

Jiaming Song

NVIDIA Email jiamings@nvidia.com
 Homepage: <https://tsong.me>
 Phone: (650) 495-6322

Education

Stanford University, Master of Science & Doctor of Philosophy 2016 - 2021
Computer Science Department
Advisor: Stefano Ermon
Thesis title: *Compression, Generation, and Inference via Supervised Learning*

Tsinghua University, Bachelor of Engineering 2012 - 2016
Department of Computer Science and Technology
Graduated with Outstanding Honor (Top 1%)

Professional Experiences

NVIDIA, Research Scientist Jun 2022 - now

Stanford University, Postdoctoral Scholar Sep 2021 - Jun 2022

Stanford University, Research Assistant Sep 2016 - Sept 2021

Facebook AI Research, Research Intern Jun 2018 - Sep 2018

OpenAI, Research Intern Jun 2017 - Sep 2017

Megvii, Research Intern Feb 2016 - Jun 2016

Information Initiative, Duke University, Visiting Researcher Jun 2015 - Sep 2015

Awards and Honors

ICLR 2022 Outstanding Paper Award 2022
For paper on “Comparing Distributions by Measuring Differences that Affect Decision Making”.

Qualcomm Innovation Fellowship (8 in total) 2018
For project on “Safe Multi-Agent Imitation Learning for Self-Driving”.

Qualcomm Scholarship (Top 1%) 2016
For Tsinghua undergraduates with exceptional research experiences.

Google Excellence Scholarship 2015
Awarded to 58 undergraduate and graduate students in China.

Outstanding Winner, Interdisciplinary Contest in Modeling (Top 0.3%) 2015
Highest award, for the paper “Organizational Churn: A Roll of the Dice?”.

Outstanding Undergraduate, China Computer Federation 2014
Awarded to 2 undergraduate students in Tsinghua University.

Zhong Shimo Scholarship (Top 0.75%) 2013
Highest scholarship in the CS Department in Tsinghua.

Publications

Refereed Conference and Journal Publications

- [47] Lantao Yu, Tianhe Yu, **Jiaming Song**, Willie Neiswanger, Stefano Ermon
Offline Imitation Learning with Suboptimal Demonstrations via Relaxed Distribution Matching
AAAI Conference on Artificial Intelligence, (AAAI 2023)
- [46] Bahjat Kawar*, **Jiaming Song***, Stefano Ermon, Michael Elad
JPEG Artifact Correction using Denoising Diffusion Restoration Models
Neural Information Processing Systems (NeurIPS) Workshop on Score-Based Methods, (NeurIPS 2022 SBM Workshop)
- [45] Bahjat Kawar, Michael Elad, Stefano Ermon, **Jiaming Song**
Denoising Diffusion Restoration Models
Neural Information Processing Systems, (NeurIPS 2022)
- [44] Chenlin Meng*, Kristy Choi*, **Jiaming Song**, Stefano Ermon
Concrete Score Matching: Generalized Score Matching for Discrete Data
Neural Information Processing Systems, (NeurIPS 2022)
- [43] Divyansh Garg, Sakanda Vaidyanath, Kuno Kim, **Jiaming Song**, Stefano Ermon
LISA: Learning Interpretable Skill Abstractions from Language
Neural Information Processing Systems, (NeurIPS 2022)
- [42] **Jiaming Song***, Lantao Yu*, Willie Neiswanger, Stefano Ermon
A General Recipe for Likelihood-free Bayesian Optimization
International Conference on Machine Learning, (ICML 2022), *Long oral presentation (Top 2.2%)*
- [41] Samarth Sinha*, **Jiaming Song***, Animesh Garg, Stefano Ermon
Experience Replay with Likelihood-free Importance Weights
4th Annual Conference on Learning for Dynamics and Control, (L4DC 2022), *Best paper award finalist*
- [40] Chenlin Meng, Yutong He, Yang Song, **Jiaming Song**, Jiajun Wu, Jun-Yan Zhu, Stefano Ermon
SDEdit: Image Synthesis and Editing with Stochastic Differential Equations
International Conference on Learning Representations, (ICLR 2022)
- [39] Shengjia Zhao, Abhishek Sinha, Yutong He, Aidan Perreault, **Jiaming Song**, Stefano Ermon
Comparing Distributions by Measuring Differences that Affect Decision Making
International Conference on Learning Representations, (ICLR 2022), *ICLR 2022 Outstanding Paper Award*
- [38] Chenlin Meng, Enci Liu, Willie Neiswanger, **Jiaming Song**, Marshall Burke, David Lobell, Stefano Ermon
IS-COUNT: Large-scale Object Counting from Satellite Images with Covariate-based Importance Sampling
AAAI Conference on Artificial Intelligence, (AAAI 2022)

- [37] Abhishek Sinha*, **Jiaming Song***, Chenlin Meng, Stefano Ermon
D2C: Diffusion-Denoising Models for Few-shot Conditional Generation
 Neural Information Processing Systems, (**NeurIPS 2021**)
- [36] Divyansh Garg, Shuvam Chakraborty, Chris Cundy, **Jiaming Song**, Stefano Ermon
IQ-Learn: Inverse soft-Q Learning for Imitation
 Neural Information Processing Systems, (**NeurIPS 2021**), *Spotlight presentation*
- [35] Lantao Yu, **Jiaming Song**, Yang Song, Stefano Ermon
Pseudo-Spherical Contrastive Divergence
 Neural Information Processing Systems, (**NeurIPS 2021**)
- [34] Yusuke Tashiro, **Jiaming Song**, Yang Song, Stefano Ermon
CSDI: Conditional Score-based Diffusion Models for Probabilistic Time Series Imputation
 Neural Information Processing Systems, (**NeurIPS 2021**)
- [33] Jiayu Chen, Yuanxin Zhang, Yuanfan Xu, Huimin Ma, Huazhong Yang, **Jiaming Song**, Yu Wang, Yi Wu
Variational Automatic Curriculum Learning for Sparse-Reward Cooperative Multi-Agent Problems
 Neural Information Processing Systems, (**NeurIPS 2021**)
- [32] Kuno Kim, Akshat Jindal, Yang Song, **Jiaming Song**, Yanan Sui, Stefano Ermon
Imitation with Neural Density Models
 Neural Information Processing Systems, (**NeurIPS 2021**)
- [31] **Jiaming Song**, Chenlin Meng, Stefano Ermon
Denoising Diffusion Implicit Models
 International Conference on Learning Representations, (**ICLR 2021**)
- [30] Abhishek Sinha*, Ayush Kumar*, **Jiaming Song***, Burak Ukzent, Hongxia Jin, Stefano Ermon
Negative Data Augmentation
 International Conference on Learning Representations, (**ICLR 2021**)
- [29] Chenlin Meng, **Jiaming Song**, Yang Song, Shengjia Zhao, Stefano Ermon
Improved Autoregressive Modeling with Distribution Smoothing
 International Conference on Learning Representations, (**ICLR 2021**), *Oral presentation*
- [28] **Jiaming Song**, Stefano Ermon
Multi-label Contrastive Predictive Coding
 Neural Information Processing Systems, (**NeurIPS 2020**), *Oral presentation*
- [27] Chenlin Meng, Lantao Yu, Yang Song, **Jiaming Song**, Stefano Ermon
Autoregressive Score Matching
 Neural Information Processing Systems, (**NeurIPS 2020**)
- [26] Jonathan Kuck, Shuvam Chakraborty, Hao Tang, Rachel Luo, **Jiaming Song**, Ashish Sabharwal, Stefano Ermon
Belief Propagation Neural Networks
 Neural Information Processing Systems, (**NeurIPS 2020**)
- [25] **Jiaming Song**, Michael Auli, Yann Dauphin, Tengyu Ma
Robust and On-the-fly Dataset Denoising for Image Classification
 European Conference on Computer Vision, (**ECCV 2020**)
- [24] Chenhao Niu, Yang Song, **Jiaming Song**, Shengjia Zhao, Aditya Grover, Stefano Ermon

- Permutation Invariant Graph Generation via Score-Based Generative Modeling**
International Conference on Artificial Intelligence and Statistics, (AISTATS 2020)
- [23] Chenlin Meng, Yang Song, **Jiaming Song**, Stefano Ermon
Gaussianization Flows
International Conference on Artificial Intelligence and Statistics, (AISTATS 2020)
- [22] Lantao Yu, Yang Song, **Jiaming Song**, Stefano Ermon
Training Deep Energy-Based Models with f-Divergence Minimization
International Conference on Machine Learning, (ICML 2020)
- [21] **Jiaming Song**, Stefano Ermon
Bridging the Gap Between f-GANs and Wasserstein GANs
International Conference on Machine Learning, (ICML 2020)
- [20] Kuno Kim, Yihong Gu, **Jiaming Song**, Shengjia Zhao, Stefano Ermon
Domain Adaptive Imitation Learning
International Conference on Machine Learning, (ICML 2020)
- [19] **Jiaming Song**, Stefano Ermon
Understanding the Limitations of Variational Mutual Information Estimators
International Conference on Learning Representations, (ICLR 2020)
- [18] Yilun Xu, Shengjia Zhao, **Jiaming Song**, Russell Stewart, Stefano Ermon
A Theory of Usable Information under Computational Constraints
International Conference on Learning Representations, (ICLR 2020), *Oral presentation*
- [17] Nate Gruver, **Jiaming Song**, Mykel Kochenderfer, Stefano Ermon
Multi-agent Adversarial Inverse Reinforcement Learning with Latent Variables
International Conference on Autonomous Agents and MultiAgent Systems (extended abstract), (AAMAS 2020)
- [16] Aditya Grover, **Jiaming Song**, Ashish Kapoor, Kenneth Tran, Alekh Agarwal, Eric Horvitz, Stefano Ermon
Bias Correction of Learned Generative Models using Likelihood-free Importance Weighting
Advances in Neural Information Processing Systems, (NeurIPS 2019)
- [15] Ali Malik, Volodymyr Kuleshov, **Jiaming Song**, Danny Nemer, Harlan Seymour, Stefano Ermon
Calibrated Model-based Deep Reinforcement Learning
International Conference on Machine Learning, (ICML 2019)
- [14] Lantao Yu, **Jiaming Song**, Stefano Ermon
Multi-agent Adversarial Inverse Reinforcement Learning
International Conference on Machine Learning, (ICML 2019)
- [13] Shengjia Zhao, **Jiaming Song**, Stefano Ermon
InfoVAE: Balancing Learning and Inference in Variational Autoencoders
AAAI Conference on Artificial Intelligence, (AAAI 2019)
- [12] **Jiaming Song**, Pratyusha Kalluri, Aditya Grover, Shengjia Zhao, Stefano Ermon
Learning Controllable Fair Representations
International Conference on Artificial Intelligence and Statistics, (AISTATS 2019)
- [11] **Jiaming Song**, Hongyu Ren, Dorsa Sadigh, Stefano Ermon
Multi-Agent Generative Adversarial Imitation Learning
Advances in Neural Information Processing Systems, (NeurIPS 2018)

- [10] Shengjia Zhao, Hongyu Ren, Arianna Yuan, **Jiaming Song**, Noah Goodman, Stefano Ermon
Bias and Generalization in Deep Generative Models: An Empirical Study
Advances in Neural Information Processing Systems, (**NeurIPS 2018**), *Spotlight presentation*
- [9] Shengjia Zhao, **Jiaming Song**, Stefano Ermon
The Information Autoencoding Family: A Lagrangian Perspective on Latent Variable Generative Models
Conference on Uncertainty in Artificial Intelligence, (**UAI 2018**), *Oral presentation*
- [8] Yang Song, **Jiaming Song**, Stefano Ermon
Accelerating Natural Gradient with Higher-Order Invariance
International Conference on Machine Learning, (**ICML 2018**)
- [7] Hongyu Ren, Russell Stewart, **Jiaming Song**, Volodymyr Kuleshov, Stefano Ermon
Adversarial Constraint Learning for Structured Prediction
International Joint Conference on Artificial Intelligence, (**IJCAI 2018**)
- [6] Hongyu Ren, Russell Stewart, **Jiaming Song**, Volodymyr Kuleshov, Stefano Ermon
Learning with weak supervision from physics and data-driven constraints
AI Magazine
- [5] **Jiaming Song**, Shengjia Zhao, Stefano Ermon
A-NICE-MC: Adversarial training for MCMC
Advances in Neural Information Processing Systems, (**NeurIPS 2017**)
- [4] Shengjia Zhao, **Jiaming Song**, Stefano Ermon
Learning Hierarchical Features from Deep Generative Models
International Conference on Machine Learning, (**ICML 2017**)
- [3] Yunzhu Li, **Jiaming Song**, Stefano Ermon
InfoGAIL: Interpretable imitation learning from visual demonstrations
Advances in Neural Information Processing Systems, (**NeurIPS 2017**)
- [2] Bei Chen, Ning Chen, Jun Zhu, **Jiaming Song**, Bo Zhang
Discriminative nonparametric latent feature relational models with data augmentation
AAAI Conference on Artificial Intelligence, (**AAAI 2016**)
- [1] **Jiaming Song**, Zhe Gan, Lawrence Carin
Factored Temporal Sigmoid Belief Networks for Sequence Learning
International Conference on Machine Learning, (**ICML 2016**)

Preprints and Technical Reports

- [7] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, **Jiaming Song**, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, Ming-Yu Liu
eDiff-I: Text-to-Image Diffusion Models with Ensemble of Expert Denoisers
arXiv preprint arXiv:2211.01324
- [6] Xuan Su, **Jiaming Song**, Chenlin Meng, Stefano Ermon
Dual Diffusion Implicit Bridges for Image-to-Image Translation
arXiv preprint arXiv:2203.08382
- [5] Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Shoeb, Abubakar Abid, Adam Fisch, Adam Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, et al.

Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models

arXiv preprint arXiv:2206.04615

- [4] Rachel Luo, Shengjia Zhao, **Jiaming Song**, Jonathan Kuck, Stefano Ermon, Silvio Savarese
Privacy Preserving Recalibration under Domain Shift
arXiv:2008.09643
- [3] **Jiaming Song**, Yang Song, Stefano Ermon
Unsupervised Out-of-Distribution Detection with Batch Normalization
arXiv:1910.09115
- [2] Shengjia Zhao, **Jiaming Song**, Stefano Ermon
Towards deeper understanding of variational autoencoding models
arXiv:1702.08658
- [1] Jun Zhu, **Jiaming Song**, Bei Chen
Max-margin Nonparametric Latent Feature Models for Link Prediction
arXiv:1602.07428

Teaching

<i>Stanford CS228: Probabilistic Graphical Models</i> TA and Lecturer on <i>Markov Chain Monte Carlo</i>	2020
<i>Stanford CS236: Deep Generative Models</i> Teaching Assistant	2018

Professional Activities

Journal Reviewer

Journal of Machine Learning Research (JMLR)
IEEE Transactions on Pattern Recognition and Machine Intelligence (TPAMI)
Journal of Artificial Intelligence Research (JAIR)
IEEE Transactions on Information Theory (TIT)
ACM Transactions on Intelligent Systems and Technology (TIST)

Conference Reviewer / Program Committee

<i>International Conference on Machine Learning (ICML)</i>	2019 - 2022
<i>Neural Information Processing Systems (NeurIPS)</i>	2019 - 2022
<i>International Conference on Learning Representations (ICLR)</i>	2018 - 2023
<i>Conference on Uncertainty in Artificial Intelligence (UAI)</i>	2019 - 2022
<i>International Conference on Artificial Intelligence and Statistics (AISTATS)</i>	2023
<i>Conference on Learning Theory (COLT)</i>	2019
<i>Conference on Computer Vision and Pattern Recognition (CVPR)</i>	2019 - 2023
<i>European Conference on Computer Vision (ECCV)</i>	2020, 2022
<i>International Conference on Computer Vision (ICCV)</i>	2019, 2021
<i>Winter Conference on Applications of Computer Vision (WACV)</i>	2021, 2022

<i>AAAI Conference on Artificial Intelligence (AAAI)</i>	2021, 2022
<i>International Joint Conference on Artificial Intelligence (IJCAI)</i>	2021
<i>Asian Conference on Machine Learning (ACML)</i>	2018, 2019
<i>Bay Area Machine Learning Symposium</i>	2018 - 2020

Workshop Organization

<i>Workshop on Information Theory and Machine Learning</i>	NeurIPS 2019
<i>Generative Models for Reinforcement Learning</i>	DALI 2018

Outreach

<i>Ermon Group Blog, Co-creator</i>	2017 - 2022
<i>Stanford CURIS program for undergraduate research, Mentor</i>	2019, 2020
<i>NeurIPS session for researchers of color, Mentor</i>	2018
<i>Stanford AI undergraduate mentoring program, Mentor</i>	2018
<i>Women in Machine Learning (WiML), Mentor</i>	2017
<i>Global NeurIPS Paper Implementation Challenge, Mentor</i>	2017