Computer Science Deparment
UC Santa Barbara, CA 93107
☎ *+1-412-251-6037*
✉ *yuxiangw@cs.ucsb.edu*
www.cs.ucsb.edu/ yuxiangw/
Google scholar profile

Yu-Xiang Wang

I am an assistant professor of computer science in UC Santa Barbara. I direct UCSB's Machine Learning Lab and have co-founded the Center of Responsible Machine Learning. My area of research is statistical machine learning with special focus on differential privacy, scalable machine learning, reinforcement learning and adaptive online learning. My work had been widely cited and adopted in various applications.

Education

2013–2017 **PhD in Statistics and Machine Learning**, *Carnegie Mellon University*.
2013–2016 **MS in Machine Learning**, *Carnegie Mellon University*.
2012–2013 **MEng in Electrical Engineering**, *National University of Singapore*.
2007–2011 **BEng in Electrical Engineering**, *National University of Singapore*.

Work experience

2018–present **Eugene Aas Chair Assistant Professor of Computer Science**, UC Santa Barbara.
2019–present **Assistant Professor**, *Statistics Department (By Courtesy)*, UC Santa Barbara, CA.
2019–present **Founding Co-Director**, *Center for Responsible Machine Learning*, UC Santa Barbara.
2017–2018 **Applied Scientist**, *Amazon AWS*, Palo Alto, CA.
2016 **Research intern**, *Microsoft Research*, NYC.

Selected Awards

2021 **NSF Award: SCALE-MoDL: Theory of Adaptivity in Deep Neural Networks**.
2021 **COLT'21 Best Student Paper Award**.
2021 **NSF CAREER Award: Exact Optimal and Data-Adaptive Methods and Tools for Differential Privacy**.
2020 **NSF Award: Towards Optimal and Adaptive Reinforcement Learning with Offline Data and Limited Adaptivity**.
2020 **NSF Award: MLWiNS: RL-based Self-driving Wireless Network Management System for QoE Optimization.**.
2020 **NSF Award: Interventional COVID-19 Response Forecasting in Local Communities Using Neural Domain Adaptation Models**.
2019 **AISTATS'19 Notable Paper Award**.
2019-2021 **Faculty research awards: Adobe Data Science Award, Amazon ML Research Award, Google Research Scholar Award, NEC Labs, Evidation Health, Appfolio**.
2018 **NSF Award: HDR Institute on Data-Driven methods for Material Discovery**.
Before 2017 **Paper awards: WSDM'16, KDD'15. Reviewer award: NeurIPS'14**.
Before 2017 **Fellowship received: CMU, Baidu. Fellowship finalist: MSR, Facebook**.

Peer-Reviewed Publications

2022

[1] Jianyu Xu, **Yu-Xiang Wang**. Towards Agnostic Feature-based Dynamic Pricing: Linear Policies vs Linear Valuation with Unknown Noise. In AISTATS 2022.

[2] Yuqing Zhu, **Yu-Xiang Wang**. Adaptive Private-K-Selection with Adaptive K and Application to Multi-label PATE. In AISTATS 2022.

[3] Yuqing Zhu, Jinshuo Dong, **Yu-Xiang Wang**. Optimal Accounting of Differential Privacy via Characteristic Function. In AISTATS 2022.

[4] Dheeraj Baby, **Yu-Xiang Wang**. Optimal Dynamic Regret in Exp-Concave Online Learning: Proper Learning Under Box Constraints. In AISTATS 2022.

[5] Peng Zhao, **Yu-Xiang Wang**, Zhi-Hua Zhou. Non-stationary Online Learning with Memory and Non-stochastic Control. In AISTATS 2022.

2021

[6] Jianyu Xu, **Yu-Xiang Wang**. Logarithmic Regret in Feature-Based Dynamic Pricing. In NeurIPS 2021.

[7] Rachel Redberg, **Yu-Xiang Wang**. Privately Publishable Per-instance Privacy. In NeurIPS 2021.

[8] Ming Yin, **Yu-Xiang Wang**. Towards Instance-Optimal Offline Reinforcement Learning with Pessimism. In NeurIPS 2021.

[9] Ming Yin, **Yu-Xiang Wang**. Optimal Uniform OPE and Model-based Offline Reinforcement Learning in Time-Homogeneous, Reward-Free and Task-Agnostic Settings. In NeurIPS 2021.

[10] Ming Yin, Yu Bai, **Yu-Xiang Wang**. Near-Optimal Offline Reinforcement Learning via Double Variance Reduction. In NeurIPS 2021.

[11] Dheeraj Baby, **Yu-Xiang Wang**. Optimal Dynamic Regret in Exp-Concave Online Learning. In COLT 2021. [Best Student Paper Award].

[12] Dheeraj Baby, Xuandong Zhao, **Yu-Xiang Wang**. An Optimal Reduction of TV-Denoising to Adaptive Online Learning. In AISTATS 2021.

[13] Chong Liu, Yuqing Zhu Kamalika Chaudhuri, **Yu-Xiang Wang**. Revisiting Model-Agnostic Private Learning: Faster Rates and Active Learning. In AISTATS 2021 and the Journal of Machine Learning Research.

[14] Ming Yin, Yu Bai, **Yu-Xiang Wang**. Near Optimal Provable Uniform Convergence in Off-Policy Evaluation for Reinforcement Learning. In AISTATS 2021.

[15] Xiaoyong Jin, **Yu-Xiang Wang**, Xifeng Yan. Inter-Series Attention Model for COVID-19 Forecasting. In SDM 2021.

[16] Hojjat Aghakhani, Dongyu Meng, **Yu-Xiang Wang**, Christopher Kruegel, and Giovanni Vigna. Bullseye Polytope: A Scalable Clean-Label Poisoning Attack with Improved Transferability.. In IEEE EuroS&P 2021.

[17] Chong Liu, **Yu-Xiang Wang**. Doubly Robust Crowdsourcing. In Journal of Artificial Intelligence Research.

2020

[18] Dheeraj Baby, **Yu-Xiang Wang**. Adaptive Online Estimation of Piecewise Polynomial Trends. In NeurIPS 2020.

[19] Yuqing Zhu, **Yu-Xiang Wang**. Improving Sparse Vector Technique with Renyi Differential Privacy. In NeurIPS 2020.

[20] Remi Tachet des Combes, Han Zhao, **Yu-Xiang Wang**, Geoff Gordon. Domain Adaptation with Conditional Distribution Matching and Generalized Label Shift. In NeurIPS 2020.

[21] Christopher DeCarolis, Mukul Ram, Seyed Esmaeili, **Yu-Xiang Wang**, Furong Huang. An end-to-end Differentially Private Latent Dirichlet Allocation Using a Spectral Algorithm. In ICML 2020.

[22] Yuqing Zhu, Xiang Yu, Manmohan Chandraker **Yu-Xiang Wang**. Private-kNN: Practical Differential Privacy for Computer Vision. In CVPR'20.

[23] Ming Yin, **Yu-Xiang Wang**. Asymptotically Efficient Off-Policy Evaluation for Tabular Reinforcement Learning. In AISTATS'20.

2019

[24] Dheeraj Baby, **Yu-Xiang Wang**. Online Forecasting of Total Variation-Bounded Sequences. In NeurIPS 2019. Short version received the [Best-Paper Honorable Mention] from ICML'19 Time-Series Workshop.

[25] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyou Zhou, Wenhu Chen, **Yu-Xiang Wang**, Xifeng Yan. Enhancing the Locality and Breaking the Memory Bottleneck of Transformer on Time Series Forecasting. In NeurIPS 2019.

[26] Tengyang Xie, Yifei Ma, **Yu-Xiang Wang**. Optimal Off-Policy Evaluation for Reinforcement Learning with Marginalized Importance Sampling. In NeurIPS 2019.

[27] Yu Bai, Tengyang Xie, Nan Jiang, **Yu-Xiang Wang**. Provably Efficient Q-Learning with Low Switching Cost. In NeurIPS 2019.

[28] Yuqing Zhu, **Yu-Xiang Wang**. Poisson Subsampled Rényi Differential Privacy. In ICML'19.

[29] **Yu-Xiang Wang**, Borja Balle, Shiva Kasiviswanathan. Subsampled Rényi Differential Privacy and Analytical Moments Accountant. In AISTATS'19. [Notable Paper Award]

[30] Yifei Ma, **Yu-Xiang Wang**, Balakrishnan Narayanaswamy. Imitation-Regularized Offline Learning. In AISTATS'19.

[31] Veeranjaneyulu Sadhanala, **Yu-Xiang Wang**, Aaditya Ramdas, Ryan Tibshirani. A Higher-Order Kolmogorov-Smirnov Test . In AISTATS'19.

[32] Xi Chen, Yining Wang, **Yu-Xiang Wang**. Non-stationary Stochastic Optimization under $L_{p,q}$-Variation Measures . In *Operations Research*, 2019.

[33] **Yu-Xiang Wang**, Huan Xu and Chenlei Leng. Provable Subspace Clustering: When LRR meets SSC. In *IEEE Transaction of Information Theory*, 2019.

2018

[34] Borja Balle, **Yu-Xiang Wang**. Improving Gaussian Mechanism for Differential Privacy: Analytical Calibration and Optimal Denoising . In ICML'18.

[35] Zack Lipton*, **Yu-Xiang Wang***, Alex Smola . Detecting and Correcting Label-Shift with Black Box Predictors . In ICML'18.

[36] Jeremy Bernstein, **Yu-Xiang Wang**, Kamyar Azizzadenesheli, Anima Anandkumar. signSGD: compressed optimisation for non-convex problems . In ICML'18.

[37] **Yu-Xiang Wang**. Revisiting differentially private linear regression: optimal and adaptive prediction and estimation in unbounded domain. in UAI'18.

[38] **Yu-Xiang Wang**. Per Instance Differential Privacy. In *Journal of Confidentiality and Privacy*, 2018.

[39] Yining Wang, **Yu-Xiang Wang** and Aarti Singh. A Theoretical Analysis for Noisy Sparse Subspace Clustering for Dimension-reduced data. In *IEEE Transaction of Information Theory*, 2018.

2016 - 2017

[40] Veeranjaneyulu Sadhanala*, **Yu-Xiang Wang***, Ryan Tibshirani, James Sharonack. Higher-Order Total Variation Classes on Grids: Minimax Theory and Trend Filtering Methods. In *Advances in Neural Information Processing Systems (NIPS)*, 2017

[41] **Yu-Xiang Wang**, Alekh Agarwal, Miro Dudik. Optimal and Adaptive Off-policy Evaluation in Contextual Bandit. In *International Conference on Machine Learning (ICML)*, 2017.

[42] Ziqi Liu, Alex Smola, Kyle Soska, **Yu-Xiang Wang**, Jun Zhu. Attributing Hacks with Survival Trend Filtering. In *Electronic Journal of Statistics, 2017 (Short version appeared at AISTATS'2017).*

[43] **Yu-Xiang Wang**, Jing Lei, Steve Fienberg. Learning with Differential Privacy: Stability, Learnability and the Sufficiency and Necessity of ERM Principle. In *JMLR, 2016.*

[44] **Yu-Xiang Wang**, James Sharpnack, Alex Smola and Ryan Tibshirani. Trend Filtering on Graphs. In *JMLR, 2016 (Short version appeared at AISTATS'2015)*

[45] **Yu-Xiang Wang** and Huan Xu. Noisy Sparse Subspace Clustering. In *JMLR, 2016 (Short version appeared at ICML'2013)*

[46] Veeranjaneyulu Sadhanala*, **Yu-Xiang Wang***, Ryan Tibshirani. Total Variation Classes Beyond 1d: Minimax Rates, and the Limitations of Linear Smoothers. In *Advances in Neural Information Processing Systems (NIPS)*, 2016

[47] **Yu-Xiang Wang**, Jing Lei, Steve Fienberg. On-Average KL-Privacy and its equivalence to Generalization for Max-Entropy Mechanisms. Privacy in Statistical Databases, 2016.

[48] **Yu-Xiang Wang**, Veeru Sadhanala, Wei Dai, Willie Neiswanger, Eric Xing. Parallel and Distributed Block-Coordinate Frank-Wolfe. In *International Conference on Machine Learning (ICML)*, 2016

[49] Veeranjaneyulu Sadhanala*, **Yu-Xiang Wang***, Ryan Tibshirani. Graph Sparsification Approaches for Large-Scale Laplacian Smoothing.. AISTATS'2016.

[50] Yining Wang, **Yu-Xiang Wang**, Aarti Singh. Graph Connectivity in Noisy Sparse Subspace Clustering. AISTATS'2016.

[51] Mu Li, Ziqi Liu, Alex Smola, **Yu-Xiang Wang**. DiFacto — Distributed Factorization Machines. In *ACM International Conference on Web Search and Data Mining (WSDM)*, 2016 [Best Paper Honorable Mention]

2016 and before

[52] Yining Wang, **Yu-Xiang Wang** and Aarti Singh. Differentially Private Subspace Clustering. In *Advances in Neural Information Processing Systems (NIPS)*, 2015

[53] **Yu-Xiang Wang**, Steve Fienberg and Alex Smola. Privacy for Free: Posterior Sampling and Stochastic Gradient Monte Carlo. In *International Conference on Machine Learning (ICML)*, 2015

[54] Ziqi Liu, **Yu-Xiang Wang**, Alex Smola. Fast Differential Private Matrix Factorization. In RecSys'15.

[55] Seth Flaxman, **Yu-Xiang Wang**, Alex Smola. Who Supported Obama in 2012? Ecological inference through distribution regression. In KDD'15. [Best Student Paper Award]

[56] Yining Wang, **Yu-Xiang Wang** and Aarti Singh. A Deterministic Analysis for Sparse Subspace Clustering for Dimension-reduced data. In *International Conference on Machine Learning (ICML)*, 2015

[57] **Yu-Xiang Wang**, Alex Smola, Ryan Tibshirani. The Falling Factorial Basis and Its Statistical Applications. In *International Conference on Machine Learning (ICML)*, 2014

[58] **Yu-Xiang Wang**, Choon Meng Lee, Loong Fah Cheong and Kim-Chuan Toh. Practical Matrix Completion and Corruptions Recovery using Alternating Robust Subspace Minimization. In *International Journal on Computer Vision (IJCV), 2014*

[59] Zhi Gao, Loong Fah Cheong and **Yu-Xiang Wang**, Block-Sparse RPCA for Salient Motion Detection. in IEEE TPAMI, 2013

[60] **Yu-Xiang Wang**, Huan Xu and Chenlei Leng. Provable Subspace Clustering: When LRR meets SSC. In *Advances in Neural Information Processing Systems (NIPS)*, 2013

[61] **Yu-Xiang Wang** and Huan Xu. Noisy Sparse Subspace Clustering. In *International Conference on Machine Learning (ICML)*, 2013.

[62] **Yu-Xiang Wang** and Huan Xu. Stability of Matrix Factorization for Collaborative Filtering. In *International Conference on Machine Learning (ICML)*, 2012

[63] **Yu-Xiang Wang**, An Efficient Algorithm for UAV Indoor Pose Estimation using Vanishing Geometry. In *IAPR Conference on Machine Vision Application (MVA)*, 2011

Invited Talks

1. 2021 Invited talk at Rutgers University on "Optimal Accounting of Differential Privacy".
2. 2021 Invited talk at Joint Statistics Meeting on "Adaptive Online Forecasting of Trends".

3. 2021 Invited talk at Google Differential Privacy Seminar: "Privately Publishable Per-Instance Differential Privacy"
4. 2021 Invited talk at Amazon Web Services: “Advances in Differential Privacy and Private Federated Learning”
5. 2021 Invited talk at State University of New York at Albany: “Per-Instance Differential Privacy and how to publish them”
6. 2021 Invited talk at Google Brain: “AutoDP”
7. 2020 Simons Institute Workshop on RL with offline data and simulation: “Uniform convergence in offline Policy evaluation”
8. 2020 RL Theory Seminar series “Uniform convergence in offline Policy evaluation”
9. 2020 Invited Discussion on "Gaussian Differential Privacy" by Dong, Su and Roth at the International Seminar on Selective Inference.
10. 2020 Invited talk at “COVID Response Seminar Series”
11. 2020 Invited talk at IEE-Microsoft Azure joint seminar series on “Offline Reinforcement Learning”
12. 2019 Invited talk at NEC Labs on “Differential privacy”
13. 2019 UCSD Talk on: “Online Forecasting”
14. 2019 Berkeley Simons Institute workshop on Privacy and the Science of Data Analysis: “Privacy Amplification by subsampling and Renyi Differential Privacy”
15. 2019 UCSB Statistics Seminar “Off-policy Evaluation and Learning in theory and in the wild”.
16. 2018 Privacy in Graphs (PiG workshop)“Per-Instance Differential Privacy”, UC Santa Cruz.
17. 2018 Caltech“Off-policy Evaluation and Learning in theory and in the wild”.
18. 2018 Harvard University “DP Deployed” meeting.
19. 2018 Berkeley Simons Institute workshop on Adaptive data analysis: “Gaussian Adaptive Data Analysis and Linear Bandits”.
20. 2018 Symposium on Statistics and Data Science: “Off policy evaluation and learning in causal inference problems”
21. 2018 UCSB Computer Science Seminar: “SignSGD and Signum algorithm for nonconvex optimization and a tutorial for effort-free deep learning with Gluon”
22. 2017 CMU Statistics Seminar: “SignSGD and Signum algorithm for nonconvex optimization and a tutorial for effort-free deep learning with Gluon”
23. 2017 CMU SMLRG: “Sequential Selective Estimation (Gaussian Adaptive Data Analysis) and linear bandits”
24. 2017 Nanyang Technological University: “Towards practical Machine Learning with Differential Privacy and Beyond”
25. 2017 National University of Singapore:“Attributing Hacks with Survival Trend Filtering”
26. 2017 National University of Singapore:“Towards practical Machine Learning with Differential Privacy and Beyond”
27. 2017 Invited talk: “Towards practical Machine Learning with Differential Privacy and Beyond” at Michigan, Cornell, UC Santa Barbara, Penn State U, UT Austin, Microsoft Research New England / NYC, Yahoo Research.

28. 2017 Penn State: “Trend Filtering on Graphs”
29. 2016 JSM Chicago: “Learning with Differential Privacy”
30. 2016 ICML Privacy Workshop: “Learning with Differential Privacy and On-Average KL-privacy”
31. 2016 Columbia U: “Trend Filtering and Optimal TV-Denoising”
32. 2016 Office of Financial Research, D.C.: “Practical machine learning with Differential Privacy and Beyond”
33. 2015 Google Pittsburgh: “Practical machine learning with Differential Privacy and Beyond”
34. 2015 Carnegie Mellon University ML seminar: “Trend Filtering, Falling Factorial Basis and Adaptive Statistical Estimation on Graphs”
35. 2015 National University of Singapore: “Trend Filtering: Some Recent Advances and Challenges”
36. 2015 Singapore Management University: “On Trend filtering and Differential Privacy”
37. 2015 CMU Math seminar: “Trend Filtering: Some Recent Advances and Challenges”
38. 2014 Center for Urban Science and Progress, NYC: “Learning with Differential Privacy”
39. 2013 Columbia U: “Noisy Sparse Subspace Clustering”

Students and Alumni

2017-2018 **Jeremy Bernstein**, *Intern at Amazon.*
“SignSGD” paper. Now PhD candidate at Caltech.

2018 **Yu Bai**, *Intern at Amazon.*
“ProxQuant” and “RL with Low-Switching Cost”. Now Senior Research Scientist at Salesforce Research.

2018 **Tengyang Xie**, *Intern at Amazon.*
“Marginalized Importance Sampling” paper. Now PhD student at UIUC.

2018-present **Dheeraj Baby**, *PhD Student.*
Works on forecasting and online learning. Received COLT’21 Best Student Paper Award. ICML’19 Time Series Workshop Best Paper Honorable Mention

2018-present **Yuqing Zhu**, *PhD Student.*
Works on differential privacy and autodp. Received Google PhD Fellowship

2018-present **Chong Liu**, *PhD Student.*
Works on active learning, ensemble learning and global optimization.

2019-present **Ming Yin**, *PhD Student.*
Works on Offline Reinforcement Learning.

2019-present **Rachel Redberg**, *PhD Student.*
Works on differential privacy.

2019-present **Xuandong Zhao**, *PhD Student.*
Works on differential privacy + NLP.

2019-present **Jianyu Xu**, *PhD Student.*
Works on dynamic pricing and other problems with structured feedback.

2020-present **Kaiqi Zhang**, *PhD Student.*
Works on theory of adaptivity in deep learning.

2020 **Peng Zhao**, *Visiting Researcher.*
“Nonstationary Non-Stochastic Control” paper. Now Assistant Professor in Nanjing University (Top 2 institute for ML research in China).

2021 **Zhiqi Bu**, *Intern at Amazon.*
Worksed on privacy in large language models. Recent PhD graduate from UPenn.

Professional activities

2019-2021 **General Chair or Program Chair**.

- NeurIPS'21 workshop on Privacy in Machine Learning (PriML)
- UCSB CRML Summit 2019 on Responsible Machine Learning

2018-Present **Area chair / Senior Program Commitee**.

- International Conference on Artificial Inteligence and Statistics (AISTATS'19, 20, 21)
- International Conference on Machine Learning (ICML'19, 20, 21)
- International Conference on Neural Information Processing Systems (NeurIPS'21)
- International Conference on Learning Representation (ICLR'21)

2020 **(Co)-organizer**.

- NeurIPS'20 workshop on Privacy-Preserving Machine Learning
- UCSB CRML Summit 2020 on AI and COVID'19
- 2020 LIMPID + IDEAS Joint Workshop on Scalable Image Informatics and the Applications of Machine Learning to Materials Discovery

2014-Present **Program committee for the following workshops**.

- CCS'20 Workshop on Theory and Practice of Differential Privacy (TPDP'20)
- NeurIPS'18 Workshop on Privacy Preserving Machine Learning
- CCS'17 Workshop on Theory and Practice of Differential Privacy (TPDP'17)
- NIPS'15 Workshop on Learning with Incomplete Information and Privacy
- AAAI'15 Workshop on AI for Cities
- CMU Machine Learning Symposium'15.

2013-Present **Reviewer / Program commitee for conferences/journals**.

- International Conference on Machine Learning (ICML)
- Neural Information Processing Systems (NIPS) [Outstanding Reviewer Award 2014]
- Conference on Learning Theory (COLT)
- International Conference on Artificial Inteligence and Statistics (AISTATS)
- IEEE Security and Privacy (S&P 2021)
- ACM-SIAM Symposium on Discrete Algorithms (SODA'21)
- Journal of Machine Learning Research
- Annals of Statistics
- Biometrika
- Annals of the Institute of Statistical Mathematics (AISM)
- Journal of the American Statistical Association (JASA)
- Electronic Journal of Statistics
- Statistica Sinica
- Journal of Computational and Graphical Statistics
- Journal of Privacy and Confidentiality
- Privacy for Statistical Databases (PSD'2016)
- IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)
- IEEE Trans. on Control of Network Systems (TCNS)
- SIAM Journal of Computing (SICOMP)
- Theory of Computing (ToC)
- Operations Research

2019-Present **Referee / Panelist for Funding Agencies**.

- National Science Foundation (2019, 2020, 2021)
- Israel Science Foundation (2020)

2019-Present **Instructors of these UCSB courses**.

- CS 291A Differential Privacy [Fall 2021]
- CS 292F Statistical Foundation of Reinforcement Learning [Spring 2021]
- CS 165A Artificial Intelligence [Winter 2019, Winter 2020, Fall 2020]
- CS 292A Convex Optimization [Spring 2019, Spring 2020]