

STEFANO ERMON

Department of Computer Science
Gates Building 228
Stanford University
Stanford, CA 94305-9040

Phone: (650) 498-9942
Email: ermon@cs.stanford.edu
<http://cs.stanford.edu/~ermon>

EMPLOYMENT

Stanford University **November 2014 – Present**
Assistant Professor, Department of Computer Science
Fellow, Woods Institute for the Environment

Atlas AI **2018–Present**
Co-Founder, Chief Technical Advisor

EDUCATION

Cornell University **August 2008 – January 2015**
Ph.D. in Computer Science, Minor in Applied Mathematics
Thesis: Decision Making And Inference Under Limited Information And High Dimensionality
Advisors: Carla P. Gomes and Bart Selman

University of Padova **September 2003 – July 2008**
M.S. in Electrical Engineering, *summa cum laude*
B.S. in Electrical Engineering, *summa cum laude*

AWARDS AND HONORS

ISSNAF Young Investigator Award, 2020.

Sloan Research Faculty Fellowship, 2020.

Microsoft Research Faculty Fellowship, 2019.

Bloomberg Data Science Research Award, 2019.

AFOSR Young Investigator Award, 2019.

IJCAI Computers and Thought Award, 2018.

Considered the premier award for artificial intelligence researchers under the age of 35.

ONR Young Investigator Award, 2018.

Hellman Faculty Fellowship, 2018.

AWS Machine Learning Award, 2018.

AAAI Outstanding Paper Award, 31st AAAI Conference on Artificial Intelligence, 2017.

One paper selected out of 638 accepted papers (and 2,590 submissions).

White House Frontiers Conference, 2016.

Invited to present at the White House Frontiers Conference, hosted by the White House and attended by President Obama.

10 World Changing Ideas of 2016, Scientific American, 2016.

Each year, Scientific American selects ten of the main discoveries/innovations of the year. Our work on poverty mapping in Africa was selected for 2016.

First Prize, World Bank Big Data Innovation Challenge, 2017.

World Bank launched a global call to find big data solutions to address hunger and food security challenges. Our crop yield prediction model won the first place among 189 submissions from 45 countries.

Finalist, NVIDIA Global Impact Award, 2016.

This award recognizes groundbreaking work that addresses the world's most important social and humanitarian problems. Our work on poverty mapping was selected as a finalist.

First Prize, INFORMS Yield Prediction Challenge Competition, 2016.

First place out of more than 30 teams at a national competition held by the Institute for Operations Research and the Management Sciences (INFORMS).

NSF CAREER Award, 2017.

Sony Faculty Innovation Award, 2017.

Early Career Spotlight,

25th International Joint Conference on Artificial Intelligence (IJCAI), 2016.

Best Student Paper Award,

31st AAAI Conference on Artificial Intelligence (Computational Sustainability Track), 2017.

Best Paper Award Runner-up and **Best Student Paper Award**,

29th Conference on Uncertainty in Artificial Intelligence (UAI), 2013.

Best Student Paper Award,

16th International Conference on Principles and Practice of Constraint Programming (CP), 2010.

Best Poster at Earth Vision CVPR Workshop, 2017.

Best Poster at WCB ICML Workshop, 2017.

McMullen Fellowship, Cornell University, 2008–2009.

PROFESSIONAL ACTIVITIES

Editorial Board, Journal of Artificial Intelligence Research (JAIR)

Journal Reviewer, Nature, Journal of Machine Learning Research (JMLR), IEEE Transactions on Information Theory, IEEE Transactions on Knowledge and Data Engineering (TKDE), Annals of Mathematics and Artificial Intelligence, Journal of the Royal Society Interface, American Chemical Society (ACS) Combinatorial Science, APL Materials, Information Processing Letters, Journal of Environmental and Resource Economics, Science China Information Sciences

Program Committee, UAI 2019, ICML 2018, AAAI 2018, AAAI 2017, ICML 2017, UAI 2017, AISTATS 2017, AAAI 2016, ICML 2016, UAI 2016, IJCAI 2016, KDD 2016, AAAI 2015, IJCAI 2015, GCAI 2015, KDD 2015, UAI 2015, AAAI 2014, AAAI 2013, IJCAI 2013, AAAI 2012, AAAI 2011, IJCAI 2011

Steering Committee, ACM SIGPLAN Machine Learning & Programming Languages, Computer Vision for Global Challenges

Local Chair, UAI 2018

Area Chair, ICML 2020, UAI 2020, AAAI 2020, NEURIPS 2020, ICLR 2019, ICML 2019, NEURIPS 2019

Conference Reviewer, ICLR 2018, NIPS 2018, NIPS 2017, NIPS 2016, NIPS 2015, ASONAM 2015, NIPS 2014, NIPS 2013, CP 2013, SAT 2013, CP 2012, SAT 2012

Grant Reviewer, National Science Foundation, Department of Energy, Department of Justice, A*STAR (Singapore), ND Renewable Energy Program, Sao Paulo Research Foundation (Brazil)

CONFERENCE and JOURNAL PUBLICATIONS (refereed and archived)

- [1] Jihyeon Lee, Dylan Grosz, Sicheng Zeng, Burak Uzkent, Marshall Burke, David Lobell, **Stefano Ermon**
Predicting Livelihood Indicators from Crowdsourced Street Level Images
In *Proc. 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.

- [2] Kumar Ayush, Burak Uzkent, Marshall Burke, David Lobell, **Stefano Ermon**
Efficient Poverty Mapping from High Resolution Remote Sensing Images
In *Proc. 35th AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [3] Jiaming Song, **Stefano Ermon**.
Multi-label Contrastive Predictive Coding.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [4] Jonathan Kuck, Shuvam Chakraborty, Hao Tang, Rachel Luo, Jiaming Song, Ashish Sabharwal, **Stefano Ermon**.
Belief Propagation Neural Networks.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [5] Yang Song, **Stefano Ermon**.
Improved Techniques for Training Score-Based Generative Models.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [6] Chenlin Meng, Lantao Yu, Yang Song, Jiaming Song, **Stefano Ermon**.
Autoregressive Score Matching.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [7] Andy Shih, **Stefano Ermon**.
Probabilistic Circuits for Variational Inference in Discrete Graphical Models.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [8] Yusuke Tashiro, Yang Song, **Stefano Ermon**.
Diversity can be Transferred: Output Diversification for White- and Black-box Attacks.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [9] Albert Gu, Tri Dao, **Stefano Ermon**, Atri Rudra, Christopher Re.
HiPPO: Recurrent Memory with Optimal Polynomial Projections.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [10] Tianyu Pang, Kun Xu, Chongxuan Li, Yang Song, **Stefano Ermon**, Jun Zhu.
Efficient Learning of Generative Models via Finite-Difference Score Matching.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [11] Tianhe Yu, Garrett Thomas, Lantao Yu, **Stefano Ermon**, James Zou, Sergey Levine, Chelsea Finn, Tengyu Ma.
MOPO: Model-based Offline Policy Optimization.
In *Proc. 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [12] Lantao Yu, Yang Song, Jiaming Song, **Stefano Ermon**
Training Deep Energy-Based Models with f-Divergence Minimization
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.
- [13] Shengjia Zhao, Tengyu Ma, **Stefano Ermon**
Individual Calibration with Randomized Forecasting
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.
- [14] Jiaming Song, **Stefano Ermon**
Bridging the Gap Between f-GANs and Wasserstein GANs
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.
- [15] Kuno Kim, Yihong Gu, Jiaming Song, Shengjia Zhao, **Stefano Ermon**
Domain Adaptive Imitation Learning
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.
- [16] Kristy Choi, Aditya Grover, Trisha Singh, Rui Shu, **Stefano Ermon**
Fair Generative Modeling via Weak Supervision
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.
- [17] Rui Shu, Tung Nguyen, Yinlam Chow, Tuan Pham, Khoat Than, Mohammad Ghavamzadeh, **Stefano Ermon**, Hung Bui
Predictive Coding for Locally-Linear Control
In *Proc. 37th International Conference on Machine Learning (ICML)*, 2020.

- [18] Christopher Yeh, Anthony Perez, Anne Driscoll, George Azzari, Zhongyi Tang, David Lobell, **Stefano Ermon**, Marshall Burke
Using Publicly Available Satellite Imagery and Deep Learning to Understand Economic Well-Being in Africa
In *Nature Communications*, 11, 2583, 2020.
- [19] Chris Cundy, **Stefano Ermon**
Flexible Approximate Inference via Stratified Normalizing Flows.
In *Proc. 36th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2020.
- [20] Kumar Ayush, Burak Uz Kent, Marshall Burke, David Lobell, **Stefano Ermon**
Generating Interpretable Poverty Maps using Object Detection in Satellite Images
In *Proc. 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 2020.
- [21] Peter M. Attia, Aditya Grover, Norman Jin, Kristen A. Severson, Todor M. Markov, Yang-Hung Liao, Michael H. Chen, Bryan Cheong, Nicholas Perkins, Zi Yang, Patrick K. Herring, Muratahan Aykol, Stephen J. Harris, Richard D. Braatz, **Stefano Ermon**, William C. Chueh.
Closed-loop Optimization of Fast-Charging Protocols for Batteries with Machine Learning.
In *Nature*, 578, 397-402, 2020.
- [22] Burak Uz Kent, **Stefano Ermon**.
Learning When and Where to Zoom with Deep Reinforcement Learning
In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [23] Joseph Duris, Dylan Kennedy, Adi Hanuka, Jane Shtalenkova, Auralee Edelen, Adam Egger, Tyler Cope, Mitchell McIntire, **Stefano Ermon**, Daniel Ratner.
Bayesian Optimization of a Free-Electron Laser
In *Physical Review Letters*, 2020.
- [24] Chenlin Meng, Yang Song, Jiaming Song, **Stefano Ermon**.
Gaussianization Flows.
In *Proc. 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [25] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, **Stefano Ermon**.
Permutation Invariant Graph Generation via Score-Based Generative Modeling.
In *Proc. 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [26] Shengjia Zhao, Christopher Yeh, **Stefano Ermon**.
A Framework for Sample Efficient Interval Estimation with Control Variates.
In *Proc. 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
- [27] Yilun Xu, Shengjia Zhao, Jiaming Song, Russell Stewart, **Stefano Ermon**.
A Theory of Usable Information under Computational Constraints.
In *Proc. 8th International Conference on Learning Representations (ICLR)*, 2020.
- [28] Jiaming Song, **Stefano Ermon**.
Understanding the Limitations of Variational Mutual Information Estimators.
In *Proc. 8th International Conference on Learning Representations (ICLR)*, 2020.
- [29] Rui Shu, Yining Chen, Abhishek Kumar, **Stefano Ermon**, Ben Poole.
Weakly Supervised Disentanglement with Guarantees.
In *Proc. 8th International Conference on Learning Representations (ICLR)*, 2020.
- [30] Vishnu Sarukkai, Anirudh Jain, Burak Uz Kent, **Stefano Ermon**.
Cloud Removal from Satellite Images Using Spatiotemporal Generator Networks.
In *Proceedings of IEEE Winter Conference of Applications of Computer Vision (WACV)*, 2020.
- [31] Burak Uz Kent, Christopher Yeh, **Stefano Ermon**.
Efficient Object Detection in Large Images Using Deep Reinforcement Learning.
In *Proceedings of IEEE Winter Conference of Applications of Computer Vision (WACV)*, 2020.
- [32] Aditya Grover, Christopher Chute, Rui Shu, Zhangjie Cao, **Stefano Ermon**.
AlignFlow: Cycle Consistent Learning from Multiple Domains via Normalizing Flows.
In *Proc. 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.

- [33] Mike Wu, Kristy Choi, Noah Goodman, **Stefano Ermon**.
Meta-Amortized Variational Inference and Learning.
In *Proc. 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
- [34] Chi-Sing Ho, Neal Jean, Catherine A. Hogan, Lena Blackmon, Stefanie S. Jeffrey, Mark Holodniy, Niaz Banaei, Amr A. E. Saleh, **Stefano Ermon**, Jennifer Dionne.
Rapid Identification of Pathogenic Bacteria using Raman Spectroscopy and Deep Learning.
Nature Communications, 30 Oct 2019, Issue 10, Number 4927.
- [35] Yang Song, **Stefano Ermon**.
Generative Modeling by Estimating Gradients of the Data Distribution.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [36] Aditya Grover, Jiaming Song, Alekh Agarwal, Kenneth Tran, Ashish Kapoor, Eric Horvitz, **Stefano Ermon**.
Bias Correction of Learned Generative Models using Likelihood-Free Importance Weighting.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [37] Yang Song, Chenlin Meng, **Stefano Ermon**.
MintNet: Building Invertible Neural Networks with Masked Convolutions.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [38] Lantao Yu, Tianhe Yu, Chelsea Finn, **Stefano Ermon**.
Meta-Inverse Reinforcement Learning with Probabilistic Context Variables.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [39] Jonathan Kuck, Tri Dao, Hamid Rezaatofghi, Ashish Sabharwa, **Stefano Ermon**.
Approximating the Permanent by Sampling from Adaptive Partitions.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [40] Sawyer Birnbaum, Volodymyr Kuleshov, Zayd Enam, Pang Wei Koh, **Stefano Ermon**.
Temporal FiLM: Capturing Long-Range Sequence Dependencies with Feature-Wise Modulations.
In *Proc. 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [41] Carla Gomes, Thomas Dietterich, Christopher Barrett, Jon Conrad, Bistra Dilkina, **Stefano Ermon**, Fei Fang, Andrew Farnsworth, Alan Fern, Xiaoli Fern, Daniel Fink, Douglas Fisher, Alexander Flecker, Daniel Freund, Angela Fuller, John Gregoire, John Hopcroft, Steve Kelling, Zico Kolter, Warren Powell, Nicole Sintov, John Selker, Bart Selman, Daniel Sheldon, David Shmoys, Milind Tambe, Weng-Keen Wong, Christopher Wood, Xiaojian Wu, Yexiang Xue, Amulya Yadav, Abdul-Aziz Yakubu, Mary Lou Zeeman.
Computational Sustainability: Computing for a Better World and a Sustainable Future.
In *Communications of the ACM*, September 2019, Vol. 62 No. 9, Pages 56-65.
- [42] Yang Song, Sahaj Garg, Jiabin Shi, **Stefano Ermon**.
Sliced Score Matching: A Scalable Approach to Density and Score Estimation.
In *Proc. 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.
- [43] Jonathan Kuck, Tri Dao, Shenjia Zhao, Burak Bartan, Ashish Sabharwal, **Stefano Ermon**.
Adaptive Hashing for Model Counting.
In *Proc. 35th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2019.
- [44] Burak Uz kent, Evan Sheehan, Chenlin Meng, Zhongyi Tang, David Lobell, Marshall Burke, **Stefano Ermon**.
Learning to Interpret Satellite Images using Wikipedia.
In *Proc. 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- [45] Michael Xie, **Stefano Ermon**.
Reparameterizable Subset Sampling via Continuous Relaxations.
In *Proc. 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- [46] Evan Sheehan, Chenlin Meng, Matthew Tan, Burak Uz kent, Neal Jean, David Lobell, Marshall Burke, **Stefano Ermon**.
Predicting Economic Development using Geolocated Wikipedia Articles.
In *Proc. 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2019.

- [47] Kristy Choi, Kedar Tatwawadi, Aditya Grover, Tsachy Weissman, **Stefano Ermon**.
Neural Joint-Source Channel Coding.
In *Proc. 36th International Conference on Machine Learning (ICML)*, 2019.
- [48] Aditya Grover, Aaron Zweig, **Stefano Ermon**.
Iterative Deep Generative Modeling of Large Graphs.
In *Proc. 36th International Conference on Machine Learning (ICML)*, 2019.
- [49] Lantao Yu, Jiaming Song, **Stefano Ermon**.
Multi-Agent Adversarial Inverse Reinforcement Learning.
In *Proc. 36th International Conference on Machine Learning (ICML)*, 2019.
- [50] Ali Malik, Volodymyr Kuleshov, Jiaming Song, Danny Nemer, Harlan Seymour, **Stefano Ermon**
Calibrated Model-Based Deep Reinforcement Learning.
In *Proc. 36th International Conference on Machine Learning (ICML)*, 2019.
- [51] Hongyu Ren, Shengjia Zhao, **Stefano Ermon**.
Adaptive Antithetic Sampling for Variance Reduction.
In *Proc. 36th International Conference on Machine Learning (ICML)*, 2019
- [52] Aditya Grover, Eric Wang, Aaron Zweig, **Stefano Ermon**.
Stochastic Optimization of Sorting Networks via Continuous Relaxations.
In *Proc. 7th International Conference on Learning Representations (ICLR)*, 2019.
- [53] Jun-Ting Hsieh, Shengjia Zhao, Stephan Eismann, Lucia Mirabella, **Stefano Ermon**.
Learning Neural PDE Solvers with Convergence Guarantees.
In *Proc. 7th International Conference on Learning Representations (ICLR)*, 2019.
- [54] Mike Wu, Noah Goodman, **Stefano Ermon**.
Differentiable Antithetic Sampling for Variance Reduction in Stochastic Variational Inference.
In *Proc. 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [55] Aditya Grover, **Stefano Ermon**.
Variational Compressive Sensing using Uncertainty Autoencoders.
In *Proc. 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [56] Rui Shu, Hung Bui, Jay Whang, **Stefano Ermon**.
Training Variational Autoencoders with Buffered Stochastic Variational Inference.
In *Proc. 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [57] Jiaming Song, Pratyusha Kalluri, Aditya Grover, Shengjia Zhao, **Stefano Ermon**.
Learning Controllable Fair Representations.
In *Proc. 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- [58] Neal Jean, Sherrie Wang, Anshul Samar, George Azzari, David Lobell, **Stefano Ermon**.
Tile2Vec: Unsupervised representation learning for spatially distributed data.
In *Proc. 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [59] Shengjia Zhao, Jiaming Song, **Stefano Ermon**.
InfoVAE: Balancing Learning and Inference in Variational Autoencoders.
In *Proc. 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [60] Jian Wei Khor, Neal Jean, Eric S Luxenberg, **Stefano Ermon**, Sindy K Y Tang.
Using Machine Learning to Discover Shape Descriptors for Predicting Emulsion Stability in a Microfluidic Channel.
In *Soft Matter*, 2018.
- [61] Shengjia Zhao, Hongyu Ren, Arianna Yuan, Jiaming Song, Noah Goodman, **Stefano Ermon**.
Bias and Generalization in Deep Generative Models: An Empirical Study.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2018.
- [62] Rui Shu, Hung Bui, Shengjia Zhao, Mykel Kochenderfer, **Stefano Ermon**.
Amortized Inference Regularization.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2018.

- [63] Neal Jean, Michael Xie, **Stefano Ermon**.
Semi-supervised Deep Kernel Learning: Regression with Unlabeled Data by Minimizing Predictive Variance.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems* (NeurIPS), 2018.
- [64] Jiaming Song, Hongyu Ren, Dorsa Sadigh, **Stefano Ermon**.
Multi-Agent Generative Adversarial Imitation Learning.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems* (NeurIPS), 2018.
- [65] Yang Song, Rui Shu, Nate Kushman, **Stefano Ermon**.
Constructing Unrestricted Adversarial Examples with Generative Models.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems* (NeurIPS), 2018.
- [66] Aditya Grover, Tudor Achim, **Stefano Ermon**.
Streamlining constraints for random k-SAT.
In *Proc. 32nd Annual Conference on Neural Information Processing Systems* (NeurIPS), 2018.
- [67] Yang Song, Jiaming Song, **Stefano Ermon**.
Accelerating Natural Gradient with Higher-Order Invariance.
In *Proc. 35th International Conference on Machine Learning* (ICML), 2018.
- [68] Volodymyr Kuleshov, Nathan Fenner, **Stefano Ermon**.
Calibrated Estimates of Predictive Uncertainty in Deep Learning.
In *Proc. 35th International Conference on Machine Learning* (ICML), 2018.
- [69] Manik Dhar, Aditya Grover, **Stefano Ermon**.
Sparse-Gen: Modeling Sparse Deviations for Compressed Sensing using Generative Models.
In *Proc. 35th International Conference on Machine Learning* (ICML), 2018.
- [70] Shengjia Zhao, Jiaming Song, **Stefano Ermon**.
A Lagrangian Perspective on Latent Variable Generative Models.
In *Proc. 34th Conference on Uncertainty in Artificial Intelligence* (UAI), 2018.
- [71] Stephan Eissman, Daniel Levy, Rui Shu, Stefan Bartzsch, **Stefano Ermon**.
Bayesian Optimization and Attribute Adjustment.
In *Proc. 34th Conference on Uncertainty in Artificial Intelligence* (UAI), 2018.
- [72] Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, **Stefano Ermon**.
Adversarial Constraint Learning for Structured Prediction.
In *Proc. 27th International Joint Conference on Artificial Intelligence* (IJCAI), 2018.
- [73] Barak Oshri, Annie Hu, Peter Adelson, Xiao Chen, Pascaline Dupas, Jeremy Weinstein, Marshall Burke, David Lobell, **Stefano Ermon**.
Infrastructure Quality Assessment in Africa using Satellite Imagery and Deep Learning.
In *Proc. 24th ACM SIGKDD Conference* (KDD), 2018.
- [74] Lijie Fan, Wenbing Huang, Chuang Gan, **Stefano Ermon**, Boqing Gong, Junzhou Huang.
End-to-End Motion Representation Learning for Video Understanding.
In *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2018.
- [75] Yang Song, Taesup Kim, Sebastian Nowozin, **Stefano Ermon**, Nate Kushman.
PixelDefend: Leveraging Generative Models to Understand and Defend against Adversarial Examples.
In *Proc. Proc. Sixth International Conference on Learning Representations* (ICLR), 2018.
- [76] Rui Shu, Hirokazu Narui, Hung Bui, **Stefano Ermon**.
A DIRT-T Approach to Unsupervised Domain Adaptation.
In *Proc. Sixth International Conference on Learning Representations* (ICLR), 2018.
- [77] Aditya Grover, Ramki Gummadi, Miguel Lazaro-Gredilla, Dale Schuurmans, **Stefano Ermon**.
Resampled Proposal Distributions for Variational Inference and Learning.
In *Proc. International Conference on Artificial Intelligence and Statistics* (AISTATS), 2018.
- [78] Aditya Grover, Todor Markov, Norman Jin, Peter Attia, Nick Perkins, Bryan Cheong, Michael Chen, Zi Yang, Stephen Harris, William Chueh, **Stefano Ermon**.
Best arm identification in multi-armed bandits with delayed and partial feedback.
In *Proc. International Conference on Artificial Intelligence and Statistics* (AISTATS), 2018.

- [79] Aditya Grover, Manik Dhar, **Stefano Ermon**.
Flow-GAN: Combining maximum likelihood and adversarial learning in generative models.
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- [80] Aditya Grover, **Stefano Ermon**.
Boosted Generative Models.
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- [81] Jonathan Kuck, **Stefano Ermon**.
Approximate Inference via Weighted Rademacher Complexity.
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- [82] Daniel Levy, **Stefano Ermon**.
Deterministic Policy Optimization by Combining Pathwise and Score Function Estimators for Discrete Action Spaces.
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- [83] Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, **Stefano Ermon**.
Learning with weak supervision from physics and data-driven constraints.
In *AI Magazine*, Spring 2018, Vol 39, No 1, pp. 27-38
- [84] William Gent, Kipil Lim, Yufeng Liang, Qinghao Li, Taylor Barnes, Sung-Jin Ahn, Kevin Stone, Mitchell McIntire, Jihyun Hong, Jay Hyok Song, Yiyang Li, Apurva Mehta, **Stefano Ermon**, Tolek Tyliczszak, Arthur Kilcoyne, David Vine, Jin-Hwan Park, Seok-Gwang Doo, Michael Toney, Wanli Yang, David Prendergast, and William Chueh.
Coupling Between Oxygen Redox and Cation Migration Explains Unusual Electrochemistry in Lithium-Rich Layered Oxides.
In *Nature Communications*.
- [85] Volodymyr Kuleshov, **Stefano Ermon**.
Neural Variational Inference and Learning in Undirected Graphical Models.
In *Proc. 31st Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [86] Jiaming Song, Shengjia Zhao, **Stefano Ermon**.
A-NICE-MC: Adversarial Training for MCMC.
In *Proc. 31st Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [87] Yunzhu Li, Jiaming Song, **Stefano Ermon**.
InfoGAIL: Interpretable Imitation Learning from Visual Demonstrations.
In *Proc. 31st Annual Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [88] Siamak Yousefi, Hirokazu Narui, Sankalp Dayal, **Stefano Ermon**, Shahrokh Valaee.
A Survey on Behaviour Recognition Using WiFi Channel State Information.
In *IEEE Communications Magazine*, 2017.
- [89] Stephen Mussmann, Daniel Levy, **Stefano Ermon**.
Fast Amortized Inference and Learning in Log-linear Models with Randomly Perturbed Nearest Neighbor Search.
In *Proc. 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- [90] Volodymyr Kuleshov, **Stefano Ermon**.
Deep Hybrid Models: Bridging Discriminative and Generative Approaches.
In *Proc. 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- [91] Shengjia Zhao, Jiaming Song, **Stefano Ermon**.
Learning Hierarchical Features from Generative Models.
In *Proc. 34th International Conference on Machine Learning (ICML)*, 2017.
- [92] Russell Stewart, **Stefano Ermon**.
Supervising Neural Networks with Physics and other Domain Knowledge.
In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
AAAI Outstanding Paper Award.
- [93] Volodymyr Kuleshov, **Stefano Ermon**.
Online Uncertainty Estimation Against an Adversary.
In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.

- [94] Jiaxuan You, Xiaocheng Li, Melvin Low, David Lobell, **Stefano Ermon**.
Deep Gaussian Process for Crop Yield Prediction Based on Remote Sensing Data.
In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
Best Student Paper Award (Computational Sustainability Track).
- [95] Colin Wei, **Stefano Ermon**.
General Bounds on Satisfiability Thresholds for Random CSPs via Fourier Analysis.
In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [96] Biagio Cosenza, Juan Durillo, **Stefano Ermon**, Ben Juurlink.
Autotuning Stencil Computations with Structural Ordinal Regression Learning.
In *Proc. IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2017.
- [97] Neal Jean, Marshall Burke, Michael Xie, Matthew Davis, David Lobell, **Stefano Ermon**.
Combining Satellite Imagery and Machine Learning to Predict Poverty.
In *Science*, 353(6301), 790-794, 2016.
Nature Research Highlight.
- [98] Xiaoyue Duan, Feifei Yang, Erin Antono, Wenge Yang, Piero Pianetta, **Stefano Ermon**, Apurva Mehta, and Yijin Liu.
Unsupervised Data Mining in Nanoscale X-ray Spectro-Microscopic Study of NdFeB Magnet.
In *Scientific Reports*, 6, 34406 (2016).
- [99] Aditya Grover, **Stefano Ermon**.
Variational Bayes on Monte Carlo Steroids.
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [100] Shengjia Zhao, Enze Zhou, Ashish Sabharwal, **Stefano Ermon**.
Adaptive Concentration Inequalities for Sequential Decision Problems.
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [101] Jonathan Ho, **Stefano Ermon**.
Generative Adversarial Imitation Learning.
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [102] Yexiang Xue, Zhiyuan Li, **Stefano Ermon**, Carla Gomes, Bart Selman.
Solving Marginal MAP Problems with NP Oracles and Parity Constraints.
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [103] Mitchell McIntire, Daniel Ratner, **Stefano Ermon**.
Sparse Gaussian Processes for Bayesian Optimization.
In *Proc. 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- [104] Jonathan Ho, Jayesh Gupta, **Stefano Ermon**.
Model-Free Imitation Learning with Policy Optimization.
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [105] Yexiang Xue, **Stefano Ermon**, Ronan Le Bras, Carla Gomes, Bart Selman.
Variable Elimination in the Fourier Domain.
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [106] Steve Mussmann, **Stefano Ermon**.
Learning and Inference via Maximum Inner Product Search.
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [107] Tudor Achim, Ashish Sabharwal, **Stefano Ermon**.
Beyond Parity Constraints: Fourier Analysis of Hash Functions for Inference.
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [108] Lun-Kai Hsu, Tudor Achim, **Stefano Ermon**.
Tight Variational Bounds via Random Projections and I-Projections.
In *Proc. 19th Conference on Artificial Intelligence and Statistics (AISTATS)*, 2016.
- [109] Michael Xie, Neal Jean, Marshall Burke, David Lobell, **Stefano Ermon**.
Transfer Learning from Deep Features for Remote Sensing and Poverty Mapping.
In *Proc. 30th AAAI Conference on Artificial Intelligence (AAAI)*, 2016.

- [110] Shengjia Zhao, Sorathan Chaturapruek, Ashish Sabharwal, **Stefano Ermon**.
Closing the Gap Between Short and Long XORs for Model Counting.
In *Proc. 30th AAAI Conference on Artificial Intelligence (AAAI)*, 2016.
- [111] Carolyn Kim, Ashish Sabharwal, **Stefano Ermon**.
Exact Sampling with Integer Linear Programs and Random Perturbations.
In *Proc. 30th AAAI Conference on Artificial Intelligence (AAAI)*, 2016.
- [112] Stefan Hadjis and **Stefano Ermon**.
Importance sampling over sets: a new probabilistic inference scheme.
In *Proc. 31st Conference on Uncertainty in Artificial Intelligence (UAI)*, 2015.
- [113] Michael Zhu and **Stefano Ermon**.
A Hybrid Approach for Probabilistic Inference using Random Projections.
In *Proc. 32nd International Conference on Machine Learning (ICML)*, 2015.
- [114] Yexiang Xue, **Stefano Ermon**, Carla Gomes, Bart Selman.
Uncovering Hidden Structure through Parallel Problem Decomposition for the Set Basis Problem with Application to Materials Discovery.
In *Proc. 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2015.
- [115] **Stefano Ermon**, Yexiang Xue, Russell Toth, Bistra Dilkina, Richard Bernstein, Theodoros Damoulas, Patrick Clark, Steve DeGloria, Andrew Mude, Christopher Barrett, and Carla Gomes.
Learning Large Scale Dynamic Discrete Choice Models of Spatio-Temporal Preferences with Application to Migratory Pastoralism in East Africa.
In *Proc. 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [116] **Stefano Ermon**, Ronan Le Bras, Santosh Suram, John M. Gregoire, Carla Gomes, Bart Selman, and Robert B. van Dover.
Pattern Decomposition with Complex Combinatorial Constraints: Application to Materials Discovery.
In *Proc. 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2015.
- [117] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Designing Fast Absorbing Markov Chains.
In *Proc. 28th AAAI Conference on Artificial Intelligence (AAAI)*, 2014.
- [118] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Low-density Parity Constraints for Hashing-Based Discrete Integration.
In *Proc. 31st International Conference on Machine Learning (ICML)*, 2014.
- [119] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Embed and Project: Discrete Sampling with Universal Hashing.
In *Proc. 27th Annual Conference on Neural Information Processing Systems (NIPS)*, 2013.
- [120] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Optimization with Parity Constraints: from Binary Codes to Discrete Integration.
In *Proc. 29th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2013.
Best student paper award. Best paper award runner-up.
- [121] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Taming the Curse of Dimensionality: Discrete Integration by Hashing and Optimization.
In *Proc. 30th International Conference on Machine Learning (ICML)*, 2013.
- [122] **Stefano Ermon**, Yexiang Xue, Carla Gomes, and Bart Selman.
Learning Policies For Battery Usage Optimization in Electric Vehicles.
In *Machine Learning*, Volume 92, Issue 1 (2013), Page 177-194.
- [123] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Density Propagation and Improved Bounds on the Partition Function.
In *Proc. 26th Annual Conference on Neural Information Processing Systems (NIPS)*, 2012.
- [124] **Stefano Ermon**, Carla Gomes, and Bart Selman.
Uniform Solution Sampling Using a Constraint Solver As an Oracle.
In *Proc. 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2012.

- [125] Liaoruo Wang, **Stefano Ermon**, and John Hopcroft.
Feature-Enhanced Probabilistic Models for Diffusion Network Inference.
In *Proc. of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2012.
- [126] **Stefano Ermon**, Yexiang Xue, Carla Gomes, and Bart Selman.
Learning Policies For Battery Usage Optimization in Electric Vehicles.
In *Proc. of European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD)*, 2012.
- [127] **Stefano Ermon**, Ronan Le Bras, Carla Gomes, Bart Selman, and Bruce van Dover.
SMT-Aided Combinatorial Materials Discovery.
In *Proc. 15th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 2012.
- [128] **Stefano Ermon**, Carla Gomes, Bart Selman, and Alexander Vladimirsky.
Probabilistic Planning with Nonlinear Utility Functions and Worst-Case Constraints.
In *Proc. 11th Intl. Conf. Autonomous Agents and Multiagent Systems (AAMAS)*, 2012.
- [129] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.
Accelerated Adaptive Markov Chain for Partition Function Computation.
In *Proc. 25th Annual Conference on Neural Information Processing Systems (NIPS)*, 2011.
- [130] **Stefano Ermon**, Carla Gomes, and Bart Selman.
A Flat Histogram Method for Computing the Density of States of Combinatorial Problems.
In *Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- [131] **Stefano Ermon**, Jon Conrad, Carla Gomes, and Bart Selman.
Risk-Sensitive Policies for Sustainable Renewable Resource Allocation.
In *Proc. 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.
- [132] **Stefano Ermon**, Carla Gomes, and Bart Selman.
A Message Passing Approach to Multiagent Gaussian Inference for Dynamic Processes.
In *Proc. 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (short paper), 2011.
- [133] **Stefano Ermon**, Carla Gomes, and Bart Selman.
Computing the Density of States of Boolean Formulas.
In *Proc. 16th International Conference on Principles and Practice of Constraint Programming (CP)*, 2010. **Best student paper award.**
- [134] **Stefano Ermon**, Jon Conrad, Carla Gomes, and Bart Selman.
Playing Games against Nature: Optimal Policies for Renewable Resource Allocation.
In *Proc. 26th Conference on Uncertainty in Artificial Intelligence (UAI)*, 2010.
- [135] **Stefano Ermon**, Carla Gomes, and Bart Selman.
Collaborative Multiagent Gaussian Inference in a Dynamic Environment Using Belief Propagation.
In *Proc. 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (short paper), 2010.

RECENT INVITED TALKS

NeurIPS Deep Learning and Inverse Problems Workshop	December, 2020
Harvard Center for Computation and Society Seminar	September, 2020
CVPR Agri-Vision Workshop	June, 2020
IAS Special Year Seminar Series on Theoretical Machine Learning	May, 2020
New Directions in Optimization, Statistics and Machine Learning	April, 2020
Oregon State University AI Seminar	March, 2020
AAAS Annual Meeting	February, 2020
ICME AI for Good	January, 2020

MIT ML Seminar	October, 2019
Deep Learning Indaba	August, 2019
AI4Vietnam	August, 2019
IJCAI AI for Social Good Workshop	August, 2019
KDD Earth Day	August, 2019
ICML I3 Workshop	June, 2019
Stanford Computer Forum Workshop	April, 2019
UPenn Applied Mathematics and Computational Science Colloquium	March, 2019
Toyota Research Institute	March, 2019

GRANT SUPPORT

Department of Energy (co-PI) INSIEME: Integrated Simulations using Exascale Multiphysics Ensembles	2020–2025
Sloan Research Fellowship (PI)	2020–2021
Samsung (PI) Detecting and Handling Anomalies with Robust Machine Learning Systems	2020–2021
Air Force Office of Scientific Research (PI) YIP: Variational Methods for Information Processing and Learning	2019–2022
Office of Naval Research (PI) YIP: Learning with Domain Knowledge: An Implicit Probabilistic Models Approach	2019–2023
Toyota Research Institute (PI) Emergent Behaviors in Multi-Agent Human-Robot Teams	2018–2020
Hellman Faculty Fellowship (PI) Weakly Supervised Machine Learning with Constraints and Prior Knowledge	2018–2019
Stanford Catalyst for Collaborative Solution (co-PI) A Microbial Culture Shift: Rapid Screening of Bacterial Bloodstream Infections and Antibiotic Susceptibility for Urban and Rural Clinical Care	2017–2019
National Science Foundation (PI) CAREER: Modeling and Inference for Large Scale Spatio-Temporal Data	2017–2022
National Science Foundation (co-PI) Efficient High-Dimensional Integration using Error-Correcting Codes	2017–2020
Schmidt Foundation (gift)	2017
Sony (PI) Learning Without Labels: Providing Supervision through Domain Knowledge	2017–2019
Siemens (PI) Generative Modeling for Structured Data Types	2017–2018
Siemens (PI) The Use of Physics and Domain Knowledge in Training Probabilistic Prediction Models from Time Series Data	2017–2018
Toyota (co-PI) High through-put Synthesis and Screening of Perovskite Electro-Catalysts for Advanced Fuel Cells	2017–2022
Toyota (co-PI) Data-Driven Design of Li-Ion Batteries	2017–2022
Intel (PI) Discovering Hidden Structure in High Dimensional Data for Anomaly Detection	2017–2019

National Science Foundation (PI) Empowering Probabilistic Reasoning with Random Projections	2016–2017
Toyota (PI) Dynamic Choice Models of Human Decision-Making	2016–2018
Ford Motor Company (PI) Inference Models for Autonomous Vehicles	2016–2018
Stanford Global Development and Poverty Initiative (PI) Closing the Data Divide: Machine Learning Approaches for Understanding Livelihoods of the Poor using Unconventional Data Sources	2016–2018
Future of Life Institute (PI) Robust Probabilistic Inference Engines for Autonomous Agents	2015–2018
Stanford Precourt Institute for Energy (co-PI) Nanoimaging-Inspired Battery Management System for Electric Vehicles: Translating Insights on Nanoscale Dynamics to Control Algorithms	2017–2019
National Science Foundation (subcontract) CompSustNet: Expanding the Horizons of Computational Sustainability	2015–2020
Swiss Re (gift)	2016

BROADER IMPACT and OUTREACH

Media Coverage

My research has been covered by numerous general and scientific press outlets including: BBC, CNBC, NY Times, Washington Post, Los Angeles Times, Nature Research Highlights, Science Magazine News.

A more equitable future, Repubblica (August 29, 2018); To Help The Poorest Of The Poor, First You Have To Find Them, NPR (December 6, 2016); Satellites Map World Poverty, Nature Research Highlights (August 24, 2016); How satellite images are helping find the worlds hidden poor, Washington Post (August 24, 2016); Artificial Intelligence Could Help End Poverty Worldwide, The Inquisitor News (August 20, 2016); How to track poverty from space, Los Angeles Times (August 19, 2016); Machine learning poverty map could help aid get to the right places in Africa, Yahoo News (August 19, 2016); The best way to predict poverty is by combining satellite images with machine learning, Quartz News (August 19, 2016); Satellite images can map poverty, Science Magazine News (August 18, 2016); Satellite images used to predict poverty, BBC News (August 18, 2016); Artificial intelligence can find, map poverty, researchers say, Business Insider (August 18, 2016); Fighting Poverty With Satellite Images and Machine Learning Wizardry, IEEE Spectrum (August 18, 2016); Satellite images of Earth help us predict poverty better than ever, The Verge (August 18, 2016); Stanford Scientists Combine Satellite Data, Machine Learning to Map Poverty, Stanford University News (August 18, 2016); Crop Analytics Challenge Offers Future Help for Seed Selection, AG Professional News (April 28, 2016); Stanford University Team Wins Syngenta Crop Challenge in Analytics, Farm Futures (April 26, 2016); Syngenta Crop Challenge in Analytics Winners Announced, PR Newswire (April 18, 2016); Five Finalists Compete for Nvidia 2016 Global Impact Award This Week, Fudzilla (April 4, 2016); Satellite Images Can Pinpoint Poverty Where Surveys Can't, New York Times (April 1, 2016); What are Some Recent Advances in Non-Convex Optimization Research?, Huffington Post (March 29, 2016); How GPUs are Helping Map Worldwide Poverty, Nvidia Press (March 25, 2016); Syngenta Crop Challenge Finalists to Use Analytics for Farm Seed Selection, Crop Protection News (March 14, 2016); Syngenta Crop Challenge Finalists Use Advanced Analytics to Optimize Farmer Seed Selection, PR Newswire (February 29, 2016); Machine Scans Millions of Satellite Images to Map Poverty, Futurity: Research News (February 25, 2016); Stanford Researchers Use Dark of Night and Machine Learning to Shed Light on Global Poverty, Stanford University News (February 24, 2016); Domsday Grants Will Advance Important AI Research, MIT Technology Review (July 2, 2015);

Outreach

AI4All: Mentor and lecturer for a nonprofit organization working to increase diversity and inclusion in AI <http://ai-4-all.org/>.

IBM Watson AI XPRIZE: Member of the judging panel for a \$5 million global competition to develop AI technologies to tackle the world's grand challenges.

Lifeboat Foundation: Advisory board member.

Machine Learning for Materials Research: Organizer of the first bootcamp in 2016.

Online Educational Materials: I am a strong proponent of open access to educational materials. I have developed online lecture notes for my Probabilistic Graphical Models (CS228) and Automated Reasoning (CS323), accessed by over **10,000 unique visitors**.

Outreach Talks: I regularly give talks to general audiences (high schools, student and professional associations, NGOs), on average 3 or 4 times a year.

Research Blog: I run a research blog with several thousand visitors per month <https://ermongroup.github.io/blog/>.

TEACHING EXPERIENCE

Instructor , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	Winter 2020-2021
Instructor , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	Winter 2019-2020
Instructor , Department of Computer Science, Stanford University CS326: Deep Generative Models	Fall 2019-2020
Instructor , Department of Computer Science, Stanford University CS325b: Data for Development	Fall 2019-2020
Instructor , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	Winter 2018-2019
Instructor , Department of Computer Science, Stanford University CS326: Deep Generative Models	Fall 2018-2019
Instructor , Department of Computer Science, Stanford University CS325b: Data for Development	Fall 2018-2019
Instructor , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	Winter 2017-2018
Instructor , Department of Computer Science, Stanford University CS325b: Data for Development	Winter 2017-2018
Instructor , Department of Computer Science, Stanford University CS221: Artificial Intelligence: Principles and Techniques	Fall 2017-2018
Instructor , Department of Computer Science, Stanford University CS325b: Data for Development	Fall 2017-2018
Instructor , Department of Computer Science, Stanford University CS323: Automated Reasoning	Spring 2016-2017
Instructor , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	Winter 2016-2017
Instructor , Department of Computer Science, Stanford University CS325: Topics in Computational Sustainability	Spring 2015-2016
Instructor , Department of Computer Science, Stanford University	Winter 2015-2016

CS228: Probabilistic Graphical Models

Instructor, Department of Computer Science, Stanford University
CS323: Automated Reasoning

Spring 2014-2015

Instructor, Department of Computer Science, Stanford University
CS228: Probabilistic Graphical Models

Winter 2014-2015

Instructor, Department of Computer Science, Cornell University
CS1130: Transition to Object-oriented Programming

Fall 2011