

# STEFANO ERMON

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

**Phone:** (650) 498-9942  
**Email:** [ermon@cs.stanford.edu](mailto: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

## EDUCATION

**Cornell University** **August 2008 – January 2015**  
Ph.D. in Computer Science, Minor in Applied Mathematics, January 2015  
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*, July 2008  
B.S. in Electrical Engineering, *summa cum laude*, July 2006

## AWARDS AND HONORS

**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,** 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,** 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,** MAPL

**Local Chair,** UAI 2018

**Area Chair,** ICLR 2019, ICML 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)

## CONFERENCE and JOURNAL PUBLICATIONS (refereed and archived)

- [1] Shengjia Zhao, Hongyu Ren, Arianna Yuan, Jiaming Song, Noah Goodman, **Stefano Ermon**. Bias and Generalization in Deep Generative Models: An Empirical Study. To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [2] Rui Shu, Hung Bui, Shengjia Zhao, Mykel Kochenderfer, **Stefano Ermon**. Amortized Inference Regularization. To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [3] Neal Jean, Michael Xie, **Stefano Ermon**. Semi-supervised Deep Kernel Learning: Regression with Unlabeled Data by Minimizing Predictive Variance. To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [4] Jiaming Song, Hongyu Ren, Dorsa Sadigh, **Stefano Ermon**. Multi-Agent Generative Adversarial Imitation Learning. To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.

- [5] Yang Song, Rui Shu, Nate Kushman, **Stefano Ermon**.  
Constructing Unrestricted Adversarial Examples with Generative Models.  
To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [6] Aditya Grover, Tudor Achim, **Stefano Ermon**.  
Streamlining constraints for random k-SAT.  
To appear in *Proc. 32nd Annual Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [7] Yang Song, Jiaming Song, **Stefano Ermon**.  
Accelerating Natural Gradient with Higher-Order Invariance.  
In *Proc. 35th International Conference on Machine Learning (ICML)*, 2018.
- [8] Volodymyr Kuleshov, Nathan Fenner, **Stefano Ermon**.  
Calibrated Estimates of Predictive Uncertainty in Deep Learning.  
In *Proc. 35th International Conference on Machine Learning (ICML)*, 2018.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] Aditya Grover, **Stefano Ermon**.  
Boosted Generative Models.  
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.

- [21] Jonathan Kuck, **Stefano Ermon**.  
Approximate Inference via Weighted Rademacher Complexity.  
In *Proc. 32nd AAAI Conference on Artificial Intelligence (AAAI)*, February 2018.
- [22] 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.
- [23] 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
- [24] 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*.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] 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.
- [29] 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.
- [30] Volodymyr Kuleshov, **Stefano Ermon**.  
Deep Hybrid Models: Bridging Discriminative and Generative Approaches.  
In *Proc. 33rd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017.
- [31] Shengjia Zhao, Jiaming Song, **Stefano Ermon**.  
Learning Hierarchical Features from Generative Models.  
In *Proc. 34th International Conference on Machine Learning (ICML)*, 2017.
- [32] 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**.
- [33] Volodymyr Kuleshov, **Stefano Ermon**.  
Online Uncertainty Estimation Against an Adversary.  
In *Proc. 31st AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [34] 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).

- [35] 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.
- [36] 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.
- [37] 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.**
- [38] 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).
- [39] Aditya Grover, **Stefano Ermon**.  
Variational Bayes on Monte Carlo Steroids.  
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [40] 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.
- [41] Jonathan Ho, **Stefano Ermon**.  
Generative Adversarial Imitation Learning.  
In *Proc. 30th Annual Conference on Neural Information Processing Systems (NIPS)*, 2016.
- [42] 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.
- [43] Mitchell McIntire, Daniel Ratner, **Stefano Ermon**.  
Sparse Gaussian Processes for Bayesian Optimization.  
In *Proc. 32nd Conference on Uncertainty in Artificial Intelligence (UAI)*, 2016.
- [44] Jonathan Ho, Jayesh Gupta, **Stefano Ermon**.  
Model-Free Imitation Learning with Policy Optimization.  
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [45] 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.
- [46] Steve Mussmann, **Stefano Ermon**.  
Learning and Inference via Maximum Inner Product Search.  
In *Proc. 33rd International Conference on Machine Learning (ICML)*, 2016.
- [47] 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.
- [48] 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.
- [49] 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.
- [50] 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.

- [51] 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.
- [52] 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.
- [53] Michael Zhu and **Stefano Ermon**.  
A Hybrid Approach for Probabilistic Inference using Random Projections.  
In *Proc. 32nd International Conference on Machine Learning (ICML)*, 2015.
- [54] 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.
- [55] **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.
- [56] **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.
- [57] **Stefano Ermon**, Carla Gomes, Ashish Sabharwal, and Bart Selman.  
Designing Fast Absorbing Markov Chains.  
In *Proc. 28th AAAI Conference on Artificial Intelligence (AAAI)*, 2014.
- [58] **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.
- [59] **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.
- [60] **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.**
- [61] **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.
- [62] **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.
- [63] **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.
- [64] **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.
- [65] 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.

- [66] **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.
- [67] **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.
- [68] **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.
- [69] **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.
- [70] **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.
- [71] **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.
- [72] **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.
- [73] **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.**
- [74] **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.
- [75] **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

Cornell AI Seminar	September, 2018
IJCAI 2018 Computers and Thought Award	July, 2018
ICML 2018 Workshop on Tractable Probabilistic Models	July, 2018
IJCAI 2018 Tutorial on Deep Generative Models	July, 2018
EPFL CS Seminar	June, 2018
Tech4Dev Data Science and Machine Learning	June, 2018
CVPR 2018 Tutorial on Generative adversarial networks	June, 2018
CVPR 2018 DeepGlobe Workshop	June, 2018
Stanford Data Science Initiative Workshop	May, 2018
Carnegie Mellon University Institute for Software Engineering Seminar	April, 2018
Big Data Analytics San Francisco	February, 2018

## GRANT SUPPORT

<b>Office of Naval Research (PI)</b> YIP: Learning with Domain Knowledge: An Implicit Probabilistic Models Approach	<b>2018–2021</b>
<b>Toyota Research Institute (PI)</b> Emergent Behaviors in Multi-Agent Human-Robot Teams	<b>2018–2020</b>
<b>Hellman Faculty Fellowship (PI)</b> Weakly Supervised Machine Learning with Constraints and Prior Knowledge	<b>2018–2019</b>
<b>Stanford Catalyst for Collaborative Solution (co-PI)</b> A Microbial Culture Shift: Rapid Screening of Bacterial Bloodstream Infections and Antibiotic Susceptibility for Urban and Rural Clinical Care	<b>2017–2019</b>
<b>National Science Foundation (PI)</b> CAREER: Modeling and Inference for Large Scale Spatio-Temporal Data	<b>2017–2022</b>
<b>National Science Foundation (co-PI)</b> Efficient High-Dimensional Integration using Error-Correcting Codes	<b>2017–2020</b>
<b>Schmidt Foundation (gift)</b>	<b>2017</b>
<b>Sony (PI)</b> Learning Without Labels: Providing Supervision through Domain Knowledge	<b>2017–2019</b>
<b>Siemens (PI)</b> Generative Modeling for Structured Data Types	<b>2017–2018</b>
<b>Siemens (PI)</b> The Use of Physics and Domain Knowledge in Training Probabilistic Prediction Models from Time Series Data	<b>2017–2018</b>
<b>Toyota (co-PI)</b> High through-put Synthesis and Screening of Perovskite Electro-Catalysts for Advanced Fuel Cells	<b>2017–2022</b>
<b>Toyota (co-PI)</b> Data-Driven Design of Li-Ion Batteries	<b>2017–2022</b>
<b>Intel (PI)</b> Discovering Hidden Structure in High Dimensional Data for Anomaly Detection	<b>2017–2019</b>
<b>National Science Foundation (PI)</b> Empowering Probabilistic Reasoning with Random Projections	<b>2016–2017</b>
<b>Toyota (PI)</b> Dynamic Choice Models of Human Decision-Making	<b>2016–2018</b>
<b>Ford Motor Company (PI)</b> Inference Models for Autonomous Vehicles	<b>2016–2018</b>
<b>Stanford Global Development and Poverty Initiative (PI)</b> Closing the Data Divide: Machine Learning Approaches for Understanding Livelihoods of the Poor using Unconventional Data Sources	<b>2016–2018</b>
<b>Future of Life Institute (PI)</b> Robust Probabilistic Inference Engines for Autonomous Agents	<b>2015–2018</b>
<b>Stanford Precourt Institute for Energy (co-PI)</b> Nanoimaging-Inspired Battery Management System for Electric Vehicles: Translating Insights on Nanoscale Dynamics to Control Algorithms	<b>2017–2019</b>
<b>National Science Foundation (subcontract)</b> CompSustNet: Expanding the Horizons of Computational Sustainability	<b>2015–2020</b>
<b>Swiss Re (gift)</b>	<b>2016</b>



## BROADER IMPACT and OUTREACH

### Media Coverage

My research has been covered by numerous general and scientific press outlets including: BBC, 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); Doomsday 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  
CS326: Deep Generative Models

**Fall 2018-2019**

<b>Instructor</b> , Department of Computer Science, Stanford University CS325b: Data for Development	<b>Fall 2018-2019</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	<b>Winter 2017-2018</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS325b: Data for Development	<b>Winter 2017-2018</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS221: Artificial Intelligence: Principles and Techniques	<b>Fall 2017-2018</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS325b: Data for Development	<b>Fall 2017-2018</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS323: Automated Reasoning	<b>Spring 2016-2017</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	<b>Winter 2016-2017</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS325: Topics in Computational Sustainability	<b>Spring 2015-2016</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	<b>Winter 2015-2016</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS323: Automated Reasoning	<b>Spring 2014-2015</b>
<b>Instructor</b> , Department of Computer Science, Stanford University CS228: Probabilistic Graphical Models	<b>Winter 2014-2015</b>
<b>Instructor</b> , Department of Computer Science, Cornell University CS1130: Transition to Object-oriented Programming	<b>Fall 2011</b>