> | 2023 |
| SILO Seminar , UW Madison <i>Detecting equivalence between iterative algorithms for optimization</i> | 2023 |
| ISL seminar , Stanford University <i>Low rank approximation for faster convex optimization</i> | 2023 |
| Two Sigma Academic Seminar (online) <i>Low rank approximation for faster convex optimization</i> | 2022 |
| Statistics Seminar , University of Chicago (online) <i>Low rank approximation for faster convex optimization</i> | 2022 |
| ORC Seminar , MIT Sloan (online) <i>Big Data is Low Rank</i> | 2022 |
| ORFE Seminar , Princeton (online) <i>Big Data is Low Rank</i> | 2022 |
| MS&E Seminar , Stanford (online) <i>Big Data is Low Rank</i> | 2022 |
| Ocurate / PredictWise Seminar , online <i>Imputing Missing Data with the Gaussian Copula</i> | 2021 |
| Intel Site Visit , Cornell University <i>Pareto optimization to pick the perfect precision</i> | 2021 |
| Data Science Seminar , Johns Hopkins (online) <i>Detecting equivalence between iterative algorithms for optimization</i> | 2021 |
| IEOR-DRO Seminar , Columbia (online) | 2021 |

| | |
|---|------|
| <i>Detecting equivalence between iterative algorithms for optimization</i> | |
| IOE Seminar , Michigan (online) | 2021 |
| <i>Detecting equivalence between iterative algorithms for optimization</i> | |
| AI seminar , Microsoft Research (online) | 2021 |
| <i>Structured Models for Automated Machine Learning</i> | |
| Industrial Engineering Seminar , Sharif University of Technology (online) | 2021 |
| <i>Big Data is Low Rank</i> | |
| Optimization and Statistical Learning Seminar , Northwestern (online) | 2021 |
| <i>Detecting equivalence between iterative algorithms for optimization</i> | |
| Statistics Seminar , University of California at Santa Barbara (online) | 2021 |
| <i>Big Data is Low Rank</i> | |
| AI seminar , Cornell University(online) | 2021 |
| <i>Automating Machine Learning</i> | |
| Computational Science and Engineering Seminar , Georgia Tech (online) | 2021 |
| <i>Big Data is Low Rank</i> | |
| Computational Mathematics and Applications Seminar , Oxford (online) | 2021 |
| <i>Big Data is Low Rank</i> | |
| Ezra Systems Seminar , Cornell (online) | 2020 |
| <i>Imputing Missing Data with the Gaussian Copula</i> | |
| Scientific Computing Seminar , Emory (online) | 2020 |
| <i>Imputing Missing Data with the Gaussian Copula</i> | |
| Mathematics of Data Science Math Seminar , Tufts (online) | 2020 |
| <i>Big Data is Low Rank</i> | |
| Applied Math Seminar , Princeton | 2020 |
| <i>Scalable Semidefinite Programming</i> | |
| Applied Math Seminar , UC Boulder | 2019 |
| <i>Optimal-Storage Semidefinite Programming using Approximate Complementarity</i> | |
| CME 300 , Stanford | 2019 |
| <i>Big Data is Low Rank</i> | |
| Johns Hopkins AMS seminar , Baltimore | 2019 |
| <i>Big Data is Low Rank</i> | |
| CAM Colloquium , Cornell University | 2019 |
| <i>Low Memory Convex Optimization</i> | |
| Rutgers Optimization Seminar , New Brunswick | 2018 |
| <i>Low Memory Convex Optimization</i> | |
| Princeton Optimization Seminar , Princeton | 2018 |
| <i>Low Memory Convex Optimization</i> | |
| UC Davis Mathematics of Data and Decisions Seminar , Davis | 2018 |
| <i>Big Data is Low Rank</i> | |
| Georgia Tech OR Colloquium , Atlanta | 2018 |
| <i>Big Data is Low Rank</i> | |
| Stanford Linear Algebra and Optimization Seminar , Stanford | 2018 |
| <i>Low Memory Convex Optimization</i> | |

| | |
|---|------|
| Ecole Polytechnique: Statistics Special Seminar , Paris <i>Big Data is Low Rank</i> | 2018 |
| Penn State OR Colloquium , State College, PA <i>Big Data is Low Rank</i> | 2018 |
| Cornell Engineering College Council , New York, <i>The New Educational Paradigm: Data Science</i> | 2017 |
| MIT ORC Seminar , Cambridge, MA <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| Capital One Tech Talk , New York <i>Low Rank Models for Automatic Machine Learning and Interpretability</i> | 2017 |
| Two Sigma Tech Talk , New York <i>Generalized Low Rank Models</i> | 2017 |
| CATALYST Academy Field Session: Operations Research , Cornell <i>Outreach session to introduce URM high school students to the discipline of OR</i> | 2017 |
| CURIE Academy Field Session: Operations Research , Cornell <i>Outreach session to introduce female high school students to the discipline of OR</i> | 2017 |
| UW Optimization Seminar , Seattle <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| NYU Numerical Analysis Seminar , New York <i>Sketchy Decisions: Convex Optimization with Optimal Storage</i> | 2017 |
| Goldman Sachs Tech Talk , New York | 2017 |
| CS Brown-Bag Colloquium , Cornell | 2017 |
| SCAN Seminar , Cornell | 2016 |
| CAM Colloquium , Cornell | 2016 |
| JPL Seminar , Pasadena | 2016 |
| Kaiser Permanente , Oakland | 2016 |
| CMI Seminar (I) , California Institute of Technology | 2015 |
| CMI Seminar (II) , California Institute of Technology | 2015 |
| Uber Tech Talk , San Francisco | 2015 |
| Applied Math Seminar , UCLA | 2015 |
| Sandia National Lab Seminar , Livermore | 2015 |
| Google Tech Talk , Mountain View | 2015 |
| Biomedical Informatics Seminar , Stanford University | 2015 |
| Palantir Tech Talk , Palo Alto | 2015 |
| Twitter Tech Talk , San Francisco | 2015 |
| ICME PhD Oral Examination , Stanford University | 2015 |
| H2O Tech Talk , Santa Clara | 2015 |
| Civis Analytics Tech Talk , Chicago | 2015 |
| TTIC Seminar , Toyota Technical Institute of Chicago | 2015 |
| IBM T. J. Watson Research Seminar , Yorktown Heights | 2015 |

| | |
|--|------|
| Hutchin Hill Capital Seminar , New York | 2015 |
| ORIE Seminar , Cornell University | 2015 |
| IEOR Seminar , UC Berkeley | 2015 |
| CMS Seminar , California Institute of Technology | 2015 |
| Heinz College Seminar , Carnegie Mellon University | 2015 |
| Mobilize Seminar , Stanford University | 2014 |
| ICME Seminar , Stanford University | 2014 |
| BlackRock SAE Tech Talk , Stanford University | 2014 |
| JuliaCon , Chicago | 2014 |
| Verizon Labs Tech Talk , Palo Alto | 2014 |
| ICME Seminar , Stanford University | 2013 |
| Marin Software Tech Talk , San Francisco | 2013 |
| Political Psychology Research Seminar , Stanford University | 2013 |
| ICME Student Seminar , Stanford University | 2010 |

Hobbies

Harp, backpacking, running, foraging, ergonomics, carbon sequestration.