

# Pan Xu

---

CONTACT INFORMATION Department of Computer Science University of California, Los Angeles  
EVI 295, 404 Westwood Plaza, Los Angeles, CA 90095 Phone: +1 (434) 806-6769  
Email: [panxu@cs.ucla.edu](mailto:panxu@cs.ucla.edu)  
<http://web.cs.ucla.edu/~panxu/>

RESEARCH INTERESTS My research is centered around **Machine Learning**, spanning areas of Artificial Intelligence, Data Science, Optimization, Reinforcement Learning, and High Dimensional Statistics. My ultimate research goal is to develop **computationally-** and **data-efficient** machine learning algorithms with both strong empirical performance and theoretical guarantees.

EDUCATION **University of California, Los Angeles (UCLA)**, CA, USA  
• Ph.D. candidate in Computer Science 2018/09 – Current  
**University of Virginia (UVA)** (Transfer Out), VA, USA  
• Ph.D. student in Computer Science 2017/07 – 2018/08  
• Ph.D. candidate in Systems and Information Engineering 2015/08 – 2017/06  
**University of Science and Technology of China (USTC)**, Hefei, China  
• Bachelor of Science in Mathematics 2011/08 – 2015/06

RESEARCH EXPERIENCE **Department of Computer Science, UCLA** 2018/09 – Current  
• Research Assistant (Advisor: Quanquan Gu)  
**Big Data Experience Lab, Adobe Research** 2019/06 – 2019/09  
• Research Intern (Mentors: Georgios Theodorou, Zheng Wen, Handong Zhao)  
**Department of Computer Science, UVA** 2017/07 – 2018/08  
• Research Assistant (Advisor: Quanquan Gu)  
**Department of Systems and Information Engineering, UVA** 2015/08 – 2017/06  
• Research Assistant (Advisor: Quanquan Gu)

HONORS AND AWARDS  
• **UCLA Outstanding Graduate Student Research Award** 2021  
• **Rising Stars in Data Science** 2021  
Host by University of Chicago, Center for Data and Computing  
• **Top Reviewer for ICML** 2020  
• **Top Reviewer for NeurIPS** 2019  
• **UCLA Graduate Division Fellowship** 2019, 2020, 2021  
• **Travel Award for UAI** 2019  
• **Travel Award for ICML** 2017, 2018  
• **Travel Award for NeurIPS** 2016, 2017, 2018, 2019, 2020  
• **UVA Presidential Fellowship in Data Science** 2016 – 2017  
• **UVA SIE Department Fellowship** 2015 – 2016  
• **USTC Excellent Thesis for Undergraduate** (85/1675, Top 5%) 2015  
*Thesis title:* High-Dimensional M-Estimator and Its Applications in the Estimation of Differential Graph Models  
• **USTC Outstanding Undergraduate Scholarship** 2012, 2013, 2014  
• **Outstanding China Young Volunteers** 2012  
• **Outstanding Freshman Scholarship** 2011

PUBLICATION (\* indicates equal contribution)  
**Peer Reviewed Journal Publications**  
[J1] Dongruo Zhou, **Pan Xu**, Quanquan Gu, Stochastic Nested Variance Reduction for Nonconvex Optimization. *Journal of Machine Learning Research (JMLR)* 21.103 (2020): 1-63.

- [J2] Dongruo Zhou, **Pan Xu**, Quanquan Gu, Stochastic Variance-Reduced Cubic Regularization Methods. *Journal of Machine Learning Research (JMLR)* 20.134 (2019): 1-47.

### Peer Reviewed Conference Proceedings

Summary: **13 first/co-first authored** papers (9 first authored, 4 co-first authored) on top-tire machine learning conferences: ICML, NeurIPS/NIPS, ICLR, UAI, AISTATS, AAAI.

- [C1] Tianyuan Jin, **Pan Xu**, Xiaokui Xiao, Quanquan Gu, Double Explore-then-Commit: Asymptotic Optimality and Beyond, In Proc. of the 34th Annual Conference on Learning Theory (**COLT 2021**), Boulder, Colorado, USA, 2021
- [C2] Difan Zou, **Pan Xu**, Quanquan Gu, Faster Convergence of Stochastic Gradient Langevin Dynamics for Non-Log-Concave Sampling, In Proc. of the 37th International Conference on Uncertainty in Artificial Intelligence (**UAI 2021**), Virtual, 2021.
- [C3] Tianyuan Jin, **Pan Xu**, Jieming Shi, Xiaokui Xiao, Quanquan Gu, MOTS: Minimax Optimal Thompson Sampling, In Proc. of the 38th International Conference on Machine Learning (**ICML 2021**), Virtual, 2021
- [C4] Tianyuan Jin, Jing Tang, **Pan Xu**, Keke Huang, Xiaokui Xiao and Quanquan Gu, Almost Optimal Anytime Algorithm for Batched Multi-Armed Bandits, In Proc. of the 38th International Conference on Machine Learning (**ICML 2021**), virtual, 2021
- [C5] Yue Wu, Weitong Zhang, **Pan Xu**, Quanquan Gu, A Finite-Time Analysis of Two Time-Scale Actor-Critic Methods, In Proc. of the 33rd Conference on Advances in Neural Information Processing Systems (**NeurIPS 2020**), Virtual, 2020. [Acceptance Rate: 20.1%]
- [C6] **Pan Xu**, Quanquan Gu, A Finite-Time Analysis of Q-Learning with Neural Network Function Approximation, In Proc. of the 37th International Conference on Machine Learning (**ICML 2020**), Virtual, 2020. [Acceptance Rate: 21.8%]
- [C7] **Pan Xu**, Felicia Gao, Quanquan Gu, Sample Efficient Policy Gradient Methods with Recursive Variance Reduction. In Proc. of the 8th International Conference on Learning Representations (**ICLR 2020**), Virtual, 2020. [Acceptance Rate: 26.5%]
- [C8] Tao Jin\*, **Pan Xu\***, Quanquan Gu, Farzad Farnoud, Rank Aggregation via Heterogeneous Thurstone Preference Models. In Proc. of the 34th Conference on Artificial Intelligence (**AAAI 2020**), New York, New York, USA, 2020. [**Oral, 348/7737 (4.5%)**]
- [C9] Difan Zou, **Pan Xu**, Quanquan Gu, Stochastic Gradient Hamiltonian Monte Carlo Methods with Recursive Variance Reduction. In Proc. of the 32nd Conference on Advances in Neural Information Processing Systems (**NeurIPS 2019**), Vancouver, Canada, 2019. [Acceptance Rate: 21.2%]
- [C10] **Pan Xu**, Felicia Gao, Quanquan Gu, An Improved Convergence Analysis of Stochastic Variance-Reduced Policy Gradient. In Proc. of the 35th International Conference on Uncertainty in Artificial Intelligence (**UAI 2019**), Tel Aviv, Israel, 2019. [**Oral, 35/450 (7.8%)**]
- [C11] Difan Zou, **Pan Xu**, Quanquan Gu, Sampling from Non-Log-Concave Distributions via Variance-Reduced Gradient Langevin Dynamics. In Proc. of the 22nd International Conference on Artificial Intelligence and Statistics (**AISTATS 2019**), Naha, Okinawa, Japan, 2019. [Acceptance Rate: 32.4%]
- [C12] **Pan Xu\***, Jinghui Chen\*, Difan Zou, Quanquan Gu, Global Convergence of Langevin Dynamics Based Algorithms for Nonconvex Optimization. In Proc. of the 31st Conference on Advances in Neural Information Processing Systems (**NeurIPS 2018**), Montréal, Canada, 2018. [**Spotlight, 168/4856 (3.5%)**]
- [C13] Dongruo Zhou, **Pan Xu**, Quanquan Gu, Stochastic Nested Variance Reduction for Nonconvex Optimization. In Proc. of the 31st Conference on Advances in Neural Information Processing Systems (**NeurIPS 2018**), Montréal, Canada, 2018. [**Spotlight, 168/4856 (3.5%)**]
- [C14] Yaodong Yu\*, **Pan Xu\***, Quanquan Gu, Third-order Smoothness Helps: Even Faster

- Stochastic Optimization Algorithms for Finding Local Minima. In Proc. of the 31st Conference on Advances in Neural Information Processing Systems (**NeurIPS 2018**), Montréal, Canada, 2018. [Acceptance Rate: 20.8%]
- [C15] Difan Zou\*, **Pan Xu\***, Quanquan Gu, Subsampled Stochastic Variance-Reduced Gradient Langevin Dynamics. In Proc. of the 34th International Conference on Uncertainty in Artificial Intelligence (**UAI 2018**), Monterey, California, USA, 2018. [Acceptance Rate: 30.9%]
- [C16] **Pan Xu\***, Tianhao Wang\*, Quanquan Gu, Continuous and Discrete-time Accelerated Stochastic Mirror Descent for Strongly Convex Functions. In Proc. of the 35th International Conference on Machine Learning (**ICML 2018**), Stockholm, Sweden, 2018. [Acceptance Rate: 25%]
- [C17] Difan Zou\*, **Pan Xu\***, Quanquan Gu, Stochastic Variance-Reduced Hamilton Monte Carlo Methods. In Proc. of the 35th International Conference on Machine Learning (**ICML 2018**), Stockholm, Sweden, 2018. [Acceptance Rate: 25%]
- [C18] Dongruo Zhou, **Pan Xu**, Quanquan Gu, Stochastic Variance-Reduced Cubic Regularized Newton Method. In Proc. of the 35th International Conference on Machine Learning (**ICML 2018**), Stockholm, Sweden, 2018. [Acceptance Rate: 25%]
- [C19] Jinghui Chen, **Pan Xu**, Lingxiao Wang, Jian Ma, Quanquan Gu, Covariate Adjusted Precision Matrix Estimation via Nonconvex Optimization. In Proc. of the 35th International Conference on Machine Learning (**ICML 2018**), Stockholm, Sweden, 2018. **[Long Talk, 212/2473 (8.6%)]**
- [C20] **Pan Xu\***, Tianhao Wang\*, Quanquan Gu, Accelerated Stochastic Mirror Descent: From Continuous-time Dynamics to Discrete-time Algorithms. In Proc. of the 21st International Conference on Artificial Intelligence and Statistics (**AISTATS 2018**), Playa Blanca, Lanzarote, Canary Islands, 2018. [Acceptance Rate: 33.1%]
- [C21] **Pan Xu**, Jian Ma, Quanquan Gu, Speeding Up Latent Variable Gaussian Graphical Model Estimation via Nonconvex Optimizations. In Proc. of the 30th Conference on Advances in Neural Information Processing Systems (**NIPS 2017**), Long Beach, CA, USA, 2017. [Acceptance Rate: 20.9%]
- [C22] Aditya Chaudhry, **Pan Xu**, Quanquan Gu, Uncertainty Assessment and False Discovery Rate Control in High-Dimensional Granger Causal Inference. In Proc. of the 34th International Conference on Machine Learning (**ICML 2017**), Sydney, Australia, 2017. [Acceptance Rate: 25.4%]
- [C23] **Pan Xu**, Tingting Zhang, Quanquan Gu, Efficient Algorithm for Sparse Tensor-variate Gaussian Graphical Models via Gradient Descent. In Proc. of the 20th International Conference on Artificial Intelligence and Statistics (**AISTATS 2017**), Fort Lauderdale, Florida, USA, 2017. [Acceptance Rate: 31.6%]
- [C24] **Pan Xu**, Quanquan Gu, Semiparametric Differential Graph Models. In Proc. of the 29th Conference on Advances in Neural Information Processing Systems (**NIPS 2016**), Barcelona, Spain, 2016. [Acceptance Rate: 22.7%]
- [C25] Lu Tian, **Pan Xu**, Quanquan Gu, Forward Backward Greedy Algorithms for Multi-Task Learning with Faster Rates. In Proc of the 32nd International Conference on Uncertainty in Artificial Intelligence (**UAI 2016**), New York / New Jersey, USA, 2016. [Acceptance Rate: 31%]

### Workshop Presentations

- [W1] Difan Zou, Lingxiao Wang, **Pan Xu**, Jinghui Chen, Weitong Zhang, Quanquan Gu, Epidemic Model Guided Machine Learning for COVID-19 Forecasts in the United States, ICLR 2021 Machine Learning for Preventing and Combating Pandemics Workshop, Virtual, 2021
- [W2] Tianyuan Jin, **Pan Xu**, Xiaokui Xiao, Quanquan Gu, Double Explore-then-Commit: Asymptotic Optimality and Beyond, NeurIPS 2020 Offline Reinforcement Learning Workshop, Virtual, 2020
- [W3] Yue Wu, Weitong Zhang, **Pan Xu**, Quanquan Gu, A Finite-Time Analysis of Two

Time-Scale Actor-Critic Methods, ICML 2020 Theoretical Foundations of Reinforcement Learning Workshop, Virtual, 2020

- [W4] **Pan Xu**, Felicia Gao, Quanquan Gu, Sample Efficient Policy Gradient Methods with Recursive Variance Reduction, NeurIPS 2019 Optimization Foundations for Reinforcement Learning Workshop, Vancouver, Canada, 2019

### Preprints

- [P1] Estee Y Cramer, **Pan Xu**, et al., Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the US, medRxiv: 2021.02.03.21250974
- [P2] **Pan Xu**, Zheng Wen, Handong Zhao, Quanquan Gu, Neural Contextual Bandits with Deep Representation and Shallow Exploration, arXiv:2012.01780
- [P3] Katriona Shea, **Pan Xu**, et al., COVID-19 Reopening Strategies at the County Level in the Face of Uncertainty: Multiple Models for Outbreak Decision Support, medRxiv: 2020.11.03.20225409
- [P4] COVID-19 Forecast Hub Consortium, **Pan Xu**, Ensemble Forecasts of Coronavirus Disease 2019 (COVID-19) in the U.S., medRxiv:2020.08.19.20177493
- [P5] **Pan Xu**, Lu Tian, Quanquan Gu, Communication-efficient Distributed Estimation and Inference for Transelliptical Graphical Models, arXiv: 1612.09297 (Presented at the ENAR 2016 Spring Meeting)

### PROJECTS

#### Combating COVID-19 using machine learning

- Description: we launched the project in March 2020 to (1) help people stay informed about the spread of coronavirus by creating data visualization of confirmed, dead, and hospitalized cases of COVID-19; and (2) provide guidance to policy makers and individuals by using a machine learning empowered epidemic model to forecast the spread of the epidemic. (Publications: [P1, P3, P4, W1])
- Website: <https://covid19.uclaml.org/>
- Impact: our model is used in the official forecast by US Centers for Disease Control and Prevention (CDC), California Department of Public Health (CDPH), The US COVID-19 Forecast Hub, German and Polish COVID-19 Forecast Hub.
- Media coverage: UCLA Newsroom, FiveThirtyEight, TPM, POLITICO, CBS News 8, New Yorker

### TALKS AND PRESENTATIONS

#### Tutorials:

- **Nonconvex Optimization for Knowledge Discovery and Data Mining**  
SIAM International Conference on Data Mining, Calgary, Canada 2019/05

#### Invited Talks:

- **Sample-Efficient Nonconvex Optimization Algorithms for Machine Learning**  
University of Michigan, Ann Arbor (Virtual) 2021/03
- **Sample-Efficient Nonconvex Optimization Algorithms for Machine Learning**  
University of Minnesota Twin Cities (Virtual) 2021/03
- **Sample-Efficient Nonconvex Optimization Algorithms for Machine Learning**  
Duke University (Virtual) 2021/03
- **Sample-Efficient Nonconvex Optimization Algorithms for Machine Learning**  
University of Pittsburgh (Virtual) 2021/03
- **Data Efficient Optimization in Reinforcement Learning**  
University of Chicago (Virtual) 2021/01
- **Sample-Efficient Optimization Algorithms for Reinforcement Learning**  
California Institute of Technology (Virtual) 2020/12
- **A Finite-Time Analysis of Two Time-Scale Actor-Critic Methods**  
Peking University (Virtual) 2020/10
- **Optimization from A Continuous-time View**  
University of California, Los Angeles 2018/10

- **Evaluating the Effects of Social Relationships on Brain Activities via High-Dimensional Differential Graphical Models**  
University of Virginia 2017/03

#### Conference Presentations:

- **A Finite-Time Analysis of Q-Learning with Neural Network Function Approximation**  
International Conference on Machine Learning (Virtual) 2020/07
- **Sample Efficient Policy Gradient Methods with Recursive Variance Reduction**
  - International Conference on Learning Representations (Virtual) 2020/04
  - Conference on Reinforcement Learning for Real Life (Virtual) 2020/06
- **An Improved Convergence Analysis of Stochastic Variance-Reduced Policy Gradient**  
International Conference on Uncertainty in Artificial Intelligence, Tel Aviv, Israel 2019/07
- **Global Convergence of Langevin Dynamics Based Algorithms for Nonconvex Optimization**  
International Conference on Advances in Neural Information Processing Systems, Montréal, Canada 2018/12
- **Stochastic Variance-Reduced Hamilton Monte Carlo Methods**  
International Conference on Machine Learning, Stockholm, Sweden 2018/07
- **Continuous and Discrete-time Accelerated Stochastic Mirror Descent for Strongly Convex Functions**  
International Conference on Machine Learning, Stockholm 2018/07
- **Uncertainty Assessment and False Discovery Rate Control in High-Dimensional Granger Causal Inference**  
International Conference on Machine Learning, Sydney, Australia 2017/07

#### PROFESSIONAL SERVICES

##### Conference/Event Organization:

- Volunteer, UAI 2019, ICLR 2020 Queer in AI workshop
- **Co-chair**, oSTEM and Queer in AI Graduate Applications Financial Aid Program, 2020
- **Co-chair**, oSTEM and Queer in AI Graduate Admissions Application Review Program, 2020
- Core organizer, Queer in AI, 2021-2022

##### Senior Program Committee

- International Joint Conference on Artificial Intelligence (IJCAI), 2021

##### Program Committee

- AAAI Conference on Artificial Intelligence (AAAI), 2020, 2021
- International Conference on Uncertainty in Artificial Intelligence (UAI), 2019, 2020, 2021
- International Joint Conference on Artificial Intelligence (IJCAI), 2020, 2021
- Asia Conference on Machine Learning (ACML), 2019, 2020
- IEEE International Conference on Big Data (BigData), 2020

##### Conference Reviewer

- International Conference on Machine Learning (ICML), 2019, 2020, 2021
- Neural Information Processing Systems (NeurIPS), 2019, 2020
- International Conference on Artificial Intelligence and Statistics (AISTATS), 2020, 2021
- ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD), 2019
- International Conference on Learning Representations (ICLR), 2020, 2021
- International Conference on Algorithmic Learning Theory (ALT), 2020
- IEEE International Conference on Big Data (BigData), 2019
- Mathematical and Scientific Machine Learning (MSML), 2021
- OPT Workshop on Optimization for Machine Learning (OPT), 2020

## Journal Reviewer

- IEEE Transactions on Knowledge and Data Engineering (TKDE)
- IEEE Transactions on Automatic Control
- IEEE Transactions on Neural Networks and Learning Systems
- IEEE Transactions on Communications
- IEEE Journal on Selected Areas in Information Theory (JSAIT)
- ACM Computing Surveys (CSUR)
- Journal of the Royal Statistical Society, Series B (JRSS-SB)
- Journal of Machine Learning Research (JMLR)
- Statistics and Computing (STCO)
- Optimization Methods and Software
- Springer Machine Learning Journal
- Neurocomputing
- PLoS ONE
- Applied Mathematical Modeling
- Social Network Analysis and Mining (SNAM)

TEACHING EXPERIENCE	<b>Teaching Assistant, Department of Computer Science, UCLA</b>	
	• CS180: Introduction to Algorithms and Complexity	Spring 2020
	• CS32: Introduction to Computer Science II	Winter 2020
	• CS31: Introduction to Computer Science I	Fall 2019
	<b>Teaching Assistant, Department of Computer Science, UVA</b>	
	• CS6316/SYS6016: Machine Learning	Spring 2018
	<b>Teaching Assistant, Department of Systems and Information Engineering, UVA</b>	
	• SYS3060: Introduction to Reinforcement Learning	Spring 2016

MENTORING	Master student:	
	• Felicia Gao (Master at UCLA)	
	<i>Project:</i> Improved analysis and new algorithms of stochastic variance reduced policy gradient methods [C10, C7]	
	• Yihe Deng (Master at UCLA)	
	<i>Project:</i> Hypergraphs representation learning with applications in genomics data	
	• Yu Yang (Master at UCLA)	
	<i>Project:</i> Hypergraphs representation learning with applications in genomics data	
	Undergraduate students:	
	• Tianhao Wang (Undergraduate at USTC, now PhD at Yale)	
	<i>Project:</i> Continuous-time dynamics discrete-time algorithms for accelerated stochastic mirror descent [C20, C16]	

PROGRAMMING SKILLS  $\LaTeX$ , MatLab, Python, C++, R, Mathematica

REFERENCES Available upon request.