Simon Shaolei Du

Paul G. Allen School of Computer Science & Engineering

ssdu@cs.washington.edu

http://simonshaoleidu.com

Research Interests

Machine Learning, Learning Theory, Deep Learning, Reinforcement Learning, Bandits, Representation Learning, Non-convex Optimization, Stochastic Optimization

Education

Carnegie Mellon University 2015 - 2019

- Ph.D in Machine Learning, School of Computer Science
- Advisors: Barnabás Póczos, Aarti Singh
- Thesis: Gradient Descent for Non-convex Problems in Modern Machine Learning

University of California, Berkeley 2011 - 2015

- B.S. in Electrical Engineering and Computer Science
- B.S. in Engineering Mathematics and Statistics
- Advisors: Ming Gu, Lei Li, Michael Mahoney, Stuart Russell

Tsinghua University Spring 2013

- Exchange student at electronic engineering department

Professional Experiences

University of Washington, School of Computer Science & Engineering *2020 - Present*

- *Assistant Professor*

Facebook AI Reseaerch *2021 - 2022*

- *Visiting Professor*

Institute for Advanced Study *2019 - 2020*

- *Postdoctoral Research Fellow* (Mentor: Sanjeev Arora)

Google Brain *Summer 2019*

- *Research Intern* (Mentors: Elad Hazan, Sham M. Kakade)

Simons Institute for the Theory of Computing *Fall 2018*

- *Visiting Student* (Mentor: Michael Mahoney)

Microsoft Research Labs *Summer 2018*

- *Research Intern* (Mentors: Alekh Agarwal, Miro Dudík, Nan Jiang, Akshay Krishnamurthy, John Langford)

Facebook AI Research *Summer 2017*

- *Research Intern* (Mentor: Yuandong Tian)

Microsoft Research Labs *Summer 2016*

- *Research Intern* (Mentors: Jianshu Chen, Lihong Li, Lin Xiao, Denny Zhou)

Accenture *Summer 2015*

- *Management Consulting Intern*

Google *Summer 2014*

- *Software Engineering Intern*

Awards/Honors

Alfred P. Sloan Research Fellowship *2024*
Awarded to 22 early-career scholars in computer science in the US and Canada.

Intel Rising Star Faculty Award *2023*
Awarded to fifteen early-career professors worldwide.

Samsung AI Researcher of the Year *2022*
Awarded to five early-career researchers in AI worldwide.

National Science Foundation CAREER Award *2022*
Supports early-career faculty who have the potential to serve as academic role models.

AAAI New Faculty Highlights *2021*
Awarded to eighteen new faculties working in AI.

Tencent AI Rhino-Bird Award *2021*

CMU School of Computer Science Distinguished Dissertation Award (Runner Up) *2019*
Top five disserations from CMU School of Computer Science.

Nomination for ACM Doctoral Dissertation Award *2019*
Two dissertations were selected from CMU.

Nomination for Joint AAAI/ACM SIGAI Doctoral Dissertation Award *2019*
One dissertation was selected from CMU.

Late-breaking Paper Award for NeurIPS Deep Reinforcement Learning Workshop *2019*

Nvidia Pioneer Award *2018*

Best Paper Award for ICML Workshop on Nonconvex Optimization *2018*

UC Berkeley EECS Honor Program *2014*

Phi Beta Kappa *2014*

Conference Publications

* denotes equal contribution or alphabetical ordering.

1. Nuoya Xiong, Lijun Ding, Simon S. Du
   How Over-Parameterization Slows Down Gradient Descent in Matrix Sensing: The Curses of Symmetry and Initialization
   In International Conference on Learning Representations (ICLR) 2024
   (spotlight presentation, 5% acceptance rate)
2. Zihan Zhang, Jason D. Lee, Yuxin Chen, Simon S. Du
   Horizon-Free Regret for Linear Markov Decision Processes
   In International Conference on Learning Representations (ICLR) 2024
   (31% acceptance rate)
3. Yuandong Tian, Yiping Wang, Zhenyu Zhang, Beidi Chen, Simon Du
   JoMA: Demystifying Multilayer Transformers via JOint Dynamics of MLP and Attention
   In International Conference on Learning Representations (ICLR) 2024
   (31% acceptance rate)
4. Zhaoyi Zhou, Chuning Zhu, Runlong Zhou, Qiwen Cui, Abhishek Gupta, Simon S. Du
   Free from Bellman Completeness: Trajectory Stitching via Model-based Return-conditioned Supervised Learning
   In International Conference on Learning Representations (ICLR) 2024
   (31% acceptance rate)
5. Haozhe Jiang, Qiwen Cui, Zhihan Xiong, Maryam Fazel, Simon S. Du
   A Black-box Approach for Non-stationary Multi-agent Reinforcement Learning

In International Conference on Learning Representations (ICLR) 2024
(31% acceptance rate)

6. Kaifeng Lyu*, Jikai Jin*, Zhiyuan Li, Simon Shaolei Du, Jason D. Lee, Wei Hu
Dichotomy of Early and Late Phase Implicit Biases Can Provably Induce Grokking
In International Conference on Learning Representations (ICLR) 2024
(31% acceptance rate)
7. Ruizhe Shi*, Yuyao Liu*, Yanjie Ze, Simon S. Du, Huazhe Xu
Unleashing the Power of Pre-trained Language Models for Offline Reinforcement Learning
In International Conference on Learning Representations (ICLR) 2024
(31% acceptance rate)
8. Yifang Chen, Yingbing Huang, Simon S. Du, Kevin Jamieson, Guanya Shi
Active Representation Learning for General Task Space with Applications in Robotics
In Conference on Neural Information Processing Systems (NeurIPS) 2023
(26% acceptance rate)
9. Yuandong Tian, Yiping Wang, Beidi Chen, Simon S. Du
Scan and Snap: Understanding Training Dynamics and Token Composition in 1-layer Transformer
In Conference on Neural Information Processing Systems (NeurIPS) 2023
(26% acceptance rate)
10. Yunchang Yang*, Han Zhong*, Tianhao Wu*, Bin Liu, Liwei Wang, Simon S. Du
A Reduction-based Framework for Sequential Decision Making with Delayed Feedback
In Conference on Neural Information Processing Systems (NeurIPS) 2023
(26% acceptance rate)
11. Angela Yuan, Chris Junchi Li, Gauthier Gidel, Michael Jordan, Quanquan Gu, Simon S. Du
Optimal Extragradient-Based Algorithms for Stochastic Variational Inequalities with Separable Structure
In Conference on Neural Information Processing Systems (NeurIPS) 2023
(26% acceptance rate)
12. Weihang Xu, Simon S. Du
Over-Parameterization Exponentially Slows Down Gradient Descent for Learning a Single Neuron
In Annual Conference of Learning Theory (COLT) 2023
(34% acceptance rate)
13. Qiwen Cui, Kaiqing Zhang, Simon S. Du
Breaking the Curse of Multiagents in a Large State Space: RL in Markov Games with Independent Linear Function Approximation
In Annual Conference of Learning Theory (COLT) 2023
(34% acceptance rate)
14. Haotian Ye*, Xiaoyu Chen*, Liwei Wang, Simon S. Du
On the Power of Pre-training for Generalization in RL: Provable Benefits and Hardness
In International Conference on Machine Learning (ICML) 2023
(Oral presentation, 3% acceptance rate)
15. Runlong Zhou, Zihan Zhang, Simon S. Du
Sharp Variance-Dependent Bounds in Reinforcement Learning: Best of Both Worlds in Stochastic and Deterministic Environments
In International Conference on Machine Learning (ICML) 2023
(28% acceptance rate)
16. Jikai Jin, Zhiyuan Li, Kaifeng Lyu, Simon S. Du, Jason D. Lee
Understanding Incremental Learning of Gradient Descent: A Fine-grained Analysis of Matrix Sensing
In International Conference on Machine Learning (ICML) 2023
(28% acceptance rate)

17. Runlong Zhou, Ruosong Wang, Simon S. Du
Horizon-Free Reinforcement Learning for Latent Markov Decision Processes
In International Conference on Machine Learning (ICML) 2023
(28% acceptance rate)

18. Yiping Wang, Yifang Chen, Kevin Jamieson, Simon Shaolei Du
Improved Active Multi-Task Representation Learning via Lasso
In International Conference on Machine Learning (ICML) 2023
(28% acceptance rate)

19. Haozhe Jiang*, Qiwen Cui*, Zhihan Xiong, Maryam Fazel, Simon S. Du
Offline Congestion Games: How Feedback Type Affects Data Coverage Requirement
In International Conference on Learning Representations (ICLR) 2023
(32% acceptance rate)

20. Shicong Cen*, Yuejie Chi*, Simon S. Du*, Lin Xiao*
Faster Last-iterate Convergence of Policy Optimization in Zero-Sum Markov Games
In International Conference on Learning Representations (ICLR) 2023
(32% acceptance rate)

21. Rui Yuan, Simon S. Du, Robert M. Gower, Alessandro Lazaric, Lin Xiao
Linear Convergence of Natural Policy Gradient Methods with Log-Linear Policies
In International Conference on Learning Representations (ICLR) 2023
(32% acceptance rate)

22. Yan Dai, Ruosong Wang, Simon S. Du
Variance-Aware Sparse Linear Bandits
In International Conference on Learning Representations (ICLR) 2023
(32% acceptance rate)

23. Yulai Zhao, Jianshu Chen, Simon S. Du
Blessing of Class Diversity in Pre-training
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2023
(notable paper, 2% acceptance rate)

24. Qiwen Cui*, Zhihan Xiong*, Maryam Fazel, Simon S. Du
Learning in Congestion Games with Bandit Feedback
In Conference on Neural Information Processing Systems (NeurIPS) 2022
(26% acceptance rate)

25. Xinqi Wang, Qiwen Cui, Simon S. Du
On Gap-dependent Bounds for Offline Reinforcement Learning
In Conference on Neural Information Processing Systems (NeurIPS) 2022
(26% acceptance rate)

26. Qiwen Cui, Simon S. Du
Provably Efficient Offline Multi-agent Reinforcement Learning via Strategy-wise Bonus
In Conference on Neural Information Processing Systems (NeurIPS) 2022
(26% acceptance rate)

27. Zhihan Xiong*, Ruoqi Shen*, Qiwen Cui*, Maryam Fazel, Simon S. Du
Near-Optimal Randomized Exploration for Tabular Markov Decision Processes
In Conference on Neural Information Processing Systems (NeurIPS) 2022
(26% acceptance rate)

28. Rui Lu, Andrew Zhao, Simon S. Du, Gao Huang
Provable General Function Class Representation Learning in Multitask Bandits and MDPs
In Conference on Neural Information Processing Systems (NeurIPS) 2022
(spotlight presentation, 3% acceptance rate)

29. Qiwen Cui, Simon S. Du
When are Offline Two-Player Zero-Sum Markov Games Solvable?
In Conference on Neural Information Processing Systems (NeurIPS) 2022

(26% acceptance rate)

30. Zihan Zhang, Xiangyang Ji, Simon S. Du
Horizon-Free Reinforcement Learning in Polynomial Time: the Power of Stationary Policies
In Annual Conference of Learning Theory (COLT) 2022
(32% acceptance rate)

31. Yifang Chen, Simon S. Du, Kevin Jamieson
Active Multi-Task Representation Learning
In International Conference on Machine Learning (ICML) 2022
(22% acceptance rate)

32. Haoyuan Cai, Tengyu Ma, Simon S. Du
Near-Optimal Algorithms for Autonomous Exploration and Multi-Goal Stochastic Shortest Path
In International Conference on Machine Learning (ICML) 2022
(22% acceptance rate)

33. Andrew Wagenmaker, Yifang Chen, Max Simchowitz, Simon S. Du, Kevin Jamieson
Reward-Free RL is No Harder Than Reward-Aware RL in Linear Markov Decision Processes
In International Conference on Machine Learning (ICML) 2022
(22% acceptance rate)

34. Tongzhou Wang, Simon S. Du, Antonio Torralba, Phillip Isola, Amy Zhang, Yuandong Tian
Denoised MDPs: Learning World Models Better Than the World Itself
In International Conference on Machine Learning (ICML) 2022
(22% acceptance rate)

35. Tianhao Wu*, Yunchang Yang*, Han Zhong*, Liwei Wang, Simon S. Du, Jiantao Jiao
Nearly Optimal Policy Optimization with Stable at Any Time Guarantee
In International Conference on Machine Learning (ICML) 2022
(22% acceptance rate)

36. Andrew Wagenmaker, Yifang Chen, Max Simchowitz, Simon S. Du, Kevin Jamieson
First-Order Regret in Reinforcement Learning with Linear Function Approximation: A Robust Estimation Approach
In International Conference on Machine Learning (ICML) 2022
(Long talk, 2% acceptance rate)

37. Yunchang Yang*, Tianhao Wu*, Han Zhong*, Evrard Garcelon, Matteo Pirotta, Alessandro Lazaric, Liwei Wang, Simon S. Du
A Reduction-Based Framework for Conservative Exploration
In International Conference on Learning Representations (ICLR) 2022
(32% acceptance rate)

38. Zhili Feng, Shaobo Han, Simon S. Du
Provable Adaptation across Multiway Domains via Representation Learning
In International Conference on Learning Representations (ICLR) 2022
(32% acceptance rate)

39. Yulai Zhao, Yuandong Tian, Jason D. Lee, Simon S. Du
Provably Efficient Policy Gradient Methods for Two-Player Zero-Sum Markov Games
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2022
(29% acceptance rate)

40. Zehao Dou, Zhuoran Yang, Zhaoran Wang, Simon S. Du
Gap-Dependent Bounds for Two-Player Markov Games
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2022
(29% acceptance rate).

41. Xiaoxia Wu, Yuge Xie, Simon S. Du, Rachel Ward
AdaLoss: A Computationally-efficient and Provably Convergent Adaptive Gradient Method
In AAAI Conference on Artificial Intelligence (AAAI) 2022

(15% acceptance rate)

42. Yifang Chen, Simon S. Du, Kevin Jamieson
Corruption Robust Active Learning
In Conference on Neural Information Processing Systems (NeurIPS) 2021
(26% acceptance rate)

43. Tian Ye, Simon S. Du
Global Convergence of Gradient Descent for Asymmetric Low-Rank Matrix Factorization
In Conference on Neural Information Processing Systems (NeurIPS) 2021
(26% acceptance rate)

44. Jean Tarbouriech*, Runlong Zhou*, Simon S. Du, Matteo Pirotta, Michal Valko, Alessandro Lazaric
Stochastic Shortest Path: Minimax, Parameter-Free and Towards Horizon-Free Regret
In Conference on Neural Information Processing Systems (NeurIPS) 2021
(Spotlight presentation, 3% acceptance rate)

45. Tongzheng Ren, Jialian Li, Bo Dai, Simon S. Du, Sujay Sanghavi
Nearly Horizon-Free Offline Reinforcement Learning
In Conference on Neural Information Processing Systems (NeurIPS) 2021
(26% acceptance rate)

46. Zihan Zhang*, Jiaqi Yang*, Xiangyang Ji, Simon S. Du
Improved Variance-Aware Confidence Sets for Linear Bandits and Linear Mixture MDP
In Conference on Neural Information Processing Systems (NeurIPS) 2021
(26% acceptance rate)

47. Zihan Zhang, Xiangyang Ji, Simon S. Du
Is Reinforcement Learning More Difficult Than Bandits? A Near-optimal Algorithm Escaping the Curse of Horizon
In Annual Conference of Learning Theory (COLT) 2021
(35% acceptance rate)

48. Haike Xu, Tengyu Ma, Simon S. Du
Fine-Grained Gap-Dependent Bounds for Tabular MDPs via Adaptive Multi-Step Bootstrap
In Annual Conference of Learning Theory (COLT) 2021
(35% acceptance rate)

49. Simon S. Du*, Wei Hu*, Zhiyuan Li*, Ruoqi Shen*, Zhao Song*, Jiajun Wu*
When is Particle Filtering Efficient for Partially Observed Linear Dynamical Systems?
In Conference on Uncertainty in Artificial Intelligence (UAI) 2021
(26% acceptance rate)

50. Simon S. Du* Sham M Kakade*, Jason D. Lee*, Shachar Lovett*, Gaurav Mahajan*, Wen Sun*, Ruosong Wang*
Bilinear Classes: A Structural Framework for Provable Generalization in RL
In International Conference on Learning Representations (ICML) 2021
(Long talk, 3% acceptance rate)

51. Zihan Zhang, Simon S. Du, Xiangyang Ji
Nearly Minimax Optimal Reward-free Reinforcement Learning
In International Conference on Learning Representations (ICML) 2021
(Long talk, 3% acceptance rate)

52. Yifang Chen, Simon S. Du, Kevin Jamieson
Improved Corruption Robust Algorithms for Episodic Reinforcement Learning
In International Conference on Learning Representations (ICML) 2021
(22% acceptance rate)

53. Tianhao Wu*, Yunchang Yang*, Simon S. Du, Liwei Wang
On Reinforcement Learning with Adversarial Corruption and Its Application to Block MDP
In International Conference on Learning Representations (ICML) 2021
(22% acceptance rate)

54. Kunhe Yang, Lin F. Yang, Simon S. Du
Q-learning with Logarithmic Regret
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2021
(29% acceptance rate)

55. Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du, Ken-ichi Kawarabayashi, Stefanie Jegelka
How Neural Networks Extrapolate: From Feedforward to Graph Neural Networks
In International Conference on Learning Representations (ICLR) 2021
(Oral presentation, 1.7% acceptance rate)

56. Simon S. Du*, Wei Hu*, Sham M. Kakade*, Jason D. Lee*, Qi Lei*
Few-Shot Learning via Learning the Representation, Provably
In International Conference on Learning Representations (ICLR) 2021
(28% acceptance rate)

57. Jiaqi Yang, Wei Hu, Jason D. Lee, Simon S. Du
Impact of Representation Learning in Linear Bandits
In International Conference on Learning Representations (ICLR) 2021
(28% acceptance rate)

58. Yining Wang, Ruosong Wang, Simon S. Du, Akshay Krishnamurthy
Optimism in Reinforcement Learning with Generalized Linear Function Approximation
In International Conference on Learning Representations (ICLR) 2021
(28% acceptance rate)

59. Ruosong Wang*, Simon S. Du*, Lin F. Yang*, Sham M. Kakade
Is Long Horizon Reinforcement Learning More Difficult Than Short Horizon Reinforcement Learning?
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(20% acceptance rate)

60. Fei Feng, Ruosong Wang, Wotao Yin, Simon S. Du, Lin F. Yang
Provably Efficient Exploration for RL with Unsupervised Learning
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(Spotlight presentation, 3% acceptance rate)

61. Simon S. Du*, Jason D. Lee*, Gaurav Mahajan*, Ruosong Wang*
Agnostic Q-learning with Function Approximation in Deterministic Systems: Tight Bounds on Approximation Error and Sample Complexity
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(20% acceptance rate)

62. Ruosong Wang*, Peilin Zhong*, Simon S. Du, Ruslan Salakhutdinov, Lin F. Yang
Planning with General Objective Functions: Going Beyond Total Rewards
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(20% acceptance rate)

63. Ruosong Wang, Simon S. Du, Lin F. Yang, Ruslan Salakhutdinov
On Reward-Free Reinforcement Learning with Linear Function Approximation
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(20% acceptance rate)

64. Yi Zhang*, Orestis Plevrakis*, Simon S. Du, Xingguo Li, Zhao Song, Sanjeev Arora
Over-parameterized Adversarial Training: An Analysis Overcoming the Curse of Dimensionality
In Conference on Neural Information Processing Systems (NeurIPS) 2020
(20% acceptance rate)

65. Sanjeev Arora*, Simon S. Du*, Sham M. Kakade*, Yuping Luo*, Nikunj Saunshi*
Provable Representation Learning for Imitation Learning via Bi-level Optimization
In International Conference on Learning Representations (ICML) 2020
(22% acceptance rate)

66. Yunbo Wang*, Bo Liu*, Jiajun Wu, Yuke Zhu, Simon S. Du, Li Fei-Fei, Joshua B. Tenenbaum
Dual Sequential Monte Carlo: Tunneling Filtering and Planning in Continuous POMDPs
International Joint Conference on Artificial Intelligence (IJCAI) 2020
(12% acceptance rate)

67. Simon S. Du*, Sham M. Kakade*, Ruosong Wang*, Lin F. Yang*
Is a Good Representation Sufficient for Sample Efficient Reinforcement Learning?
In International Conference on Learning Representations (ICLR) 2020
(Spotlight presentation, 6% acceptance rate)

68. Sanjeev Arora*, Simon S. Du*, Ruslan Salakhutdinov*, Ruosong Wang*, Dingli Yu*
Harnessing the Power of Infinitely Wide Deep Nets on Small-data Tasks
In International Conference on Learning Representations (ICLR) 2020
(Spotlight presentation, 6% acceptance rate)

69. Keyulu Xu, Jingling Li, Mozhi Zhang, Simon S. Du, Ken-ichi Kawarabayashi, Stefanie Jegelka
What Can Neural Networks Reason About?
In International Conference on Learning Representations (ICLR) 2020
(Spotlight presentation, 6% acceptance rate)

70. Simon S. Du*, Yuping Luo*, Ruosong Wang*, Hanrui Zhang*
Provably Efficient Q-learning with Function Approximation via Distribution Shift Error Checking Oracle
In Conference on Neural Information Processing Systems (NeurIPS) 2019
(21% acceptance rate)

71. Simon S. Du*, Kangcheng Hou*, Barnabs Pczos*, Ruslan Salakhutdinov*, Ruosong Wang*, Keyulu Xu*
Graph Neural Tangent Kernel: Fusing Graph Neural Networks with Graph Kernels
In Conference on Neural Information Processing Systems (NeurIPS) 2019
(21% acceptance rate)

72. Sanjeev Arora*, Simon S. Du*, Wei Hu*, Zhiyuan Li*, Ruslan Salakhutdinov*, Ruosong Wang*
On Exact Computation with an Infinitely Wide Neural Net
In Conference on Neural Information Processing Systems (NeurIPS) 2019
(Spotlight presentation, 3% acceptance rate)

73. Bin Shi, Simon S. Du, Weijie Su, Michale I. Jordan
Acceleration via Symplectic Discretization of High-Resolution Differential Equations
In Conference on Neural Information Processing Systems (NeurIPS) 2019
(21% acceptance rate)

74. Tianyi Liu*, Minshuo Chen*, Mo Zhou, Simon S. Du, Enlu Zhou and Tuo Zhao
Towards Understanding the Importance of Shortcut Connections in Residual Networks
In Conference on Neural Information Processing Systems (NeurIPS) 2019
(21% acceptance rate)

75. Simon S. Du*, Jason D. Lee*, Haochuan Li*, Liwei Wang*, Xiyu Zhai*
Gradient Descent Finds Global Minima of Deep Neural Networks
In International Conference on Machine Learning (ICML) 2019
(23% acceptance rate)

76. Simon S. Du, Akshay Krishnamurthy, Nan Jiang, Alekh Agarwal, Miroslav Dudík, John Langford
Provably efficient RL with Rich Observations via Latent State Decoding
In International Conference on Learning Representations (ICML) 2019
(23% acceptance rate)

77. Sanjeev Arora*, Simon S. Du*, Wei Hu*, Zhiyuan Li*, Ruosong Wang*
Fine-Grained Analysis of Optimization and Generalization for Overparameterized Two-Layer Neural Networks
In International Conference on Learning Representations (ICML) 2019
(23% acceptance rate)

78. Simon S. Du*, Wei Hu*.

Width Provably Matters in Optimization for Deep Linear Neural Networks
In International Conference on Learning Representations (ICML) 2019
(23% acceptance rate)

79. Simon S. Du*, Xiyu Zhai*, Barnabás Póczos, Aarti Singh.
Gradient Descent Provably Optimizes Over-paramterized Neural Networks
In International Conference on Learning Representations (ICLR) 2019
(33% acceptance rate)

80. Simon S. Du*, Wei Hu*
Linear Convergence of the Primal-Dual Gradient Method for Convex-Concave Saddle Point Problems without Strong Convexity
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2019
(32% acceptance rate)

81. Simon S. Du*, Wei Hu*, Jason D. Lee*
Algorithmic Regularization in Learning Deep Homogeneous Models: Layers are Automatically Balanced
In Conference on Neural Information Processing Systems (NeurIPS) 2018
(24% acceptance rate)

82. Simon S. Du*, Yining Wang*, Xiyu Zhai, Sivaraman Balakrishnan, Ruslan Salakhutdinov, Aarti Singh
How Many Samples are Needed to Learn a Convolutional Neural Network?
In Conference on Neural Information Processing Systems (NeurIPS) 2018
(24% acceptance rate)

83. Simon S. Du, Jason D. Lee, Yuandong Tian, Barnabás Póczos, Aarti Singh
Gradient Descent Learns One-hidden-layer CNN: Don't be Afraid of Spurious Local Minima
In International Conference of Machine Learning (ICML) 2018
(Long talk, 8% acceptance rate)

84. Simon S. Du, Jason D. Lee
On the Power of Over-parametrization in Neural Networks with Quadratic Activation
In International Conference of Machine Learning (ICML) 2018
(25% acceptance rate)

85. Xiao Zhang*, Simon S. Du*, Quanquan Gu
Fast and Sample Efficient Inductive Matrix Completion via Multi-Phase Procrustes Flow
In International Conference of Machine Learning (ICML) 2018
(25% acceptance rate).

86. Yi Wu, Siddharth Srivastava, Nick Hay, Simon S. Du, Stuart Russell
Discrete-Continuous Mixtures in Probabilistic Programming: Generalized Semantics and Inference Algorithms.
In International Conference of Machine Learning (ICML) 2018
(25% acceptance rate)

87. Simon S. Du, Jason D. Lee, Yuandong Tian
When is a Convolutional Filter Easy To Learn?
In International Conference on Learning Representations (ICLR) 2018
(34% acceptance rate)

88. Yining Wang, Simon S. Du, Sivaraman Balakrishnan, Aarti Singh
Stochastic Zeroth-order Optimization in High Dimensions
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2018
(Oral presentation, 5% acceptance rate).

89. Simon S. Du, Chi Jin, Jason D. Lee, Michael I. Jordon, Barnabás Póczos, Aarti Singh
Gradient Descent Can Take Exponential Time to Escape Saddle Points
In Conference on Neural Information Processing Systems (NIPS) 2017
(Spotlight presentation, 4% acceptance rate)

90. Simon S. Du, Yining Wang, Aarti Singh

On the Power of Truncated SVD for General High-rank Matrix Estimation Problems
In Conference on Neural Information Processing Systems (NIPS) 2017
(21% acceptance rate)

91. Simon S. Du, Jayanth Koushik, Aarti Singh, Barnabás Póczos
Hypothesis Transfer Learning via Transformation Functions
In Conference on Neural Information Processing Systems (NIPS) 2017
(21% acceptance rate)

92. Srinivasan Vijayarangan, Paloma Sodhi, Prathamesh Kini, James Bourne, Simon S. Du, Hanqi Sun, Barnabás Póczos, Dimitrios Apostolopoulos, David Wettergreen
High-throughput Robotic Phenotyping of Energy Sorghum Crops
Conference on Field and Service Robotics (FSR) 2017

93. Simon S. Du, Jianshu Chen, Lihong Li, Lin Xiao, Dengyong Zhou
Stochastic Variance Reduction Methods for Policy Evaluation
In International Conference of Machine Learning (ICML) 2017
(25% acceptance rate)

94. Simon S. Du, Sivaraman Balakrishnan, Aarti Singh
Computationally Efficient Robust Estimation of Sparse Functionals
In Annual Conference of Learning Theory (COLT) 2017
(32% acceptance rate, merged with another paper)

95. Shashank Singh, Simon S. Du, Barnabás Póczos
Efficient Nonparametric Smoothness Estimation
In Conference on Neural Information Processing Systems (NIPS) 2016
(23% acceptance rate)

96. Maria-Florina Balcan*, Simon S. Du*, Yining Wang*, Adams Wei Yu*
An Improved Gap-Dependency Analysis of the Noisy Power Method
In Annual Conference on Learning Theory (COLT) 2016
(32% acceptance rate).

97. David G. Anderson*, Simon S. Du*, Michael W. Mahoney*, Christopher Melgaard*, Kunming Wu*, Ming Gu*
Spectral Gap Error Bounds for Improving CUR Matrix Decomposition and Nyström Method
In International Conference on Artificial Intelligence and Statistics (AISTATS) 2015
(27% acceptance rate)

Journal Publications

* denotes equal contribution or alphabetical ordering.

1. Jifan Zhang*, Yifang Chen*, Gregory Canal, Stephen Mussmann, Arnav M. Das, Gantavya Bhatt, Yinglun Zhu, Jeffrey Bilmes, Simon Shaolei Du, Kevin Jamieson, Robert D Nowak
LabelBench: A Comprehensive Framework for Benchmarking Adaptive Label-Efficient Learning
Journal of Data-centric Machine Learning Research (DMLR) 2024
2. Runlong Zhou, Zelin He, Yuandong Tian, Yi Wu, Simon S. Du
Understanding Curriculum Learning in Policy Optimization for Solving Combinatorial Optimization Problems
Transactions on Machine Learning Research (JMLR) 2024
3. Wenqing Zheng*, Hao Yang*, Jiarui Cai, Peihao Wang, Xuan Jiang, Simon S. Du, Yinhai Wang, Zhangyang Wang
Integrating Traffic Science with Representation Learning for City-wide Network Congestion Prediction
Information Fusion 2023
4. Shusheng Xu, Yancheng Liang, Yunfei Li, Simon S. Du, Yi Wu
Beyond Information Gain: An Empirical Benchmark for Low-Switching-Cost Reinforcement Learning

Transactions on Machine Learning Research (TMLR) 2023

5. Bin Shi, Simon S. Du, Michale I. Jordan, Weijie Su
Understanding the Acceleration Phenomenon via High-Resolution Differential Equations
Mathematical Programming (MP) 2021

6. Yining Wang, Yi Wu, Simon S. Du
Near-Linear Time Local Polynomial Nonparametric Estimation with Box Kernels
Informs Journal on Computing 2020

7. Xi Chen*, Simon S. Du*, Xin T. Tong*
On Stationary-Point Hitting Time and Ergodicity of Stochastic Gradient Langevin Dynamics
Journal of Machine Learning Research (JMLR) 2020

Preprints and Technical Reports

* denotes equal contribution or alphabetical ordering.

1. Chuning Zhu, Xinqi Wang, Tyler Han, Simon S. Du, Abhishek Gupta
Transferable Reinforcement Learning via Generalized Occupancy Models

2. Runlong Zhou, Simon S. Du, Beibin Li
Reflect-RL: Two-Player Online RL Fine-Tuning for LMs

3. Avinandan Bose, Simon S. Du, Maryam Fazel
Offline Multi-task Transfer RL with Representational Penalization

4. Qiwen Cui, Maryam Fazel, Simon S. Du
Learning Optimal Tax Design in Nonatomic Congestion Games

5. Yan Dai, Qiwen Cui, Simon S. Du
Refined Sample Complexity for Markov Games with Independent Linear Function Approximation

6. Yiping Wang, Yifang Chen, Wendan Yan, Kevin Jamieson, Simon Shaolei Du
Variance Alignment Score: A Simple But Tough-to-Beat Data Selection Method for Multimodal Contrastive Learning

7. Gantavya Bhatt*, Yifang Chen*, Arnav M. Das*, Jifan Zhang*, Sang T. Truong, Stephen Mussmann, Yinglun Zhu, Jeffrey Bilmes, Simon S. Du, Kevin Jamieson, Jordan T. Ash, Robert D. Nowak
An Experimental Design Framework for Label-Efficient Supervised Finetuning of Large Language Models

8. Zihan Zhang, Wenhao Zhan, Yuxin Chen, Simon S. Du, Jason D. Lee
Optimal Multi-Distribution Learning

9. Meixin Zhu, Simon S. Du, Xuesong Wang, Hao Yang, Ziyuan Pu, Yinhai Wang
TransFollower: Long-Sequence Car-Following Trajectory Prediction through Transformer

10. Minbo Gao*, Tianle Xie*, Simon S. Du, Lin F. Yang
A Provably Efficient Algorithm for Linear Markov Decision Process with Low Switching Cost

11. Xiang Wang, Xinlei Chen, Simon S. Du, Yuandong Tian
Towards Demystifying Representation Learning with Non-contrastive Self-supervision

12. Zhiyuan Li*, Ruosong Wang*, Dingli Yu*, Simon S. Du, Wei Hu, Ruslan Salakhutdinov, Sanjeev Arora
Enhanced Convolutional Neural Tangent Kernels

13. Simon S. Du*, Ruosong Wang*, Mengdi Wang*, Lin F. Yang*
Continuous Control with Contexts, Provably

14. Simon S. Du*, Surbhi Goel*
Improved Learning of One-hidden-layer Convolutional Neural Networks with Overlaps

15. Simon S. Du, Yining Wang, Sivaraman Balakrishnan, Pradeep Ravikumar, Aarti Singh
Robust Nonparametric Regression under Huber's ϵ-contamination Model

Teaching Experiences

Instructor:

- Machine Learning (UW CSE 446), Spring 2024
- Deep Learning (UW CSE 543), Fall 2023
- Machine Learning (UW CSE 446/546, co-teach with Kevin Jamieson), Spring 2023
- Deep Learning (UW CSE 543), Winter 2023
- Deep Learning (UW CSE 543/599I), Spring 2022
- Machine Learning (UW CSE 446/546, co-teach with Jamie Morgenstern), Autumn 2021
- Machine Learning (UW CSE 446/546, co-teach with Sewoong Oh), Spring 2021
- Theoretical Deep Learning (UW CSE 599D1), Winter 2021

Teaching Assistant:

- Introduction to Machine Learning (CMU 10-601, Instructor: Matt Gormley) Spring 2017
- Introduction to Machine Learning (CMU 10-601, Instructors: Maria-Florina Balcan and Matt Gormley) Fall 2016

Guest Lecturer:

- Statistical Methods in Computer Science (University of Washington), Fall 2023
- Introduction to Computer Science (Tsinghua University), Fall 2023
- Theoretical Deep Learning (Princeton University), Fall 2019
- Introduction to Artificial Intelligence (Tsinghua University), Fall 2019

Mentoring Experiences

Ph.D. Student Advising:

- Yifang Chen — 2020 - present
  UW CSE Ph.D. (co-advised with Kevin Jamieson)
- Qiwen Cui — 2021 - present
  UW CSE Ph.D.
- Vector Runlong Zhou — 2022 - present
  UW CSE Ph.D.
- Yancheng Liang — 2023 - present
  UW CSE Ph.D. (co-advsied with Natasha Jaques)
- Siting Li — 2023 - present
  UW CSE Ph.D.
- Divyansh Pareek — 2023 - present
  UW CSE Ph.D. (co-advised with Sewoong Oh)
- Xinqi Wang — 2023 - present
  UW CSE Ph.D.
- Weihang Xu — 2023 - present
  UW CSE Ph.D. (co-advised with Maryam Fazel)
- Yiping Wang — 2023 - present
  UW CSE Ph.D.

Postdoc Advising:

- Zihan Zhang — 2023 - present
  co-advised with Yuxin Chen and Jason D. Lee

Former Ph.D. Students and Postdocs:

- Ruosong Wang 2022 - 2023
  UW Postdoc → Assistant Professor at PKU CFCS

Former Undergraduate Student Researchers:

- Yan Dai 2022 - 2024
  THU IIIS Undergraduate→ Visit at UW → MIT ORC Ph.D.
- Animesh Jha 2022 - 2024
  IIT Kharagpur CSE Undergraduate → Virtual visit at UW → Stanford CS Master
- Eric Jiang 2022 - 2024
  THU IIIS Undergraduate→ Visit at UW → Berkeley EECS Ph.D.
- Liming Liu 2023 - 2024
  PKU Math Undergraduate→ Visit at UW → Georgia Tech ISYE Ph.D.
- Nuoya Xiong 2023 - 2024
  THU IIIS Undergraduate→ Visit at UW → CMU MLD Ph.D.
- Jiarui Yao 2023 - 2024
  THU IIIS Undergraduate→ Visit at UW → UIUC CS Ph.D.
- Zhaoyi Zhou 2023 - 2024
  THU IIIS Undergraduate→ Visit at UW → CMU ECE Ph.D.
- Stanley Wei 2022 - 2023
  UT Austin CS Undergraduate → Virtual Visit at UW → Princeton ECE Ph.D.
- Haotian Ye 2022 - 2023
  PKU Math Undergraduate→ Virtual Visit at UW → Stanford CS Ph.D.
- Jikai Jin 2022 - 2023
  PKU Math Undergraduate→ Virtual Visit at UW → Stanford ICME Ph.D.
- Xinqi Wang 2022 - 2023
  THU IIIS Undergraduate→ Visit at UW → UW CSE Ph.D.
- Yiping Wang 2022 - 2023
  ZJU CS Undergraduate→ Virtual Visit at UW → UW CSE Ph.D.
- Siting Li 2022 - 2023
  THU IIIS Undergraduate→ Virtual Visit at UW → UW CSE Ph.D.
- Weihang Xu 2022 - 2023
  THU IIIS Undergraduate→ Visit at UW → UW CSE Ph.D.
- Qing Yang 2022 - 2023
  SJTU CS Undergraduate → Visit at UW → USC CS Master
- Vector Runlong Zhou 2021 - 2022
  THU IIIS Undergraduate→ Virtual Visit at UW → UW CSE Ph.D.
- Haike Xu 2021
  THU IIIS Undergraduate→ Virtual Visit at UW → MIT EECS Ph.D.
- Yulai Zhao 2021 - 2022
  THU EE Undergraduate→ Virtual Visit at UW → Princeton ECE Ph.D.
- Haoyuan Cai 2021 - 2022
  THU IIIS Undergraduate→ Virtual Visit at UW → Princeton ECE Ph.D.
- Hetvi Jethwani 2020 - 2021
  IIT Delhi Math Undergraduate → Virtual Visit at UW → StatML Ph.D. based at Imperial and Oxford
- Alex DeWeese 2021
  UC Berkeley Undergraduate→ Virtual Visit at UW → CMU ECE Ph.D.
- Tian Ye 2020 - 2021
  THU IIIS Undergraduate→ Virtual Visit at UW → CMU MLD Ph.D.
- Kunhe Yang 2020 - 2021
  THU IIIS Undergraduate → Virtual Visit at UW → UC Berkeley EECS Ph.D.
- Jiaqi Yang 2020 - 2021
  THU IIIS Undergraduate → Virtual Visit at UW → UC Berkeley EECS Ph.D.
- Tianhao Wu 2020 - 2021
  PKU Math Undergraduate → Virtual Visit at UW → UC Berkeley EECS Ph.D.
- Rui Lu 2020 - 2021
  THU IIIS Undergraduate → Virtual Visit at UW → THU Automation Ph.D.

- Kangcheng Hou 2019
  ZJU CS Undergraduate→ Visit at CMU → UCLA CS Ph.D.
- Haochuan Li 2018
  PKU Physics Undergraduate→ Virtual Visit at CMU → MIT EECS Ph.D.
- Hanqi Sun 2016 - 2017
  CMU CS Undergraduate → NVIDIA
- Langxuan Su 2016 - 2017
  CMU Math Undergraduate → Duke Math Ph.D.

High School Student Mentoring:

- Reed Evertt (Tesla STEM High School) 2020 - 2021
  2nd place in the Robotics and Intelligent Machines at WSSEF

Thesis Committees:

- Chuning Zhu
  UW CSE Ph.D. Qualifying exam Spring 2024
- Ravil Mussabayev
  UW Math Ph.D. Final exam Spring 2024
  UW Math Ph.D. General exam Winter 2022
- Andrew Wagenmaker
  UW CSE Ph.D. Final exam Spring 2024
  UW CSE Ph.D. General exam Spring 2022
- Ruoqi Shen
  UW CSE Ph.D. Final exam Spring 2024
  UW CSE Ph.D. General exam Fall 2021
- Chengyuan Ma
  UW Math Ph.D. Final exam Spring 2024
  UW Math Ph.D. General exam Spring 2022
- Jiawei Yao
  UW Tacoma CSS Ph.D. General exam Spring 2024
- Daogao Liu
  UW CSE Ph.D. General exam Winter 2024
- Haoran Zhu
  UW Physics Ph.D. General exam Summer 2023
- Frank Yang
  UW CEE Ph.D. Final exam Summer 2023
  UW CEE Ph.D. General exam Winter 2023
- Qiwen Cui
  UW CSE Ph.D. Qualifying exam Spring 2023
- Vector Runlong Zhou
  UW CSE Ph.D. Qualifying exam Spring 2023
- Meng-Ju Tsai
  UW CEE Ph.D. Final exam Spring 2023
  UW CEE Ph.D. General exam Winter 2023
- Yuming Hubert Liu
  UW Economics Ph.D. Final exam Fall 2022
  UW Economics Ph.D. General exam Spring 2022
- Haotian Jiang
  UW CSE Ph.D. Final exam Fall 2022
  UW CSE Ph.D. General exam Fall 2020
- Arnav Das
  UW ECE Ph.D. Qualifying exam Fall 2022
- Mayuree Binjolkar
  UW CEE Ph.D. General exam Fall 2022
- Farzam Ebrahimnejad
  UW CSE Ph.D. General exam Spring 2022

- Zhihan Xiong
  UW CSE Ph.D. Qualification exam — Spring 2022
- Jiarui Cai
  UW ECE Ph.D.Final exam — Spring 2022
  UW ECE Ph.D. General exam — Spring 2021
- Yifang Chen
  UW CSE Ph.D. Qualification exam — Spring 2022
- Farzam Ebrahimnejad
  UW CSE Ph.D. General exam — Spring 2022
- Sabiha Rustam
  UW Chemical Engineering Ph.D. General exam — Fall 2021
- John E. Ash
  UW CEE Ph.D. Final exam — Winter 2021
- Biraj Pandey
  UW Math Ph.D. General exam — Fall 2020

Grants & Gifts

Sloan Research Fellowship *2024*

- PI: Simon S. Du (UW)
- Total funded amount: $75,000

Intel Rising Star Faculty Award *2023*

- PI: Simon S. Du (UW)
- Total funded amount: $50,000

UW + Amazon Science Faculty Research Award *2023*

- Title: Theoretically principled representation learning for multi-task reinforcement learning
- PI: Simon S. Du (UW); Co-PI: Maryam Fazel (UW)
- Total funded amount: $100,000

National Science Foundation (NSF-IIS AI Institute #2229881) *2023*

- Title: AI Institute for Societal Decision Making
- PI: Aarti Singh (CMU); Co-PIs: Jeff Schneider (CMU), Cleotilde Gonzalez (CMU), Terri Adams-Fuller (Howard), Sham Kakade (Harvard)
- Simon S. Du: Senior Personnel
- Total funded amount: $19,879,659; Personal share: $750,000 (2023 - 2028)

National Science Foundation (NSF-CIF Medium #2212261) *2022*

- Title: Toward a Mathematical Foundation of Deep Reinforcement Learning
- PI: Simon S. Du (UW); Co-PIs: Maryam Fazel (UW), Jason D. Lee (Princeton), Tengyu Ma (Stanford)
- Total funded amount: $1,200,000; Personal share: $300,000 (2022 - 2026)

National Science Foundation (NSF CAREER # 2143493) *2022*

- Title: Toward a Foundation of Over-Parameterization
- PI: Simon S. Du (UW)
- Total funded amount: $570,000 (2022 - 2027)

eScience Institute UW Azure Cloud Computing Credits *2021*

- PI: Simon S. Du (UW)
- Total funded amount: $20,000

Tencent AI Lab Rhino-Bird Gift Fund *2021*

- PI: Simon S. Du (UW)
- Total funded amount: $50,000

National Science Foundation (NSF SCALE MoDL # 2134106) *2021*

- Title: Adaptivity of Deep Neural Networks
- PI: Yu-Xiang Wang (UCSB); Co-PIs: Simon S. Du (UW), Yingbin Liang (OSU)
- Total funded amount: $900,000; Personal share: $300,000 (2021 - 2024)

National Science Foundation (NSF-IIS Small #2110170) *2021*

- Title: Theoretical Foundations of Reinforcement Learning: From Tabula Rasa to Function Approximation
- PI: Simon S. Du (UW)
- Total funded amount: $500,000 (2021 - 2024)

NEC Lab America Gift Fund *2021*

- PI: Simon S. Du (UW)
- Total funded amount: $5,000

Invited Talks

The Paradigm Shifts in AI

- Materials Research Society Summit, April 2024, Seattle, USA

Pre-Training Data Selection for Representation Learning

- University of Illinois at Urbana-Champaign Department of CS, Feb 2024, Champaign, USA
- University of Washington Biostatistics Department, Feb 2024, Seattle, USA

How Over-Parameterization Slows Down Gradient Descent

- Kyoto University, March 2024, Kyoto, Japan
- Carnegie Mellon University, Feb 2024, Pittsburgh, USA
- University of Illinois at Urbana-Champaign Department of ECE, Feb 2024, Champaign, USA
- New York University, Nov 2023, New York City, USA
- Google Efficient ML Workshop, Nov 2023, New York City, USA
- Big Data and Machine Learning in Econometrics, Finance, and Statistics, Oct 2023, Chicago, USA

Passive and Active Multi-Task Representation Learning

- (Virtual) Intel, Nov 2023
- (Virtual) Transportation Research Board, Aug 2023
- SIAM Conference on Optimization, Jun 2023, Seattle, USA
- Fudan University, May 2023, Shanghai, China
- (Virtual) RWTH Aachen University, Apr 2023, Aachen, Germany
- (Virtual) Microsoft Research Asia Theory Talk Series, Jan 2023, Beijing, China
- Third Workshop on Seeking Low Dimensionality in Deep Neural Networks, Jan 2023, Abu Dhabi, United Arab Emirates
- (Virtual) Google Research Learning Theory Seminar, Nov 2022, New York City, USA
- Samsung Advanced Insitute of Technology, Nov 2022, Suwon, Korea
- (Virtual) Institute for Foundations of Machine Learning, Oct 2022, Austin, USA
- Tencent AI Lab, Sep 2022, Bellevue, USA
- (Virtual) University of Maryland, Sep 2022, Baltimore, USA
- Representation Learning Theory Workshop at TTIC, Aug 2022, Chicago, USA

Toward a Theoretical Foundation of Deep Learning

- Samsung Artificial Intelligence Forum, Nov 2022, Seoul, Korea

When are Offline Multi-Agent Games Solvable?

- Carnegie Mellon University, Feb 2024, Pittsburgh, USA
- University of Washington CS Theory Talk, Jan 2023, Seattle, USA
- University of Washington Industrial and System Engineering, Nov 2022, Seminar, Seattle, USA
- (Virtual) The 1st International Workshop on Safe Reinforcement Learning Theory and its Applications, Sep 2022, Cran Field, UK

- New Directions in Theoretical Machine Learning, Sep 2022, Krün, Germany
- NSF AI in Transportation Workshop, June 2022, Seattle, USA
- New Advances in Statistics and Data Science, May 2022, Honolulu, USA
- Simons Institute, May 2022, Berkeley, USA
- (Virtual) Institute for Foundations of Data Science, April 2022

Horizon-Free Reinforcement Learning

- (Virtual) Berkeley Laboratory for Information and System Sciences (BLISS) Seminar, Sep 2021
- (Virtual) Qingyuan Salon for Young AI Scientist Keynote Talks, Apr 2021
- (Virtual) Reinforcement Learning Theory Virtual Seminar, Nov 2020

Nearly Minimax Optimal Reward-free Reinforcement Learning

- (Virtual) Simons Institute, Dec 2020, Berkeley, USA

Provable Representation Learning

- (Virtual) International Joint Conference on Theoretical Computer Science, Aug 2021, Beijing, China
- (Virtual) Instituto Superior Técnico, July 2021, Lisbon, Portugal
- (Virtual) Amazon, Jun 2021
- (Virtual) Conference on Deployable AI, Jun 2021, Chennai, India
- (Virtual) Facebook AI Research, Apr 2021
- (Virtual) Beijing Academy of Artificial Intelligence, Feb 2021, Beijing, China
- (Virtual) University of Washington, Nov 2020, Seattle, USA
- (Virtual) Tencent AI Lab, Aug 2020, Bellevue, USA

On Reinforcement Learning with Large State Space and Long Horizon

- (Virtual) University of Washington Statistics Seminar, Nov 2020, Seattle, USA
- (Virtual) Microsoft Research, Jul 2020, Redmond, USA

Ultra-wide Neural Networks and Kernels

- Radix Trading, Mar 2022, Chicago, USA
- (Virtual) Shanghai University of Finance and Economics, Aug 2021, Shanghai, China
- (Virtual) AAAI New Faculty Highlights, Feb 2021
- (Virtual) Two Sigma, New York City, Jan 2021
- (Virtual) TechBeat, Jan 2021
- (Virtual) INFORMS Annual Meeting, Session on Bridging Deep Learning with Stochastic Analysis and Mean-Field Theory, Nov 2020
- (Virtual) Bytedance, Aug 2020, Beijing, China
- (Virtual) Max Planck Institute, Aug 2020, Leipzig, Germany
- (Virtual) SIAM Conference on Mathematics of Data Science, Jun 2020, Cincinnati, USA
- Georgia Institute of Technology ISyE Statistics Seminar Series, Nov 2019, Atlanta, USA
- IIIS-Haihua Distinguished Seminar Series in Artificial Intelligence, Oct 2019, Beijing, China
- Institute for Advanced Study Workshop on Theory of Deep Learning, Oct 2019, Princeton, USA
- Institute for Advanced Study Talk Series, Sep 2019, Princeton, USA

Foundations of Learning Systems with (Deep) Function Approximators

- (Virtual) Facebook Artificial Intelligence Research, Oct 2020
- (Virtual) University of Texas at Austin Computer Science Department, Apr 2020, Austin, USA
- (Virtual) University of California at San Diego Data Science Institute, Mar 2020, San Diego, USA
- (Virtual) Georgia Institute of Technology School of Computer Science, Mar 2020, Atlanta, USA
- (Virtual) University of Southern California Department of Computer Science, Mar 2020, Los Angeles, USA
- University of California at Los Angeles Computer Science Department, Mar 2020, Los Angeles, USA
- (Virtual) University of Washington School of Computer Science & Engineering, Mar 2020, Seattle, USA
- Cornell University Department of Computer Science, Feb 2020, Ithaca, USA
- University of Illinois at Urbana-Champaign Department of ECE, Feb 2020, Champaign, USA
- University of Chicago Computer Science Department, Feb 2020, Chicago, USA
- Yale University Department of Statistics and Data Science, Feb 2020, New Haven, USA
- Duke University Computer Science Department, Jan 2020, Durham, USA
- University of Pennsylvania Wharton Business School, Jan 2020, Philadelphia, USA

- University of Michigan Computer Science & Engineering Division, Jan 2020, Ann Arbor, USA

Provably Efficient Reinforcement Learning with Function Approximation

- (Virtual) INFORMS Annual Meeting, Session on Information, Learning and Inference, Nov 2020
- Simons Institute, Dec 2019, Berkeley, USA
- Princeton University Alg-ML Seminar, Jul 2019, Princeton, USA
- Peking University Machine Learning Theory Workshop, Jun 2019, Beijing, China

Understanding Optimization and Generalization in Deep Learning: A Trajectory-based Analysis

- CMU AI Lunch, Feb 2019, Pittsburgh, USA
- Harbin Institute of Technology, Jan 2019, Harbin, China
- Peking University, Dec 2018, Beijing, China
- Institute for Interdisciplinary Information Sciences, Tsinghua University, Dec 2018, Beijing, China

On the Power of Randomly Initialized Gradient Descent for Learning Convolutional Neural Networks

- Peking University, Jan 2018, Beijing, China
- Microsoft Research Asia, Jan 2018, Beijing, China
- Institute of Computing Technology, Chinese Academy of Sciences, Jan 2018, Beijing, China
- Tencent AI Lab, Jan 2018, Shenzhen, China
- Kuaishou, Jan 2018, Beijing, China
- FLAIR (Future Leaders of AI Retreat), Dec 2017, Shanghai, China
- AT&T Labs Graduate Student Symposium, Dec 2017, New York City, USA
- CMU Deep Learning Seminar, Sep 2017, Pittsburgh, USA
- Facebook AI Research, Aug 2017, Menlo Park, USA

Stochastic Variance Reduction Methods for Policy Evaluation

- CMU SELECT Seminar, Feb 2017, Pittsburgh, USA
- Microsoft Research Machine Learning Lunch, Aug 2016, Redmond, USA

Efficient Robust Sparse Estimation in High Dimensions

- CMU Statistical Machine Learning Seminar, Oct 2017, Pittsburgh, USA

Professional Activities

Area Chair or Senior Program Committee:

- Conference on Artificial Intelligence (AAAI) 2020, 2021
- International Conference on Algorithmic Learning Theory (ALT) 2024
- Annual Conference on Learning Theory (COLT) 2021, 2022, 2023, 2024
- International Conference on Learning Representations (ICLR) 2021, 2022, 2023, 2024
- International Conference on Machine Learning (ICML) 2020, 2021, 2022, 2023, 2024
- International Joint Conferences on Artificial Intelligence (IJCAI) 2021
- Conference on Neural Information Processing Systems (NeurIPS) 2021, 2022, 2023, 2024

Paper Reviewer:

- Journal:
  - Artificial Intelligence Review (AIRE)
  - Annals of Applied Probability (AAP)
  - Annals of Statistics (AOS)
  - Electronic Journal of Statistics (EJS)
  - Communications in Computational Physics (CiCP)
  - Cybernetics and Systems
  - Journal of the American Statistical Association (JASA)
  - Journal of Machine Learning Research (JMLR, *editorial board*)
  - Journal of Scientific Computing (JOMP)
  - Mathematics of Operations Research (MathOR)

  - Machine Learning Journal (MLJ)
  - Mathematical Programming (MP)
  - Neural Computation
  - Proceedings of the IEEE
  - Operation Research (OR)
  - Proceedings of the National Academy of Sciences of the United States of America (PNAS)
  - IEEE Transactions on Automatic Control (TAC)
  - IEEE Transactions on Artificial Intelligence (TAI)
  - IEEE Transactions on Image Processing (TIP)
  - IEEE Transactions on Information Theory (TIT)
  - Transactions of Machine Learning Research (TMLR)
  - IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
  - SIAM Journal on Mathematics of Data Science (SIMODS)
  - SIAM Journal on Numerical Analysis (SINUM)
  - Transactions on Signal Processing (TSP)
- Conference:
  - Conference on Artificial Intelligence (AAAI)
  - International Conference on Artificial Intelligence and Statistics (AISTATS)
  - International Conference on Algorithmic Learning Theory (ALT)
  - IEEE Conference on Decision and Control (CDC)
  - Annual Conference on Learning Theory (COLT)
  - European Workshop on Reinforcement Learning (EWRL)
  - International Conference on Learning Representations (ICLR)
  - International Conference on Machine Learning (ICML)
  - IEEE International Symposium on Information Theory (ISIT)
  - International Joint Conferences on Artificial Intelligence (IJCAI)
  - ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)
  - Conference on Systems and Machine Learning (MLSys)
  - Conference on Neural Information Processing Systems (NeurIPS)
  - Optimization for Machine Learning (OPT)
  - Symposium on Discrete Algorithms (SODA)
  - Symposium on the Theory of Computing (STOC)
  - Conference on Uncertainty in Artificial Intelligence (UAI)

Other Academic Services:

- Israel Science Foundation proposal reviewer 2024
- NSF IIS-RI Panel 2023
- NSF IIS-III Panel 2023
- NSF CCF-CIF Panel 2023
- NSF IIS-RI Panel 2022
- Poland National Science Centre proposal reviewer 2023
- Swiss NSF proposal reviewer 2023
- NSF DMS Ad Hoc proposal reviewer 2022
- NeurIPS 2021 Workshop reviewer
- ICML 2021 Workshop reviewer
- ICLR 2021 Session Chair
- COLT 2021 Session Chair

Organizer:

- The Frontiers of AI-Empowered Methods and Solutions to Urban Transportation Challenges 2022

- NeurIPS 2021 Workshop on Ecological Theory of Reinforcement Learning
- ICML 2021 Reinforcement Learning Theory Workshop
- One World Seminar Series on the Mathematics of Machine Learning (founding member) 2020 - 2021
- CMU ML Blog Editor (founding member) 2018-2019

University Services:

- UW CSE Faculty Recruiting Committee 2024
- UW Royalty Research Fund Reviewer 2022
- UW CSE Ph.D. Admission Area Chair 2022, 2023
- UW CSE Prospective Student Visit Day Machine Learning Session Organizer 2021
- CMU Machine Learning Department Ph.D. Admission Committee 2019
- CMU Machine Learning Department Master Admission Committee 2017

Outreach Activities:

- Mentor for Learning Theory Mentorship Workshop 2022
- Talk on machine learning and UW CSE to UW STARS students
- Advise UW STARS student about research and career
- AISTATS 2020 Mentor
- Bellevue School District STEM Career Fair Panelist