

CHRISTOPHER DE SA

Gates Hall 426, 107 Hoy Rd, Ithaca, NY 14853
(607) 255-9207

cmd353@cornell.edu

<https://www.cs.cornell.edu/~cdesa/>

RESEARCH INTERESTS

My research interests include algorithmic, software, and hardware techniques for high-performance data analytics, with a focus on fast parallel variants of stochastic algorithms such as stochastic gradient descent (SGD) and Gibbs sampling. I am also interested in using these techniques to construct data analytics and machine learning frameworks, including for deep learning, that are efficient, parallel, and distributed.

APPOINTMENTS

Assistant Professor, Department of Computer Science, Cornell University, 2017–

EDUCATION

Stanford University, Stanford California

Ph.D. in Electrical Engineering, 2017

Thesis: Fast Stochastic Algorithms for Machine Learning

Advisors: Professors Kunle Olukotun and Christopher Ré

M.S. in Electrical Engineering, 2011

B.S. in Electrical Engineering, 2011, with Distinction

PUBLICATIONS

A. Feder Cooper, Katherine Lee, Madiha Choksi, Solon Barocas, Christopher De Sa, James Grimmelmann, Jon Kleinberg, Siddhartha Sen, Baobao Zhang. “Arbitrariness of Prediction in Fair Classification.” In *The 38th Annual AAAI Conference on Artificial Intelligence*, February 2024. ***Best Student Paper (Honorable Mention)***.

Zilu Li, Guandao Yang, Xi Deng, Christopher De Sa, Bharath Hariharan, Steve Marschner. “Neural Caches for Monte Carlo Partial Differential Equation Solvers.” In *SIGGRAPH Asia 2023*, December 2023.

Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, Christopher De Sa. “QuIP: 2-Bit Quantization of Large Language Models With Guarantees.” In *Proceedings of the 36th Neural Information Processing Systems Conference*, December 2023. ***Spotlight***.

A. Feder Cooper, Wentao Guo, Khiem Pham, Tiancheng Yuan, Charlie F. Ruan, Yucheng Lu, Christopher De Sa. “CD-GraB: Coordinating Distributed Example Orders for Provably Accelerated Training.” In *Proceedings of the 36th Neural Information Processing Systems Conference*, December 2023.

Albert Tseng, Tao Yu, Toni J.B. Liu, Christopher De Sa. “Coneheads: Hierarchy Aware Attention.” In *Proceedings of the 36th Neural Information Processing Systems Conference*, December 2023.

Kush Bhatia, Avaniika Narayan, Christopher De Sa, Christopher Re. “TART: A plug-and-play Transformer module for task-agnostic reasoning.” In *Proceedings of the 36th Neural Information Processing Systems Conference*, December 2023.

Isay Katsman, Eric Ming Chen, Sidhanth Holalkere, Anna Asch, Aaron Lou, Ser-Nam Lim, Christopher De Sa. “Riemannian Residual Neural Networks.” In *Proceedings of the 36th Neural Information Processing Systems Conference*, December 2023.

Oliver Richardson, Joe Halpern, Christopher De Sa. “Inference for Probabilistic Dependency Graphs.” In *UAI: the 39th Conference on Uncertainty in Artificial Intelligence*, August 2023.

Jue Wang, Yucheng Lu, Binhang Yuan, Beidi Chen, Percy Liang, Christopher De Sa, Christopher Re, Ce Zhang. “CocktailSGD: Fine-tuning Foundation Models over 500Mbps Networks.” In *ICML: the Thirty-ninth International Conference on Machine Learning*, July 2023.

Yingheng Wang, Yair Schiff, Aaron Gokaslan, Weishen Pan, Fei Wang, Christopher De Sa, Volodymyr Kuleshov . “InfoDiffusion: Representation Learning Using Information Maximizing Diffusion Models.” In *ICML: the Thirty-ninth International Conference on Machine Learning*, July 2023.

Yucheng Lu, Amir Yazdanbakhsh, Shivani Agrawal, Suvinay Subramanian, Oleg Rybakov, Christopher De Sa. “STEP: Learning N:M Structured Sparsity Masks from Scratch with Pre-condition.” In *ICML: the Thirty-ninth International Conference on Machine Learning*, July 2023.

Yucheng Lu, Christopher De Sa. “Decentralized Learning: Theoretical Optimality and Practical Improvements.” In *JMLR: Journal of Machine Learning Research*, April 2023.

Yucheng Lu, Conglong Li, Minjia Zhang, Christopher De Sa, Yuxiong He. “Maximizing Communication Efficiency for Large-scale Training via 0/1 Adam.” In *ICLR: The Eleventh International Conference on Learning Representations*, May 2023.

Tao Yu, Christopher De Sa. “Random Laplacian Features for Learning with Hyperbolic Space.” In *ICLR: The Eleventh International Conference on Learning Representations*, May 2023.

Yucheng Lu, Wentao Guo, Christopher De Sa. “GraB: Finding Provably Better Data Permutations than Random Reshuffling.” In *NeurIPS: Proceedings of the 35th Neural Information Processing Systems Conference*, December 2022.

Tao Yu, Yichi Zhang, Zhiru Zhang, Christopher De Sa. “Understanding Hyperdimensional Computing for Parallel Single-Pass Learning.” In *NeurIPS: Proceedings of the 35th Neural Information Processing Systems Conference*, December 2022.

Jerry Chee, Megan Renz, Anil Damle, Christopher De Sa. “Model Preserving Compression for Neural Networks.” In *NeurIPS: Proceedings of the 35th Neural Information Processing Systems Conference*, December 2022.

Christopher De Sa, Satyen Kale, Jason D. Lee, Ayush Sekhari, Karthik Sridharan. “From Gradient Flow on Population Loss to Learning with Stochastic Gradient Descent.” In *NeurIPS: Proceedings of the 35th Neural Information Processing Systems Conference*, December 2022.

A. Feder Cooper, Jonathan Frankle, Christopher De Sa. “Non-Determinism and the Lawlessness of ML Code.” In *CSLAW: 2nd ACM Symposium on Computer Science and Law*, November 2022. **Oral.**

Ruqi Zhang, Andrew Wilson, Christopher De Sa. “Low-Precision Stochastic Gradient Langevin Dynamics.” In *ICML: the Thirty-ninth International Conference on Machine Learning*, July 2022. **Spotlight.**

Yucheng Lu, Si Yi Meng, and Christopher De Sa. “A General Analysis of Example-Selection for Stochastic Gradient Descent.” In *ICLR: Proceedings of the Tenth International Conference on Learning Representations*, April 2022. **Spotlight.**

Chengrun Yang, Ziyang Wu, Jerry Chee, Christopher De Sa, and Madeleine Udell. “How Low Can We Go: Trading Memory for Error in Low-Precision Training.” In *ICLR: Proceedings of the Tenth International Conference on Learning Representations*, April 2022.

A. Feder Cooper, Yucheng Lu, and Christopher De Sa. “Hyperparameter Optimization Is Deceiving Us, and How to Stop It.” In *NeurIPS: Proceedings of the 34th Neural Information Processing Systems Conference*, December 2021.

Tao Yu, Christopher De Sa. “Representing Hyperbolic Space Accurately using Multi-Component Floats.” In *NeurIPS: Proceedings of the 34th Neural Information Processing Systems Conference*, December 2021.

Isay Katsman, Aaron Lou, Derek Lim, Qingxuan Jiang, Ser-Nam Lim, Christopher De Sa. “Equivariant Manifold Flows.” In *NeurIPS: Proceedings of the 34th Neural Information Processing Systems Conference*, December 2021.

A. Feder Cooper, Karen Levy, Christopher De Sa. “Accuracy-Efficiency Trade-Offs and Accountability in Distributed ML Systems.” In *EAAMO: Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (to appear)*, October 2021. **Oral.**

Yucheng Lu, Christopher De Sa. “Optimal Complexity in Decentralized Training.” In *ICML: the Thirty-eighth International Conference on Machine Learning*, July 2021. **Outstanding Paper (Honorable Mention).**

Yucheng Lu, Youngsuk Park, Lifan Chen, Yuyang Wang, Christopher De Sa, Dean Foster. “Variance Reduced Training with Stratified Sampling for Forecasting Models.” In *ICML: the Thirty-eighth International Conference on Machine Learning*, July 2021.

Johan Björck, Xiangyu Chen, Christopher De Sa, Carla Gomes, Kilian Weinberger. “Low-Precision Reinforcement Learning: Running Soft Actor-Critic in Half Precision.” In *ICML: the Thirty-eighth International Conference on Machine Learning*, July 2021.

Isay Katsman, Aaron Lou, Derek Lim, Qingxuan Jiang, Ser-Nam Lim, Christopher De Sa. “Equivariant Manifold Flows.” In *ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, July 2021.

A. Feder Cooper, Yucheng Lu, and Christopher De Sa. “Hyperparameter Optimization Is Deceiving Us, and How to Stop It.” In *ICLR 2021, Workshop on Robust ML (RML)*, May 2021.

Jessica Zosa Forde*, A. Feder Cooper*, Kweku Kwegyir-Aggrey, Christopher De Sa, Michael Littman. “Model Selection’s Disparate Impact in Real-World Deep Learning Applications.” In *ICLR 2021, Workshop on the Science and Engineering of Deep Learning (SEDL)*, May 2021. **Oral.**

Ruqi Zhang, Yingzhen Li, Christopher De Sa, Sam Devlin, Cheng Zhang. “Meta-Learning Divergences for Variational Inference.” In *AISTATS: Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, April 2021.

Bowen Yang, Jian Zhang, Jonathan Li, Christopher Ré, Christopher R. Aberger, Christopher De Sa. “PipeMare: Asynchronous Pipeline Parallel DNN Training.” In *MLSys: Proceedings of the 4th Conference on Machine Learning and Systems*, April 2021.

Ruqi Zhang, A. Feder Cooper, Christopher De Sa. “Asymptotically Optimal Exact Minibatch Metropolis-Hastings.” In *NeurIPS: Proceedings of the 33rd Neural Information Processing Systems Conference*, December 2020. **Spotlight.**

Christopher De Sa. “Random Reshuffling is Not Always Better.” In *NeurIPS: Proceedings of the 33rd Neural Information Processing Systems Conference*, December 2020. **Spotlight.**

Aaron Lou, Derek Lim, Isay Katsman, Leo Huang, Qingxuan Jiang, Ser-Nam Lim, Christopher De Sa. “Neural Manifold Ordinary Differential Equations.” In *NeurIPS: Proceedings of the 33rd Neural Information Processing Systems Conference*, December 2020.

Yucheng Lu, Christopher De Sa. “Moniqua: Modulo Quantized Communication in Decentralized SGD.” In *ICML: the Thirty-seventh International Conference on Machine Learning*, July 2020.

Aaron Lou, Isay Katsman, Qingxuan Jiang, Serge Belongie, Ser Nam Lim, Christopher De Sa. “Differentiating through the Fréchet Mean.” In *ICML: the Thirty-seventh International Conference on Machine Learning*, July 2020.

A. Feder Cooper, Karen Levy, Christopher De Sa. “Regulating Accuracy-Efficiency Trade-Offs in Distributed Machine Learning Systems.” In *LML 2020: ICML Workshop on Law and Machine Learning*, July 2020. **Oral.**

Aaron Lou, Derek Lim, Isay Katsman, Leo Huang, Qingxuan Jiang, Ser-Nam Lim, Christopher De Sa. “Neural Manifold Ordinary Differential Equations.” In *INNF+ 2020: ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*, July 2020.

Ruqi Zhang, A. Feder Cooper, Christopher De Sa. “AMAGOLD: Amortized Metropolis Adjustment for Efficient Stochastic Gradient MCMC.” In *AISTATS: The 23rd International Conference on Artificial Intelligence and Statistics*, June 2020.

Weizhe Hua, Yuan Zhou, Christopher De Sa, Zhiru Zhang, G. Edward Suh. “Channel Gating Neural Networks.” In *NeurIPS: Proceedings of the 32nd Neural Information Processing Systems Conference*, December 2019.

Tao Yu, Christopher De Sa. “Numerically Accurate Hyperbolic Embeddings Using Tiling-Based Models.” In *NeurIPS: Proceedings of the 32nd Neural Information Processing Systems Conference*, December 2019. **Spotlight.**

Ruqi Zhang, Christopher De Sa. “Poisson-Minibatching for Gibbs Sampling with Convergence Rate Guarantees.” In *NeurIPS: Proceedings of the 32nd Neural Information Processing Systems Conference*, December 2019. **Spotlight.**

Zheng Li, Christopher De Sa. “Dimension-Free Bounds for Low-Precision Training.” In *NeurIPS: Proceedings of the 32nd Neural Information Processing Systems Conference*, December 2019.

Tianyi Zhang, Zhiqiu Lin, Guandao Yang, Christopher De Sa. “QPyTorch: A Low-Precision Arithmetic Simulation Framework.” In *EMC2: Workshop on Energy Efficient ML and Cognitive Computing, at NeurIPS*, December 2019.

Weizhe Hua, Yuan Zhou, Christopher De Sa, Zhiru Zhang, G. Edward Suh. “Boosting the Performance of CNN Accelerators with Dynamic Fine-Grained Channel Gating.” In *MICRO '52 Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, October 2019.

Ken Birman, Bharath Hariharan, Christopher De Sa. “Cloud-Hosted Intelligence for Real-time IoT Applications.” In *SIGOPS Operating Systems Review 53*, July 2019.

Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Christopher De Sa, Zhiru Zhang. “Improving Neural Network Quantization without Retraining using Outlier Channel Splitting.” In *ICML: the Thirty-sixth International Conference on Machine Learning*, June 2019.

Jayadev Acharya, Christopher De Sa, Dylan J. Foster, Karthik Sridharan. “Distributed Learning with Sublinear Communication.” In *ICML: the Thirty-sixth International Conference on Machine Learning*, June 2019. **Long oral.**

Tri Dao, Albert Gu, Alexander J. Ratner, Virginia Smith, Christopher De Sa, Christopher Ré. “A Kernel Theory of Modern Data Augmentation.” In *ICML: the Thirty-sixth International Conference on Machine Learning*, June 2019.

Guandao Yang, Tianyi Zhang, Polina Kirichenko, Junwen Bai, Andrew Gordon Wilson, Christopher De Sa. “SWALP : Stochastic Weight Averaging in Low Precision Training.” In *ICML: the Thirty-sixth International Conference on Machine Learning*, June 2019.

Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Christopher De Sa, Zhiru Zhang. “Building Efficient Deep Neural Networks with Unitary Group Convolutions.” In *CVPR: The Conference on Computer Vision and Pattern Recognition*, June 2019.

Christopher De Sa, Ihab F. Ilyas, Benny Kimelfeld, Christopher Ré, Theodoros Rekatsinas. “A Formal Framework For Probabilistic Unclean Databases.” In *ICDT: 22nd International Conference on Database Theory*, March 2019.

Ilse M. Van Meerbeek, Jose A. Barreiros, Robert F. Shepherd, Christopher M. De Sa. “Addressing sensor drift in a proprioceptive optical foam system.” In *Proc. SPIE 10970: Sensors and Smart Structures Technologies for Civil, Mechanical, and Aerospace Systems 2019*, March 2019.

Ilse Van Meerbeek, Christopher De Sa, Robert Shepherd. “Soft optoelectronic sensory foams with proprioception.” In *Science Robotics*, November 2018.

Christopher De Sa, Vincent Chen, Wing Wong. “Minibatch Gibbs Sampling on Large Graphical Models.” In *ICML: Proceedings of the 35rd International Conference on Machine Learning*, July 2018. **Long oral.**

Frederic Sala, Christopher De Sa, Albert Gu, Christopher Ré. “Representation Tradeoffs for Hyperbolic Embeddings.” In *ICML: Proceedings of the 35rd International Conference on Machine Learning*, July 2018. **Long oral.**

Dan Alistarh, Christopher De Sa, Nikola Konstantinov. “The Convergence of Stochastic Gradient Descent in Asynchronous Shared Memory.” In *PODC: Principles of Distributed Computing*, July 2018.

Christopher De Sa, Bryan He, Ioannis Mitliagkas, Christopher Ré, Peng Xu. “Accelerated Stochastic Power Iteration.” In *AISTATS: The 21st International Conference on Artificial Intelligence and Statistics*, April 2018.

Christopher De Sa, Albert Gu, Rohan Puttagunta, Christopher Ré, Atri Rudra. “A Two Pronged Progress in Structured Dense Matrix Multiplication.” In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, January 2018.

Tri Dao, Christopher De Sa, Christopher Ré. “Gaussian Quadrature for Kernel Features.” In *NeurIPS: Proceedings of the 30th Neural Information Processing Systems Conference*, December 2017. **Spotlight.**

Christopher De Sa, Matt Feldman, Christopher Ré, Kunle Olukotun. “Understanding and Optimizing Asynchronous Low-Precision Stochastic Gradient Descent.” In *ISCA: 44th International Symposium on Computer Architecture*, June 2017.

Paroma Varma, Dan Iter, Christopher De Sa, Christopher Ré. “Flipper: A Systematic Approach to Debugging Training Sets.” In *HILDA: Proceedings of the 2nd Workshop on Human-In-the-Loop Data Analytics*, at SIGMOD, May 2017.

Alex Ratner, Christopher De Sa, Sen Wu, Daniel Selsam, Christopher Ré. “Data Programming: Creating Large Training Sets, Quickly.” In *NeurIPS: Proceedings of the 29th Neural Information Processing Systems Conference*, December 2016.

Bryan He, Christopher De Sa, Ioannis Mitliagkas, Christopher Ré. “Scan Order in Gibbs Sampling: Models in Which it Matters and Bounds on How Much.” In *NeurIPS: Proceedings of the 29th Neural Information Processing Systems Conference*, December 2016.

Paroma Varma, Rose Yu, Dan Iter, Christopher De Sa, Christopher Ré. “Socratic Learning: Empowering the Generative Model.” In *FiLM-NIPS: Future of Interactive Learning Machines at NIPS*, December 2016.

Christopher De Sa, Kunle Olukotun, Christopher Ré. “Ensuring Rapid Mixing and Low Bias for Asynchronous Gibbs Sampling.” In *ICML: Proceedings of the 33rd International Conference on Machine Learning*, June 2016. **Best Paper Award.**

Jian Zhang, Christopher De Sa, Ioannis Mitliagkas, Christopher Ré. “Parallel SGD: When does Averaging Help?.” In *OptML: Optimization Methods for the Next Generation of Machine Learning, workshop at ICML*, June 2016.

Christopher De Sa, Alex Ratner, Christopher Ré, Jaeho Shin, Feiran Wang, Sen Wu, Ce Zhang. “DeepDive: Declarative Knowledge Base Construction.” In *SIGMOD Record, Research Highlight*, April 2016.

Raghu Prabhakar, David Koeplinger, Kevin J. Brown, HyoukJoong Lee, Christopher De Sa, Christos Kozyrakis, Kunle Olukotun. “Generating Configurable Hardware from Parallel Patterns.” In *ASPLOS: 21st Int’l Conference on Architectural Support for Programming Languages and Operating Systems*, April 2016.

Kevin J. Brown, HyoukJoong Lee, Tiark Rompf, Arvind K. Sujeeth, Christopher De Sa, Christopher Aberger, Kunle Olukotun. “Have Abstraction and Eat Performance, Too: Optimized Heterogeneous Computing with Parallel Patterns.” In *CGO: International Symposium on Code Generation and Optimization*, March 2016.

Christopher De Sa, Ce Zhang, Kunle Olukotun, Christopher Ré. “Rapidly Mixing Gibbs Sampling for a Class of Factor Graphs Using Hierarchy Width.” In *NIPS: Proceedings of the 28th Neural Information Processing Systems Conference*, December 2015. **Spotlight.**

Christopher De Sa, Ce Zhang, Kunle Olukotun, Christopher Ré. “Taming the Wild: A Unified Analysis of Hogwild!-Style Algorithms.” In *NIPS: Proceedings of the 28th Neural Information Processing Systems Conference*, December 2015.

Jaeho Shin, Sen Wu, Feiran Wang, Ce Zhang, Christopher De Sa, Christopher Ré. “Incremental Knowledge Base Construction Using DeepDive.” In *VLDB: Proceedings of the 41st International Conference on Very Large Data Bases*, September 2015. **Best of Issue.**

Christopher De Sa, Kunle Olukotun, Christopher Ré. “Global Convergence of Stochastic Gradient Descent for Some Nonconvex Matrix Problems.” In *ICML: Proceedings of the 32nd International Conference on Machine Learning*, July 2015.

MANUSCRIPTS

Yucheng Lu, Jack Nash, Christopher De Sa. “MixML: A Unified Analysis of Weakly Consistent Parallel Learning.” *Manuscript updated May 2020.*

Ritchie Zhao, Christopher De Sa, Zhiru Zhang. “Overwrite Quantization: Opportunistic Outlier Handling for Neural Network Accelerators.” *Manuscript updated October 2019*.

Christopher De Sa, Megan Leszczynski, Jian Zhang, Alana Marzoev, Christopher R. Aberger, Kunle Olukotun, Christopher Ré. “High-Accuracy Low-Precision Training.” *Manuscript updated December 2018*.

Alexander Ratner et. al.. “SysML: The New Frontier of Machine Learning Systems.” *On arxiv March 2019*.

PHD STUDENTS

Ruqi Zhang — PhD 2021, Statistics (now an Assistant Professor at Purdue CS)

Yucheng Lu — PhD Student, Computer Science (now at Together)

A. Feder Cooper — PhD Student, Computer Science

Jerry Chee — PhD Student, Computer Science

Tao Yu — PhD Student, Computer Science

Yaohui Cai — PhD Student, Electrical and Computing Engineering

Si Yi (Cathy) Meng — PhD Student, Computer Science

Albert Tseng — PhD Student, Computer Science

AWARDS, FELLOWSHIPS, AND GRANTS

NSF CAREER Grant

Large-Scale Markov Chain Monte Carlo for Reliable Machine Learning, 2021.

NSF RII Small Grant

Reliable Machine Learning in Hyperbolic Spaces, 2020.

ICML Outstanding Paper Award (Honorable Mention), 2021

Awarded for “Optimal Complexity in Decentralized Training.”

Mr. & Mrs. Richard F. Tucker Excellence in Teaching Award, 2020

Awarded for excellent teaching in the College of Engineering.

ICML Best Paper Award, 2016

Awarded for “Ensuring Rapid Mixing and Low Bias for Asynchronous Gibbs Sampling.”

Stanford Graduate Fellowship, Stanford University, 2013–2015

Awarded to outstanding doctoral students in science and engineering

Numerical Technologies Co-Founders Fellowship, Stanford University, 2012–2013

Awarded to the top scorer on the Stanford EE Qualifying Exam

TALKS

Invited Talk at Together, January 2023. — GraB: Better Example Orders than Random Reshuffling.

Invited Talk at Microsoft Research, January 2023. — Better Data Ordering for Stochastic Gradient Descent.

Invited Talk at Google Brain, November 2022. — Understanding (and Surpassing) the Limits of Hyperdimensional Computing.

CSL Invited Speakers Series, November 2022. — Understanding (and Surpassing) the Limits of Hyperdimensional Computing.

Invited talk at Snorkel AI, August 2022 — Beyond Random Reshuffling: Example Order Selection in Machine Learning.

Information Theory and Applications Workshop (ITA) 2022 — Beyond Random Reshuffling: Example Order Selection in Machine Learning.

Cornell FIND Seminar 2022 — Example Selection Methods for Stochastic Gradient Descent

Columbia CCSS Tech Talk 2022 — Machine Learning in Hyperbolic Space

Emory CS Seminar 2020 — Machine Learning Systems for the Data Tsunami

SOSP 2019 — Vector Institute Workshop on Machine Learning Systems

ISCA 2017 — Workshop on Architecture for Graph Processing

STOC 2017 — Plenary Talk

STOC 2017 — Workshop on New Challenges in Machine Learning

**TEACHING
EXPERIENCE**

Instructor for 4780/5780 — Machine Learning

Cornell — 2022, 2018

Instructor for 4787/5777 — Principles of Large-Scale Machine Learning

Cornell — 2023, 2022, 2021, 2020, 2019

Instructor for 6787 — Advanced Machine Learning Systems

Cornell — 2021, 2020, 2019, 2018, 2017

Instructor for 7792 — Special Topics in Machine Learning

Cornell — 2023

SERVICE

Reviews: NeurIPS, IEEE Micro, IEEE TKDE, JASA, FOCS, JMLR, IBM AI Systems Day, AAAI, AISTATS, SIGMOD, VLDB, 2021 Cornell Institute for Digital Agriculture RIF student summer stipend application review committee, NSF Panel 2021, 2020

Program Committee Member: AAAI 2021, 2020, 2019; AIOT Workshop @ KDD 2022; CIDA Digital Agriculture Hackathon Chair 2022; Eurosys 2019, 2021; ICLR 2021; ICML 2021, 2020, 2019; MLSys 2023, 2021, 2020, 2019, 2018; MLSys Workshop and Tutorial Chair 2022; SIGMOD 2021; UAI 2020

Cornell CS Faculty Recruiting Committee — 2020, 2021, 2022

Cornell CS PhD Admissions Committee — 2019, 2018

Cornell Institute for Digital Agriculture Executive Committee — 2022–2023

Cornell Institute for Digital Agriculture Hackathon Co-chair — 2022