

# SIMON SHAOLEI DU

Paul G. Allen School of Computer Science & Engineering

ssdu@cs.washington.edu

<http://simonshaoleidu.com>

---

## RESEARCH INTERESTS

Large Language Models, Agent, Data Selection, Reinforcement Learning, Learning Theory, 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

**XAI** 2025 - present

- *Member of Technical Staff*

**University of Washington, School of Computer Science & Engineering** 2025 - Present

- *Associate Professor*

**University of Washington, School of Computer Science & Engineering** 2020 - 2025

- *Assistant Professor*

**Facebook AI Research** 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

---

<b>MIT Technology Review Innovators Under 35 Asia Pacific Award</b> Awarded to 35 innovators under 35 in the Asia Pacific region.	<i>2025</i>
<b>IEEE AI's 10 to Watch</b> Awarded to 10 rising stars in the field of Artificial Intelligence worldwide.	<i>2024</i>
<b>Schmidt Science AI2050 Early Career Fellowship</b> Awarded to 20 early-career scholars to pursue bold and ambitious work on hard problems in AI.	<i>2024</i>
<b>Alfred P. Sloan Research Fellowship</b> Awarded to 22 early-career scholars in computer science in the US and Canada.	<i>2024</i>
<b>Intel Rising Star Faculty Award</b> Awarded to fifteen early-career professors worldwide.	<i>2023</i>
<b>Samsung AI Researcher of the Year</b> Awarded to five early-career researchers in AI worldwide.	<i>2022</i>
<b>National Science Foundation CAREER Award</b> Supports early-career faculty who have the potential to serve as academic role models.	<i>2022</i>
<b>AAAI New Faculty Highlights</b> Awarded to eighteen new faculties working in AI.	<i>2021</i>
<b>Tencent AI Rhino-Bird Award</b>	<i>2021</i>
<b>CMU School of Computer Science Distinguished Dissertation Award (Runner Up)</b> Top five dissertations from CMU School of Computer Science.	<i>2019</i>
<b>Nomination for ACM Doctoral Dissertation Award</b> Two dissertations were selected from CMU.	<i>2019</i>
<b>Nomination for Joint AAAI/ACM SIGAI Doctoral Dissertation Award</b> One dissertation was selected from CMU.	<i>2019</i>
<b>Late-breaking Paper Award for NeurIPS Deep Reinforcement Learning Workshop</b>	<i>2019</i>
<b>Nvidia Pioneer Award</b>	<i>2018</i>
<b>Best Paper Award for ICML Workshop on Nonconvex Optimization</b>	<i>2018</i>
<b>UC Berkeley EECS Honor Program</b>	<i>2014</i>
<b>Phi Beta Kappa</b>	<i>2014</i>

CONFERENCE PUBLICATIONS

---

\* denotes equal contribution or alphabetical ordering.

1. Paresh R. Chaudhary, Yancheng Liang, Daphne Chen, Simon Shaolei Du, Natasha Jaques  
**Improving Human-AI Coordination through Online Adversarial Training and Generative Models**  
*In International Conference on Learning Representations (ICLR) 2026*  
(28% acceptance rate)
2. Yiran Zhang, Weihang Xu, Mo Zhou, Maryam Fazel, Simon Shaolei Du  
**Convergence Dynamics of Over-Parameterized Score Matching for a Single Gaussian**  
*In International Conference on Learning Representations (ICLR) 2026*  
(28% acceptance rate)

3. Shuyue Stella Li\*, Avinandan Bose\*, Faeze Brahman, Simon Shaolei Du, Pang Wei Koh, Maryam Fazel, Yulia Tsvetkov  
**Personalized Reasoning: Just-in-time Personalization and Why LLMs Fail at It**  
*In International Conference on Learning Representations (ICLR) 2026*  
(28% acceptance rate)
4. Divyansh Pareek, Sewoong Oh, Simon S. Du  
**Understanding the Gain from Data Filtering in Multimodal Contrastive Learning**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
5. Shulun Chen, Runlong Zhou, Zihan Zhang, Maryam Fazel, Simon S. Du  
**Gap-Dependent Variance-Aware Regret Bounds for Tabular MDPs**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
6. Siting Li, Xiang Gao, Simon S. Du  
**Highlighting What Matters: Promptable Embeddings for Attribute-Focused Image Retrieval**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
7. Yiping Wang, Qing Yang, Zhiyuan Zeng, Liliang Ren, Lucas Liu, Baolin Peng, Hao Cheng, Xuehai He, Kuan Wang, Jianfeng Gao, Weizhu Chen, Shuohang Wang\*, Simon S. Du\*, Yelong Shen\*  
**Reinforcement Learning for Reasoning in Large Language Models with One Training Example**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
8. Zihan Zhang, Yuxin Chen, Jason D. Lee, Simon Shaolei Du, Lin Yang, Ruosong Wang  
**Deployment Efficient Reward-Free Exploration with Linear Function Approximation**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
9. Liming Liu, Zixuan Zhang, Simon Du, Tuo Zhao  
**A Minimalist Example of Edge-of-Stability and Progressive Sharpening**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2025*  
(25% acceptance rate)
10. Avinandan Bose, Zhihan Xiong, Yuejie Chi, Simon S. Du, Lin Xiao, Maryam Fazel  
**LoRe: Personalizing LLMs via Low-Rank Reward Modeling**  
*In Conference on Language Modeling (COLM) 2025*  
(32% acceptance rate)
11. Runlong Zhou, Maryam Fazel, Simon S. Du  
**Extragradient Preference Optimization (EGPO): Beyond Last-Iterate Convergence for Nash Learning from Human Feedback**  
*In Conference on Language Modeling (COLM) 2025*  
(32% acceptance rate)
12. Xiyu Zhai, Runlong Zhou, Liao Zhang, Simon S. Du  
**Transformers are Efficient Compilers, Provably**  
*In Conference on Language Modeling (COLM) 2025*  
(32% acceptance rate)
13. Siting Li, Pang Wei Koh, Simon Shaolei Du  
**Exploring How Generative MLLMs Perceive More Than CLIP with the Same Vision Encoder**  
*In Annual Meeting of the Association for Computational Linguistics (ACL) 2025*  
(Main conference)
14. Zihan Zhang, Jason D. Lee, Simon S. Du, Yuxin Chen  
**Anytime Acceleration of Gradient Descent**  
*In Annual Conference of Learning Theory (COLT) 2025*  
(31% acceptance rate)

15. Zihan Zhang, Yuxin Chen, Jason D. Lee, Simon S. Du, Ruosong Wang  
**Minimax Optimal Regret Bound for Reinforcement Learning with Trajectory Feedback**  
*In International Conference on Machine Learning (ICML) 2025*  
 (27% acceptance rate)
16. Kunal Jha, Wilka Carvalho, Yancheng Liang, Simon S. Du, Max Kleiman-Weiner\*, Natasha Jaques\*  
**Cross-environment Cooperation Enables Zero-shot Multi-agent Coordination**  
*In International Conference on Machine Learning (ICML) 2025*  
 (spotlight, 3% acceptance rate)
17. Yiping Wang, Xuehai He, Kuan Wang, Luyao Ma, Jianwei Yang, Shuohang Wang, Simon S. Du, Yelong Shen  
**Is Your World Simulator a Good Story Presenter? A Consecutive Events-Based Benchmark for Future Long Video Generation**  
*In Conference on Computer Vision and Pattern Recognition (CVPR) 2025*  
 (22 % acceptance rate)
18. Ruizhe Shi\*, Runlong Zhou\*, Simon S. Du  
**The Crucial Role of Samplers in Online Direct Preference Optimization**  
*In International Conference on Learning Representations (ICLR) 2025*  
 (32% acceptance rate)
19. Avinandan Bose, Simon S. Du, Maryam Fazel  
**Offline Multi-task Transfer RL with Representational Penalization**  
*In International Conference on Artificial Intelligence and Statistics (AISTATS) 2025*  
 (31% acceptance rate)
20. Yiping Wang\*, Yifang Chen\*, Wendan Yan, Alex Fang, Wenjing Zhou, Kevin Jamieson, Simon Shaolei Du  
**CLIPLoss and Norm-Based Data Selection Methods for Multimodal Contrastive Learning**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2024*  
 (Spotlight, 2% acceptance rate)
21. Yancheng Liang, Daphne Chen, Abhishek Gupta, Simon S. Du\*, Natasha Jaques\*  
**Learning to Cooperate with Humans using Generative Agents**  
*Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)
22. Divyansh Pareek, Simon S. Du, Sewoong Oh  
**Understanding the Gains from Repeated Self-Distillation**  
*In Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)
23. Ruizhe Shi, Yifang Chen, Yushi Hu, Alisa Liu, Hannaneh Hajishirzi, Noah A. Smith, Simon S. Du  
**Decoding-Time Language Model Alignment with Multiple Objectives**  
*In Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)
24. Weihang Xu, Maryam Fazel, Simon S. Du  
**Toward Global Convergence of Gradient EM for Over-Parameterized Gaussian Mixture Models**  
*In Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)
25. Chuning Zhu, Xinqi Wang, Tyler Han, Simon S. Du, Abhishek Gupta  
**Transferable Reinforcement Learning via Generalized Occupancy Models**  
*In Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)
26. Qiwen Cui, Maryam Fazel, Simon S. Du  
**Learning Optimal Tax Design in Nonatomic Congestion Games**  
*In Information Processing Systems (NeurIPS) 2024*  
 (25% acceptance rate)

27. Zihan Zhang, Yuxin Chen, Jason D. Lee, Simon S. Du  
**Settling the Sample Complexity of Online Reinforcement Learning**  
*In Annual Conference of Learning Theory (COLT) 2024*  
(35% acceptance rate)
28. Yan Dai, Qiwen Cui, Simon S. Du  
**Refined Sample Complexity for Markov Games with Independent Linear Function Approximation**  
*In Annual Conference of Learning Theory (COLT) 2024*  
(35% acceptance rate)
29. Zihan Zhang, Wenhao Zhan, Yuxin Chen, Simon S. Du, Jason D. Lee  
**Optimal Multi-Distribution Learning**  
*In Annual Conference of Learning Theory (COLT) 2024*  
(35% acceptance rate)
30. Runlong Zhou, Simon S. Du, Beibin Li  
**Reflect-RL: Two-Player Online RL Fine-Tuning for LMs**  
In Annual Meeting of the Association for Computational Linguistics (ACL) 2024  
(Main conference)
31. 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**  
In Annual Meeting of the Association for Computational Linguistics (ACL) 2024  
(Findings)
32. Chenhao Lu, Ruizhe Shi\*, Yuyao Liu\*, Kaizhe Hu, Simon S. Du, Huazhe Xu  
**Rethinking Transformers in Solving POMDPs**  
*In International Conference on Machine Learning (ICML) 2024*  
(27% acceptance rate)
33. 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)
34. 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)
35. 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)
36. 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)
37. 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)
38. 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)
39. 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)
  40. 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)
  41. 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)
  42. 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)
  43. 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)
  44. 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)
  45. 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)
  46. 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)
  47. 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)
  48. 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)
  49. 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)
  50. 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)
51. 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)
52. 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)
53. 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)
54. Yan Dai, Ruosong Wang, Simon S. Du  
**Variance-Aware Sparse Linear Bandits**  
*In International Conference on Learning Representations (ICLR) 2023*  
(32% acceptance rate)
55. 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)
56. 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)
57. 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)
58. 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)
59. 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)
60. 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)
61. 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)
62. 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)

63. Yifang Chen, Simon S. Du, Kevin Jamieson  
**Active Multi-Task Representation Learning**  
*In International Conference on Machine Learning (ICML) 2022*  
 (22% acceptance rate)
64. 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)
65. 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)
66. 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)
67. 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)
68. 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)
69. 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)
70. 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)
71. 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)
72. 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).
73. 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)
74. Yifang Chen, Simon S. Du, Kevin Jamieson  
**Corruption Robust Active Learning**  
*In Conference on Neural Information Processing Systems (NeurIPS) 2021*  
 (26% acceptance rate)

75. 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)
76. Jean Tarbouriech\*, Runlong Zhou\*, Simon S. Du, Matteo Pirodda, 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)
77. 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)
78. 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)
79. 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)
80. 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)
81. 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)
82. 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)
83. 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)
84. 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)
85. 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)
86. 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)
87. 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)
88. 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)
89. 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)
90. 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)
91. 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)
92. 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)
93. 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)
94. 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)
95. 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)
96. 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)
97. 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)
98. 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)
99. 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)
100. 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)
101. 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)
102. 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)
103. 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)
104. 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)
105. 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)
106. 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)
107. 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)
108. 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)
109. 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)
110. 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)
111. Simon S. Du\*, Xiyu Zhai\*, Barnabás Póczos, Aarti Singh.  
**Gradient Descent Provably Optimizes Over-parameterized Neural Networks**

- In International Conference on Learning Representations (ICLR) 2019*  
(33% acceptance rate)
112. 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)
113. 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)
114. 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)
115. 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)
116. 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)
117. 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).
118. 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)
119. 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)
120. 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).
121. Simon S. Du, Chi Jin, Jason D. Lee, Michael I. Jordan, 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)
122. 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)
123. 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)

124. 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*
125. 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)
126. 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)
127. Shashank Singh, Simon S. Du, Barnabás Póczos  
**Efficient Nonparametric Smoothness Estimation**  
*In Conference on Neural Information Processing Systems (NIPS) 2016*  
(23% acceptance rate)
128. 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).
129. 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. Zihan Zhang, Wenhao Zhan, Yuxin Chen, Simon S. Du, Jason D. Lee  
**Optimal Multi-Distribution Learning**  
*Journal of the ACM (JACM) 2025*
2. Zihan Zhang, Yuxin Chen, Jason D. Lee, Simon S. Du  
**Settling the Sample Complexity of Online Reinforcement Learning**  
*Journal of the ACM (JACM) 2025*
3. 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*
4. 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 (TMLR) 2024*
5. Wenqing Zheng\*, Hao Yang\*, Jiarui Cai, Peihao Wang, Xuan Jiang, Simon S. Du, Yin Hai Wang, Zhangyang Wang  
**Integrating Traffic Science with Representation Learning for City-wide Network Congestion Prediction**  
*Information Fusion 2023*

6. 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*
7. Bin Shi, Simon S. Du, Michale I. Jordan, Weijie Su  
**Understanding the Acceleration Phenomenon via High-Resolution Differential Equations**  
*Mathematical Programming (MP) 2021*
8. Yining Wang, Yi Wu, Simon S. Du  
**Near-Linear Time Local Polynomial Nonparametric Estimation with Box Kernels**  
*Infoms Journal on Computing 2020*
9. 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*

---

## TEACHING EXPERIENCES

### Instructor:

- Deep Learning (UW CSEP 590), Winter 2026
- Deep Learning (UW CSE 543), Fall 2025
- Deep Learning (UW CSE 543), Fall 2024
- 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:

- Reinforcement Learning (University of Washington), Fall 2024
- 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:

- |   |                |
|---|----------------|
| • Vector Runlong Zhou<br>UW CSE Ph.D.                             | 2022 - present |
| • Yancheng Liang<br>UW CSE Ph.D. (co-advsiad with Natasha Jaques) | 2023 - present |
| • Siting Li<br>UW CSE Ph.D.                                       | 2023 - present |
| • Divyansh Pareek<br>UW CSE Ph.D. (co-advised with Sewoong Oh)    | 2023 - present |

- 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.
- Runjia Li 2025 - present  
UW CSE Ph.D.
- Ruizhe Shi 2025 - present  
UW CSE Ph.D. (co-advised with Banghua Zhu)
- Avinandan Bose 2025 - present  
UW CSE Ph.D. (co-advised with Maryam Fazel)

**Postdoc Advising:**

- Mo Zhou 2024 - present  
co-advised with Maryam Fazel
- Xiyu Zhai 2024 - present
- Tom Yan 2025 - present  
co-advised with Natasha Jaques

**Former Ph.D. Students and Postdocs:**

- Yifang Chen 2020 - 2025  
UW CSE Ph.D. (co-advised with Kevin Jamieson) → Researcher at Microsoft AI
- Qiwen Cui 2021 - 2025  
UW CSE Ph.D. → Quantitative Researcher at Two Sigma
- Zihan Zhang 2023 - 2025  
UW/Princeton Postdoc (co-advised with Yuxin Chen and Jason D. Lee) → Assistant Professor at HKUST CSE
- Ruosong Wang 2022 - 2023  
UW Postdoc → Assistant Professor at PKU CFCS

**Former Undergraduate Student Researchers:**

- Ruizhe Shi 2023 - 2025  
THU IIIS Undergraduate → Visit at UW → UW CSE Ph.D.
- Natalia Zhang 2024  
THU IIIS Undergraduate → Visit at UW → Caltech CMS Ph.D.
- Mingrun Jiang 2024  
THU IIIS Undergraduate → Visit at UW → Duke CS Ph.D.
- Shiyuan Zhang 2024  
THU IIIS Undergraduate → Visit at UW → UCLA Ph.D.
- 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 Master → UCLA CS PhD
- 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:**

- Vector Runlong Zhou Fall 2025  
UW CSE Ph.D. General exam exam
- Arnab Maiti Spring 2023  
UW CSE Ph.D. Qualifying exam
- Hanwen Xu Fall 2025  
UW CSE Ph.D. General exam exam

UW CSE Ph.D. General exam exam	Fall 2025
• Zixuan Liu	
UW CSE Ph.D. General exam exam	Fall 2025
• Siting Li	
UW CSE Ph.D. Qualifying exam exam	Fall 2025
• Arnav Das	
UW ECE Ph.D. General exam exam	Fall 2025
UW ECE Ph.D. Qualifying exam	Fall 2022
• Xinqi Wang	
UW CSE Ph.D. Qualifying exam	Fall 2025
• Yiping Wang	
UW CSE Ph.D. Qualifying exam	Fall 2025
• Yancheng Liang	
UW CSE Ph.D. Qualifying exam	Spring 2025
• Ryan Huang	
UW ISE Ph.D. General exam	Winter 2025
• Tianhao Wu	
UC Berkeley Qualification exam	Winter 2025
• Qiwen Cui	
UW CSE Final exam	Winter 2025
UW CSE Ph.D. General exam	Winter 2024
UW CSE Ph.D. Qualifying exam	Spring 2023
• Trent Vonich	
UW ATOM Ph.D. General exam	Fall 2024
• Weihang Xu	
UW CSE Ph.D. Qualifying exam	Fall 2024
• Yifang Chen	
USE CSE Ph.D. General exam	Summer 2024
UW CSE Ph.D. Qualification exam	Spring 2022
• Xiyang Liu	
UW CSE Ph.D. Final exam	Summer 2024
UW CSE Ph.D. General exam	Spring 2024
• Farzam Ebrahimnejad	
UW CSE Ph.D. Final exam	Spring 2024
UW CSE Ph.D. General exam	Spring 2022
• Chenxi Liu	
UW CEE Ph.D. Final exam	Summer 2024
UW CEE Ph.D. General exam	Spring 2024
• Avinandan Bose	
UW CSE Ph.D. Qualifying exam	Spring 2024
• Addie Woicik	
UW CSE Ph.D. Final exam	Spring 2024
• 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. Final exam	Fall 2024
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
• 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
• Mayuree Binjolkar	
UW CEE Ph.D. General exam	Fall 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
• 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

---

Google Research	2025
• PI: Simon S. Du (UW)	
• Total funded amount: \$80,000 + \$70,000 computing credit	
Apple	2025
• PI: Simon S. Du (UW)	
• Total funded amount: \$300,000	
Cross-Pacific AI Initiative	2025
• Title: Inferring Strategic Behavior to Ensure Trustworthy Multi-agent AI Systems	
• PIs: Natasha Jaques (UW), Max Kleiman-Weiner (UW), Lillian Ratliff (UW), Simon S. Du (UW) and Kevin Jamieson (UW)	
• Total funded amount: \$720,000; Personal share: \$ 144,000 (2025 - 2027)	
Google Cloud	2025
• PI: Simon S. Du (UW)	
• Total funded amount: \$50,000	
Amazon Robotics	2025
• PI: Simon S. Du (UW)	

- Total funded amount: \$200,000 (2025 - 2026)
- Schmidt Science AI2050 Early Career Fellowship *2024*
- Title: Mathematical Foundations and Algorithms for Human-AI Cooperation
  - PI: Simon S. Du (UW)
  - Total funded amount: \$450,000
- 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

---

### **Global Convergence of Gradient EM for Over-Parameterized Gaussian Mixtures**

- IFDS Workshop on Theoretical Foundations of Applied AI, August 2025, Seattle, USA
- Tsinghua Statistics + AI Frontier Summit, July 2025, Beijing, China

### **Reinforcement Learning for Reasoning in Large Language Models with One Training Example**

- Chinese University of Hong Kong - Shenzhen, May 2025, Shenzhen, China
- Toyota Technological Institute of Chicago, May 2025, Pittsburgh, USA
- Carnegie Mellon University, May 2025, Pittsburgh, USA

### **Transformers are Efficient Compilers**

- Google Theory and Practice of Foundation Models Workshop, Nov 2024, Mountain View, USA

### **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**

- (Virtual) Mathematics of Information Processing Seminar, LMU Munich, June 2025
- Hong Kong Univeristy, May 2025, Hong Kong, China
- (Virtual) One World Seminar Series on the Mathematics of Machine Learning, June 2024
- 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?**

- Second Reinforcement Learning Theory Workshop, Edmonton, Canada
- 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

---

##### **Senior Area Chair:**

- International Conference on Learning Representations (ICLR) 2026
- International Conference on Machine Learning (ICML) 2026

##### **Area Chair or Senior Program Committee:**

- Conference on Artificial Intelligence (AAAI) 2020, 2021
- International Conference on Algorithmic Learning Theory (ALT) 2024, 2025
- Annual Conference on Learning Theory (COLT) 2021, 2022, 2023, 2024, 2025
- International Conference on Learning Representations (ICLR) 2021, 2022, 2023, 2024, 2025
- International Conference on Machine Learning (ICML) 2020, 2021, 2022, 2023, 2024, 2025
- International Joint Conferences on Artificial Intelligence (IJCAI) 2021
- Conference on Neural Information Processing Systems (NeurIPS) 2021, 2022, 2023, 2024, 2025

##### **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
- Information and Inference
- 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)
- SIAM Journal on Optimization (SIOPT)
- Statistical Science (STS)
- Transactions on Signal Processing (TSP)
- Conference:
  - Conference on Artificial Intelligence (AAAI)
  - Annual Meeting of the Association for Computational Linguistics (ACL)
  - 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:**

- NSF IIS-III Panel 2025

- NSF IIS-RI Panel 2024
- 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