Madeleine Udell

227 Frank H.T. Rhodes Hall
136 Hoy Road, Ithaca NY 14853
415-515-7872
udell@cornell.edu
https://people.orie.cornell.edu/mru8/

Academic Employment	**Stanford University**	Stanford, CA
	Assistant Professor	*July 2022 –*
	Gabilan Fellow	
	Department of Management Science and Engineering	

Cornell University — Ithaca, NY

Associate Professor with tenure (on leave) — *July 2022 – June 2024*
Assistant Professor — *July 2016 – June 2022*
Richard and Sybil Smith Sesquicentennial Fellow
Department of Operations Research and Information Engineering
Graduate field member in Operations Research, Computer Science, Applied Mathematics, Data Science, Statistics, Electrical and Computer Engineering, and Systems Engineering

California Institute of Technology — Pasadena, CA

Postdoctoral Fellow, Center for the Mathematics of Information — *June 2015 – June 2016*

Education

Stanford University — Stanford, CA

Ph.D. in Computational and Mathematical Engineering — *June 2015*
GPA: 4.0.
Thesis: *Generalized Low Rank Models.* Advisor: Stephen P. Boyd.

Yale University — New Haven, CT

B.S. in Mathematics and Physics — *June 2009*
Summa cum laude, with honors in mathematics and honors in physics.
GPA: 3.95.
Thesis: *Local Parametrizations via Laplacian Eigenfunctions.* Advisor: Peter W. Jones.

Papers

In the pipeline

11. 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, 2022, 2204.07615

10. Y. Zhao and M. Udell. gcimpute: A package for missing data imputation, 2022, 2203.05089

9. S. Zhao, Z. Frangella, and M. Udell. NysADMM: faster composite convex optimization via low-rank approximation, 2022, 2202.11599

8. C.-H. Chang, J. Yoon, S. Arik, M. Udell, and T. Pfister. Data-efficient and interpretable tabular anomaly detection. *arXiv preprint arXiv:2203.02034*, 2022

7. V. S. Lokhande, K. Sohn, J. Yoon, M. Udell, C. Lee, and T. Pfister. Towards group robustness in the presence of partial group labels. *CoRR*, abs/2201.03668, 2022, 2201.03668

6. Z. Frangella, J. A. Tropp, and M. Udell. Randomized Nyström preconditioning. *Submitted*, 2021, 2110.02820

5. S. Zhao, L. Lessard, and M. Udell. An automatic system to detect equivalence between iterative algorithms. *Submitted*, 2021, 2105.04684

4. J. Fan, L. Ding, C. Yang, and M. Udell. Low-rank tensor recovery with Euclidean-norm-induced Schatten-p quasi-norm regularization. *Submitted*, 2020, 2012.03436

3. L. Ding, J. Fan, and M. Udell. *k*FW: A Frank-Wolfe style algorithm with stronger subproblem oracles. *Submitted*, 2020, 2006.16142

2. L. Ding and M. Udell. A strict complementarity approach to error bound and sensitivity of solution of conic programs. *Submitted*, 2020, 2012.00183

1. N. Sengupta, M. Udell, N. Srebro, and J. Evans. Matrix factorization for missing value imputation and sparse data reconstruction. *Accepted at Sociological Methodology*, 2022

Refereed Journal Articles

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

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

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. LePendu, 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

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 N. Y. A. 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, 2012.03436

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.

13. N. Singh, B. Kates, J. Mentch, A. Kharkar, M. Udell, and I. Drori. Privileged zero-shot AutoML, 2021, 2106.13743

12. C. Yang, J. Fan, Z. Wu, and M. Udell. Efficient AutoML pipeline search with matrix and tensor factorization, 2020, 2006.04216

11. I. Drori, L. Liu, Q. Ma, J. Deykin, B. Kates, and M. Udell. Real-time AutoML, 2020

10. B. Liu and M. Udell. Impact of accuracy on model interpretations. 2020, 2011.09903

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

Grants	**Current**

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

ONR (PI): *Young Investigator Award: DREAMI: Dimension Reduction for Efficient Automated Machine Intelligence*, $528,890, June 2020 to June 2023.

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

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.)

Pending

DOE (PI): *Young Investigator Award: Scalable Scientific Computing with Low Dimensional Structure*, $750,000, October 1, 2021 to September 30, 2026.

NSF (co-PI): *HDR Institute: Integrating data driven methods and science-based models for simulation, engineering, and design.*

Past

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	**Microsoft Research Faculty Fellowship Finalist**,	*2021*
	Alfred P. Sloan Foundation Research Fellowship,	*2021*
	Cornell Engineering Research Excellence Award, Cornell University	*2020*
	INFORMS Optimization Society Student Research Prize, INFORMS	*2019*

First place for student Lijun Ding's paper on Storage-Optimal Semidefinite Programming

INFORMS Undergraduate Operations Research Prize, INFORMS   2018
Honorable mention for student Song Zhou's paper on Limited Memory Kelley's Method

Douglas Whitney '61 Engineering Teaching Excellence Award, Cornell University
2018

Doing Good with Good OR Student Paper Competition, INFORMS   2017
Second place, for Optimal design of efficient rooftop photovoltaic arrays.

Center for the Mathematics of Information Postdoctoral Fellowship   2015
California Institute of Technology

Gerald J. Lieberman Fellowship, Stanford University   2014

Awarded to doctoral students demonstrating the potential to become academic leaders.
(12 Lieberman Fellows are selected among all doctoral candidates at Stanford each year.)

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   2009
Awarded for the best senior thesis research paper in the sciences.

US Physics Olympics Team Member   2005

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

Ishani Karmakar, ICME (rotation)   *expected graduation June 2026*

Ali Ahmaditeshnizi, MS&E (rotation)   *expected graduation June 2026*

Mike Van Ness, MS&E (chair)   *expected graduation June 2025*

Zachary Frangella, MS&E (chair)   *expected graduation August 2024*

PhD committee member Robin Alexandra Brown (ICME)   *expected graduation June 2024*

Cornell advising At Cornell, I advised 3 postdoctoral researchers; 7 PhDs (5 graduated as of June 2022); 16 undergraduates; and 26 Master's students. I served on three PhD committees and was a member of seven graduate fields: ORIE, CS, CAM, Statistics, ECE, Systems, and Data Science.

Richard Lanas Phillips, Computer Science (chair)   *expected graduation May 2024*

Shipu Zhao, Systems Engineering (chair)   *expected graduation May 2023*

Yuxuan Zhao, Statistics (chair)   *graduated May 2022*
Gaussian copula for mixed data with missing values: model estimation and imputation

Chengrun Yang, Electrical and Computer Engineering (chair)   *graduated May 2022*
Automated machine learning under resource constraints

Lijun Ding, Operations Research (co-chair)   *graduated August 2021*
Large scale semidefinite programming: simplicity, conditioning, and an efficient algorithm.

Xiaojie Mao, Statistics (co-chair)   *graduated May 2021*
Machine Learning Methods for Data-driven Decision Making: Contextual Optimization, Causal Inference, and Algorithmic Fairness.

Yiming Sun, Statistics (co-chair) *graduated October 2019*
High Dimensional Data Analysis with Dependency and Under Limited Memory.

Cornell PhD committee member Tianyi Shi (Applied Math) *graduated May 2022*

Zhengze Zhou (Statistics) *graduated May 2021*

Yingjie Bi (Electrical and Computer Engineering) *graduated January 2020*

Teaching	**ORIE 4741: Learning with Big Messy Data**	Cornell University

Instructor *Fall 2016, 2017, 2019, 2020, 2021*
Undergraduate level elective course in data analysis; 2016: 85 students. 2017: 117 students. 2019: 147 students. 2020: 117 students. 2021: 155 students.

ORIE 3120: Practical Tools for Operations Research,
Data Science, and Machine Learning Cornell University
Instructor *Spring 2020*
Undergraduate level required course; 222 students

ORIE 7191: Optimization for Machine Learning Cornell University
Instructor *Spring 2019*
PhD level elective course; 15 students.

CS+ORIE+STSCI 1380: Data Science for All Cornell University
Co-instructor *Spring 2018*
Undergraduate level elective course in data analysis; 30 students.

ORIE 6326: Convex Optimization Cornell University
Instructor *Spring 2017*
PhD level elective course; 47 students.

EE 364b: Convex Optimization II Stanford University
Teaching assistant *Spring 2014*

CVX 101: Convex Optimization EdX Stanford
Head teaching assistant *Winter 2014*
Taught 10,000 students worldwide.

EE 364a: Convex Optimization I Stanford University
Instructor *Summer 2013*

EE 364a: Convex Optimization I Stanford University
Teaching Assistant *Winter 2012*

CME Refresher Course: Discrete Math and Algorithms Stanford University
Instructor *September 2011, September 2012*

CME 305: Discrete Mathematics and Algorithms Stanford University
Teaching Assistant *Winter 2011*

Service	**SciML Scientific Machine Learning Advisory Committee**	NumFocus

Committee Member *2020–*

ORIE Curriculum Committee Cornell University
Committee Member *2020–2021*

INFORMS AI Strategy Advisory Committee INFORMS
Committee Member *2019*
Advise the INFORMS board to develop synergies and opportunities in artificial intelligence

INFORMS Session: Large Scale Semidefinite Programming Seattle
Co-organizer (with Lijun Ding) *2019*

CAM Colloquium Committee Cornell University
Committee Member *2018–2019*

CAM PhD Admissions Committee Cornell University
Committee Member *2018, 2021*

BoydFest: Stephen Boyd's 60th Birthday Conference Stanford University
Co-organizer (with Maryam Fazel and Mung Chiang) *2018*

ORIE Curriculum Review Committee Cornell University
Committee Member *2017–2018*

ICDM Workshop: Data Driven Discovery of Models New Orleans
Co-organizer (with Christophe Giraud-Carrier and Ishanu Chattopadhyay) *2017*

ORIE PhD Admissions Committee Cornell University
Committee Member *2016, 2017*

SIAM Annual: mini-symposium on Robust Low-Rank Models and Applications
New Orleans
Co-organizer (with Tamara Kolda) *2016*

MOPTA: session on Large-scale Distributed Convex Optimization Lehigh
Organizer *2015*

Committee on the Future of the School of Engineering Stanford University
Committee Member *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.

JuliaOpt Github
Co-owner *Fall 2014 –*
The JuliaOpt organization curates high quality optimization software in the Julia language.

C^2 Computational Consulting Stanford University
Consultant *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 Stanford University
Committee Member *Fall 2014 – Spring 2015*
Student member on Electrical Engineering broad area search committee.

Information Systems Laboratory Colloquium Stanford University
Coordinator *Winter 2012 – Spring 2013*
Invited and hosted academic speakers for weekly seminar series.

Committee on Graduate Studies Stanford University

Committee Member   *Fall 2011 – Spring 2013*
Debated and decided policies for all graduate students at Stanford.
Approved and reauthorized interdisciplinary graduate programs.

Graduate Student Housing   Stanford University
Community Associate   *Winter 2011 – Spring 2013*
Planned and led events for 800 graduate students.

Judicial Affairs   Stanford University
Juror   *Fall 2009 – Spring 2010*

Editorial   **Technical Editor**   Mathematical Programming Computation

2021–

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);

Patents   **M. Udell and O. Toole. Optimal Design of Residential Photovoltaic Arrays.**

Application No. 62/400,542, filed on September 27, 2016.

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

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   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.

Talks and posters

Statistics Seminar, University of Chicago (online)   *2022*
Low rank approximation for faster convex optimization

ORC Seminar, MIT Sloan (online)   *2022*
Big Data is Low Rank

ORFE Seminar, Princeton (online)   *2022*
Big Data is Low Rank

MS&E Seminar, Stanford (online)   *2022*
Big Data is Low Rank

Ocurate / PredictWise Seminar, online   *2021*
Imputing Missing Data with the Gaussian Copula

Intel Site Visit, Cornell University   *2021*
Pareto optimization to pick the perfect precision

Data Science Seminar, Johns Hopkins (online)   *2021*
Detecting equivalence between iterative algorithms for optimization

IEOR-DRO Seminar, Columbia (online) *2021*
Detecting equivalence between iterative algorithms for optimization

IOE Seminar, Michigan (online) *2021*
Detecting equivalence between iterative algorithms for optimization

Workshop on Automated Data Science, ECML (online), *2021*
Structured Models for Automated Machine Learning

AI seminar, Microsoft Research (online) *2021*
Structured Models for Automated Machine Learning

Industrial Engineering Seminar, Sharif University of Technology (online) *2021*
Big Data is Low Rank

Workshop on Low Rank models and Applications, Fields Institute (online) *2021*
Imputing Missing Data with the Gaussian Copula

Complexity of Matrix Computation Seminar, (online) *2021*
Panelist: What does it mean to compute a low rank approximation of a matrix?

Optimization and Statistical Learning Seminar, Northwestern (online) *2021*
Detecting equivalence between iterative algorithms for optimization

E-NLA Numerical Linear Algebra Seminar, (online) *2021*
Scalable Semidefinite Programming

Mathematical Foundations and Algorithms for Tensor Computations IPAM, UCLA
(online) *2021*
Low Rank Tucker Approximation of a Tensor from Streaming Data

Statistics Seminar, University of California at Santa Barbara (online) *2021*
Big Data is Low Rank

AI seminar, Cornell University(online) *2021*
Automating Machine Learning

Keynote, East Coast Optimization Meeting (online) *2021*
Scalable Semidefinite Programming

Women in Data Science Workshop, WiDS Global Conference (online) *2021*
Automating Machine Learning

Computational Science and Engineering Seminar, Georgia Tech (online) *2021*
Big Data is Low Rank

Computational Mathematics and Applications Seminar, Oxford University (online)
2021
Big Data is Low Rank

Ezra Systems Seminar, Cornell (online) *2020*
Imputing Missing Data with the Gaussian Copula

OPTML++ Seminar, MIT (online) *2020*
Big Data is Low Rank

Scientific Computing Seminar, Emory (online) *2020*
Imputing Missing Data with the Gaussian Copula

QMnet Seminar, Melbourne (online) *2020*
Missing Data Imputation with Low Rank Models

The Art of Learning with Missing Values (ARTEMISS) Workshop, ICML (online)
2020
Imputing Missing Data with the Gaussian Copula

SIAM Mathematics of Data Science, (online)   *2020*
Imputing Missing Data with the Gaussian Copula

Mathematics of Data Science Math Seminar, Tufts (online)   *2020*
Big Data is Low Rank

Applied Math Seminar, Princeton   *2020*
Scalable Semidefinite Programming

Science on Tap, Ithaca, NY   *2020*
Filling in Missing Data: Elections, _____, Healthcare.

Low-rank models winter school, Villars-Sur-Ollon, Switzerland   *2020*
Low Rank Models for Missing Data and Optimization

Statistics and Computation, Alan Turing Institute, London   *2020*
Big Data is Low Rank

Reunion Conference on Foundations of Data Science, Simons Institute   *2019*
Missing Value Imputation for Mixed Data Through Gaussian Copula

NeurIPS, Vancouver   *2019*
Factor Group-Sparse Regularization for Efficient Low-Rank Matrix Recovery

INFORMS, Seattle   *2019*
Low Rank Tucker Approximation of a Tensor from Streaming Data

Knowledge Discovery and Data Mining (KDD), Anchorage   *2019*
Oboe: Collaborative Filtering for AutoML Initialization

JuliaCon, Baltimore   *2019*
Keynote: Big Data is Low Rank `using LowRankModels`

Applied Math Seminar, UC Boulder   *2019*
Optimal-Storage Semidefinite Programming using Approximate Complementarity

Learning for Dynamics and Control (L4DC), MIT   *2019*
Oboe: Collaborative Filtering for AutoML Initialization (poster)

Machine Learning for Health (ML4H), Vector Institute, Toronto   *2019*
Representation Learning, Patient Similarity, and Subtyping

Low Rank Optimization Workshop, Leipzig MPI for Mathematics in the Sciences *2019*
Low Rank Tucker Approximation of a Tensor from Streaming Data

Optimization and Statistical Learning, Les Houches   *2019*
Optimal-Storage Semidefinite Programming using Approximate Complementarity

Women and Mathematics (WAM) Ambassador Program, Cornell University *2019*
Filling in Missing Data: Elections, _____, Healthcare.

CME 300, Stanford   *2019*
Big Data is Low Rank

Women in Data Science, Stanford   *2019*
Plenary: Big Data is Low Rank
100,000 conference attendees worldwide!

Johns Hopkins AMS seminar, Baltimore   *2019*

Big Data is Low Rank

CAM Colloquium, Cornell University   *2019*
Low Memory Convex Optimization

NeurIPS workshop on AI in financial services, Montreal   *2018*

Moderated Industry Panel

NeurIPS workshop on AI in financial services, Montreal   *2018*
Fairness under Unawareness

NeurIPS spotlight talk, Montreal   *2018*
Limited Memory Kelley's Method Converges for Composite Convex and Submodular Optimization

Rutgers Optimization Seminar, New Brunswick   *2018*
Low Memory Convex Optimization

Princeton Optimization Seminar, Princeton   *2018*
Low Memory Convex Optimization

UC Davis Mathematics of Data and Decisions Seminar, Davis   *2018*
Big Data is Low Rank

Georgia Tech OR Colloquium, Atlanta   *2018*
Big Data is Low Rank

Stanford Linear Algebra and Optimization Seminar, Stanford   *2018*
Low Memory Convex Optimization

ISMP, Bordeaux   *2018*
Sketchy Decisions: Convex Optimization with Optimal Storage

Ecole Polytechnique: Statistics Special Seminar, Paris   *2018*
Big Data is Low Rank

DARPA D3M Workshop, Arlington   *2018*
Composable Robust Structured Data Inference: AutoML, Causal Inference, Big Data is Low Rank

AI in advancement, Cornell   *2018*
Panel Discussion

Penn State OR Colloquium, State College, PA   *2018*
Big Data is Low Rank

Cornell Engineering College Council, New York,   *2017*
The New Educational Paradigm: Data Science

INFORMS, Houston   *2017*
Optimal Design of Rooftop Photovoltaic Arrays

SIMONS Institute, Berkeley   *2017*
Sketchy Decisions: Convex Optimization with Optimal Storage

MIT ORC Seminar, Cambridge, MA   *2017*
Sketchy Decisions: Convex Optimization with Optimal Storage

Capital One Tech Talk, New York   *2017*
Low Rank Models for Automatic Machine Learning and Interpretability

Schonfeld Quantitative Conference, New York   *2017*
Convex Optimization Modeling

STRATA, New York   *2017*
Generalized Low Rank Models

Two Sigma Tech Talk, New York   *2017*
Generalized Low Rank Models

CATALYST Academy Field Session: Operations Research, Cornell   *2017*

Outreach session to introduce URM high school students to the discipline of OR

CURIE Academy Field Session: Operations Research, Cornell   2017
Outreach session to introduce female high school students to the discipline of OR

JuliaCon, Berkeley   2017
Julia: the Type of Language for Mathematical Programming

LCCC workshop on Distributed Optimization (Invited), Lund   2017
Sketchy Decisions: Convex Optimization with Optimal Storage

UW Optimization Seminar, Seattle   2017
Sketchy Decisions: Convex Optimization with Optimal Storage

SIOPT, Vancouver   2017
Sketchy Decisions: Convex Optimization with Optimal Storage

DARPA D3M Kickoff, Arlington   2017
Composable Robust Structured Data Inference

Optimization Under Uncertainty Workshop, Duke   2017
Sketchy Decisions: Convex Optimization with Optimal Storage

Yale Alumni in Science and Engineering Talk, New York   2017
Filling in Missing Data: Elections, _____, Healthcare.

NYU Numerical Analysis Seminar, New York   2017
Sketchy Decisions: Convex Optimization with Optimal Storage

Goldman Sachs Tech Talk, New York   2017

CS Brown-Bag Colloquium, Cornell   2017

MIIS (Tutorial and Invited Talk), Chinese University of Hong Kong, Shenzhen   2016

NIPS, Barcelona   2016

INFORMS, Nashville   2016

SCAN Seminar, Cornell   2016

CAM Colloquium, Cornell   2016

ICCOPT, Tokyo   2016

SIAM Annual Meeting, Boston   2016

JPL Seminar, Pasadena   2016

DARPA ISAT Workshop on the Future of Storage, New York   2016

Kaiser Permanente, Oakland   2016

TDA 2016, Leuven   2016

CMI Seminar (I), California Institute of Technology   2015

CMI Seminar (II), California Institute of Technology   2015

DARPA SIMPLEX program meeting, Stanford University   2015

H2O World, Santa Clara   2015

Uber Tech Talk, San Francisco   2015

INFORMS, Philadelphia   2015

Applied Math Seminar, UCLA   2015

Sandia National Lab Seminar, Livermore   2015

ISMP, Pittsburgh		*2015*
Optimization in Julia, JuliaCon, Cambridge		*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*
Distributed Machine Learning Workshop, NIPS, Montreal		*2014*
HPTCDL Workshop, SC14, New Orleans		*2014*
INFORMS, San Francisco		*2014*
ICME Seminar, Stanford University		*2014*
Bay Area Julia Users Meetup, San Francisco		*2014*
BlackRock SAE Tech Talk, Stanford University		*2014*
Modern Massive Data Sets (MMDS), UC Berkeley		*2014*
JuliaCon, Chicago		*2014*
Verizon Labs Tech Talk, Palo Alto		*2014*
IPAM Workshop on Mathematics of Politics, UCLA		*2013*
Workshop on Large Matrices, NIPS, Lake Tahoe		*2013*
IPAM Workshop on Optimization, UCLA		*2013*
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*

Numbers **Erdös number:** 3

Bacon number: 3

Erdös-Bacon number: 6

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