

Jiaming Song

✉ tsong@cs.stanford.edu

🌐 tsong.me

Education

2016 – **Stanford University**, Palo Alto, CA.
Ph.D. Program in Computer Science. Advisor: [Stefano Ermon](#).

2012 – 2016 **Tsinghua University (THU)**, Beijing, China.
B.Eng. in Computer Science and Technology. Graduated with Outstanding Honor (Top 1%).

Publications

- December 2018 **Multi-Agent Generative Adversarial Imitation Learning**
Jiaming Song, Hongyu Ren, Dorsa Sadigh, Stefano Ermon.
To appear in *the 31th Neural Information Processing Systems (NIPS 2018)*.
Abridged version in *the 6th International Conference on Learning Representations (ICLR 2018) Workshop Track, 1st Workshop on Goal Specifications for Reinforcement Learning*.
- December 2018 **Bias and Generalization in Deep Generative Models: An Empirical Study**
Shengjia Zhao, Hongyu Ren, Arianna Yuan, Jiaming Song, Noah Goodman, Stefano Ermon.
To appear in *the 31th Neural Information Processing Systems (NIPS 2018)*. Spotlight Presentation.
- August 2018 **Learning with Weak Supervision from Physics and Data-Driven Constraints**
Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, Stefano Ermon.
In *AI Magazine* 39(1): 27-38.
- August 2018 **A Lagrangian Perspective on Latent Variable Generative Models**
Shengjia Zhao, Jiaming Song, Stefano Ermon.
In *the 2018 Conference on Uncertainty in Artificial Intelligence (UAI 2018)*. Oral Presentation.
Abridged version in *2018 Bay Area Machine Learning Symposium, ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- July 2018 **Accelerating Natural Gradient with Higher-Order Invariance**
Yang Song, Jiaming Song, Stefano Ermon.
In *the 35th International Conference on Machine Learning (ICML 2018)*.
- July 2018 **Adversarial Constraint Learning for Structured Prediction**
Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, Stefano Ermon.
In *2018 International Joint Conference on Artificial Intelligence (IJCAI 2018)*.
Abridged version in *NIPS Workshop on Learning with Limited Data*.
- December 2017 **A-NICE-MC: Adversarial Training for MCMC**
Jiaming Song, Shengjia Zhao, Stefano Ermon.
In *the 30th Neural Information Processing Systems (NIPS 2017)*.
Abridged version in *ICML 2017 Workshop on Implicit Models*.
- December 2017 **InfoGAIL: Interpretable Imitation Learning from Visual Demonstration**
Yunzhu Li, Jiaming Song, Stefano Ermon.
In *the 30th Neural Information Processing Systems (NIPS 2017)*.
- August 2017 **Learning Hierarchical Features from Generative Models**
Shengjia Zhao, Jiaming Song, Stefano Ermon.
In *the 34th International Conference on Machine Learning (ICML 2017)*.

- June 2016 **Factored Sigmoid Belief Networks for Sequence Learning**
Jiaming Song, Zhe Gan, Lawrence Carin.
In *33rd International Conference on Machine Learning (ICML 2016)*.
- June 2015 **Organizational Churn: A Roll of the Dice?**
Canyao Liu*, Jiaming Song*, Chuan Yu*.
In *Undergraduate Mathematics and its Applications (UMAP), Issue 36.2*. Invited paper.
- June 2015 **Discriminative Nonparametric Latent Feature Relational Models with Data Augmentation**
Bei Chen, Ning Chen, Jun Zhu, Jiaming Song, Bo Zhang.
In *10th Association for the Advancement of Artificial Intelligence Conference (AAAI 2016)*.

Workshop Papers and Manuscripts

- July 2018 **InfoVAE: Information Maximizing Variational Autoencoders**
Shengjia Zhao, Jiaming Song, Stefano Ermon.
In *ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- July 2018 **Dual Optimization for Latent Variable Generative Models**
Shengjia Zhao*, Jiaming Song*, Stefano Ermon.
In *ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- July 2018 **Markov Chain Monte Carlo for Learning Belief Networks**
Laetitia Shao*, Jiaming Song*, Aditya Grover, Stefano Ermon.
In *ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- July 2018 **Learning Controllable Fair Representations via Latent Variable Generative Models**
Jiaming Song, Pratyusha Kalluri, Aditya Grover, Shengjia Zhao, Stefano Ermon.
In *ICML 2018 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.
- December 2017 **Structured Prediction with Adversarial Constraint Learning**
Hongyu Ren, Russell Stewart, Jiaming Song, Volodymyr Kuleshov, Stefano Ermon.
In *NIPS 2017 Workshop on Learning with Limited Data*.
- December 2017 **A Lagrangian Perspective on Latent Variable Generative Modeling**
Shengjia Zhao, Jiaming Song, Stefano Ermon.
In *NIPS 2017 Workshop on Bayesian Deep Learning*.
- December 2017 **An Empirical Study of the Generalization Behavior of Generative Adversarial Networks**
Hongyu Ren, Shengjia Zhao, Jiaming Song, Lijie Fan, Stefano Ermon.
In *NIPS 2017 Workshop on Deep Learning: Bridging Theory and Practice*.
- April 2017 **Generative Adversarial Learning of Markov Chains**
Jiaming Song, Shengjia Zhao, Stefano Ermon.
In *the 5th International Conference on Learning Representations (ICLR 2017) Workshop*.
- February 2016 **Max-Margin Nonparametric Latent Feature Relational Models for Link Prediction**
Jun Zhu, Jiaming Song, Bei Chen.
In *arXiv preprint arXiv:1602.07428*.
- February 2017 **Towards Deeper Understanding of Variational Autoencoding Models**
Shengjia Zhao, Jiaming Song, Stefano Ermon.
In *arXiv preprint arXiv:1702.08658*.

Professional Experiences

- June 2018 - **Research Intern**, Facebook AI Research. Mentors: Michael Auli, Yann Dauphin and Tengyu Ma.
Sept 2018 Research on deep learning and optimization.

- June 2017 - **Research Intern**, OpenAI. Mentors: Rocky Duan and John Schulman.
Sept 2017 Research on deep reinforcement learning.
- April 2016 - **Detection, Tracking and Reidentification Group**, Megvii Inc. Mentor: Chi Zhang
July 2016 Developed a scalable framework to provide distant supervision for unlabeled data, which allows model distillation and merging network structures for different tasks, such as detection and parsing.
[Megvii Inc.](#) is a leading unicorn start-up in China, with emphasis on machine learning and computer vision.
- July 2015 - **Information Initiative @ Duke (iiD)**, Duke University. Advisor: Prof. [Lawrence Carin](#).
September 2015 Worked on conditional factored deep generative models using recent Neural Variational Inference methods, which allows for semi-supervised deep learning and sequence generation with side information.
- November 2014 **Statistical AI & Learning (TSAIL) Group**, Tsinghua University. Advisor: Prof. [Jun Zhu](#).
- June 2015 Explored stochastic variational methods for link prediction problems. Proposed an efficient method that would train on a network with over 3 million nodes, a significant improvement over original methods.
- July 2014 - **Visual Computing Group**, Microsoft Research Asia. Advisor: [Jingdong Wang](#).
October 2014 Implemented a convolutional neural network for multiple label image annotation with [Caffe](#).

Honors and Awards

- June 2018 **Qualcomm Innovation Fellowship (US) Winner**, issued by Qualcomm.
Safe Multi-Agent Imitation Learning for Self-Driving (top 4.5%)
- June 2016 **Qualcomm Scholarship**, issued by Qualcomm.
Offered to Tsinghua undergraduates with exceptional research experiences (top 1%).
- June 2015 **Google Excellence Scholarship**, issued by Google.
This scholarship is offered to Chinese undergraduate and graduate students who possess remarkable academic achievements and project experiences. 58 students are selected nationwide (6 in Tsinghua University).
- April 2015 **Outstanding Winner**, Interdisciplinary Contest in Modeling 2015.
Highest award (9 out of 2317) of the contest. Published a paper which models organizational churn using Bayesian-inspired methods and network science. See github.com/jiamings/icm2015 for more details.
- April 2015 **Third Prize**, 33rd Tsinghua Challenge Cup, issued by Tsinghua University.
Our project implements fast, scalable video segmentation and classification which utilizes deep activation features. Please see jiamings.github.io/projects/decaf-video for details.
- October 2014 **Outstanding Undergraduate**, issued by the China Computer Federation (CCF).
Only 4 students in Tsinghua, and 100 in China are awarded each year.
- May 2014 **Spark Program for Technological Innovation**, Tsinghua University.
Among top 50/3000 students for achievements in scientific and technological innovations.
- December 2013 **Zhong Shimo Scholarship**, issued by Dept. of Computer Science and Technology.
Highest scholarship in the CS Department. (top 0.75%)
- July 2011 **Bronze Prize, National Olympiad in Informatics**, issued by China Computer Federation (CCF).

Professional Services and Outreach

- Reviewer ICLR 2018, BayLearn 2018, ACML 2018
- Fall 2017 - Now **Ermon Group Blog**, Maintaining an academic blog post (<http://ermongroup.github.io/blog/>), with several thousand visitors per month.
- April 2018 **Data Learning and Inference 2018**, Organizer
Organizing the Generative Models for Reinforcement Learning workshop.

December 2017 **Women in Machine Learning (WiML) 2017**, Mentor.
Invited to discussion panel on “A-NICE-MC: Adversarial Training for MCMC”.

December 2017 **Global NIPS Paper Implementation Challenge**, Mentor.
Provide guidance to reproduce results in NIPS 2017 papers.

Teaching

Fall 2018 CS236, Deep Generative Models, Stanford University