

Yang Liu

CURRICULUM VITAE — JULY 2023

(734) 546-7375
<https://yliuu.com>
yangliu@ucsc.edu

ACADEMIC POSITIONS	University of California, Santa Cruz Assistant Professor, Computer Science & Engineering Faculty Affiliate, NSF Institute for Foundations of Data Science (IFDS) Program Faculty, UCSC Natural Language Processing Master Program at Silicon Valley	2018 – PRESENT
	RIKEN, Japan Visiting Faculty (remote)	2023 – TO START
	Harvard University Postdoctoral Fellow, Center for Research on Computation and Society	2016 – 2018
RESEARCH INTERESTS	Areas: Human-Centered Machine Learning, Sequential Decision Making, Optimization Topics: Fairness in Machine Learning, Weakly Supervised Learning, Crowdsourcing	
EDUCATION	University of Michigan, Ann Arbor Ph.D. in Electrical Engineering and Computer Science M.Sc. in EECS, Mathematics	2010 – 2015
	Shanghai Jiao Tong University B.Sc. in Information Security	2006 – 2010
GRANTS & AWARDED PROJECTS	NSF CAREER: Human-Centered Machine Learning: Robustness, Fairness and Dynamics (\$526,935)	2022 – 2027
	NSF & Amazon FAI: Fairness in Machine Learning with Human in the Loop, Lead PI (\$1M)	2021 – 2024
	NSF CORE: Wisdom of Crowds with Machines in the Loop, PI (\$233,399)	2020 – 2023
	NSF IFDS: Institute for Foundations of Data Science, Senior Personnel (\$2.23M)	2020 – 2025
	ONR Basic AI Research: Sequential Decision Making with Human Biases, co-PI (\$453,000)	2020 – 2023
	DARPA SCORE: Replication Markets, co-PI (\$177,551)	2019 – 2020
AWARDS	UCSC Office of Research: Special Research Grant, PI (\$4,000)	2019 – 2020
	NSF CAREER Award	2022
	Best Paper Award, ICML 2022 workshop on New Frontiers in Adversarial Machine Learning	2022
	Best Paper Award, AAMAS 2022 workshop on Learning with Strategic Agents	2022
	Best Paper Award, ICML 2021 workshop on Algorithmic Recourse	2021
	Best Paper Award, IJCAI 2021 workshop on Weakly Supervised Representation Learning	2021
	Best Application Paper Award, IEEE International Conference on Data Science and Advanced Analytics	2014
	Towner Price (outstanding Ph.D. research) Finalist, University of Michigan	2014
	Best Poster Award of Engineering Graduate Symposium, University of Michigan	2011
	Dean's fellowship, EECS:Systems, University of Michigan	2010 – 2012
	Li & Fung Scholar, University of Hong Kong	2008
	Tencent Scholarship	2008
	National First Prize in Mathematics Olympic Games	2006

PUBLICATIONS 2023

Google
Scholar

α - β
Alphabetical
 Δ
Students at
UCSC

1. [Tier Balancing: Towards Dynamic Fairness over Underlying Causal Factors](#)
Zeyu Tang, Yatong Chen, **Yang Liu** and Kun Zhang
ICLR – *International Conference on Learning Representations*, 2023
2. [Distributionally Robust Post-hoc Classifiers under Prior Shifts](#)
Jiaheng Wei, Harikrishna Narasimhan, Ehsan Amid, Wen-Sheng Chu, **Yang Liu** and Abhishek Kumar
ICLR – *International Conference on Learning Representations*, 2023
- 2022
3. [Fairness Transferability Subject to Bounded Distribution Shift](#)
Yatong Chen Δ , Reilly Raab Δ , Jialu Wang Δ and **Yang Liu**
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2022
4. [Certifying Some Distributional Fairness with Subpopulation Decomposition](#)
Mintong Kang, Linyi Li, Maurice Weber, **Yang Liu**, Ce Zhang, Bo Li
NeurIPS (spotlight) – *Annual Conference on Neural Information Processing Systems*, 2022
5. [Adaptive Data Debiasing Through Bounded Exploration](#)
Yifan Yang, **Yang Liu** and Parinaz Naghizadeh
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2022
6. [To Smooth or Not? When Label Smoothing Meets Noisy Labels](#)
Jiaheng Wei Δ , Hangyu Liu, Tongliang Liu, Gang Niu, Masashi Sugiyama, and **Yang Liu**
ICML (Long talk) – *International Conference on Machine Learning*, 2022
7. [Understanding Instance-Level Impact of Fairness Constraints](#)
Jialu Wang Δ , Xin Wang and **Yang Liu**
ICML – *International Conference on Machine Learning*, 2022
8. [Metric-Fair Classifier Derandomization](#)
Jimmy Wu, Yatong Chen Δ and **Yang Liu**
ICML – *International Conference on Machine Learning*, 2022
9. [Detecting Corrupted Labels Without Training a Model to Predict](#)
Zhaowei Zhu Δ , Zihao Dong Δ and **Yang Liu**
ICML – *International Conference on Machine Learning*, 2022
10. [Label Noise Transition Matrix Estimation for Tasks with Lower-Quality Features](#)
Zhaowei Zhu Δ , Jialu Wang Δ and **Yang Liu**
ICML – *International Conference on Machine Learning*, 2022
11. [Estimating Instance-dependent Label-noise Transition Matrix using a Deep Neural Network](#)
Shuo Yang, Erkun Yang, Bo Han, **Yang Liu**, Min Xu, Gang Niu, and Tongliang Liu
ICML – *International Conference on Machine Learning*, 2022
12. [Learning with Noisy Labels Revisited: A Study Using Real-World Human Annotations](#)
Jiaheng Wei Δ , Zhaowei Zhu Δ , Hao Cheng Δ , Tongliang Liu, Gang Niu, and **Yang Liu**
ICLR – *International Conference on Learning Representations*, 2022
<http://noisylabels.com>
13. [The Rich Get Richer: Disparate Impact of Semi-Supervised Learning](#)
Zhaowei Zhu Δ , Tianyi Luo Δ and **Yang Liu**
ICLR – *International Conference on Learning Representations*, 2022
14. [Forecasting the publication and citation outcomes of Covid-19 preprints](#)
Michael Gordon, Michael Bishop, Yiling Chen, Anna Dreber, Brandon Goldfedder, Felix Holzmeister, Magnus Johannesson, **Yang Liu**, Louisa Tran, Charles Twardy, Juntao Wang and Thomas Pfeiffer

RSOS – *Royal Society Open Science*, 2022

15. [Machine Truth Serum](#)
Tianyi Luo[△] and Yang Liu
MLJ – *Machine Learning*, 2022
16. [Assessing Multilingual Fairness in Pre-trained Multimodal Representations](#)
Jialu Wang[△], Yang Liu and Xin Eric Wang
ACL – *Findings of Association for Computational Linguistics*, 2022
17. [Fair Classification with Instance-dependent Label Noise](#)
Songhua Wu, Mingming Gong, Bo Han, Yang Liu and Tongliang Liu
CLEaR – *Causal Learning and Reasoning*, 2022
- 2021
18. [Unintended Selection: Persistent Qualification Rate Disparities and Interventions](#)
Reilly Raab[△] and Yang Liu
NeurIPS (Spotlight) – *Annual Conference on Neural Information Processing Systems*, 2021
19. [Can Less be More? When Increasing-to-Balancing Label Noise Rates Considered Beneficial](#)
Yang Liu and Jialu Wang[△]
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2021
20. [Policy Learning Using Weak Supervision](#)
Jingkang Wang*, Hongyi Guo*, Zhaowei Zhu*[△] and Yang Liu
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2021
21. [Bandit Learning with Delayed Impact of Actions](#)
Wei Tang, Chien-Ju Ho and Yang Liu
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2021
22. [Understanding Instance-Level Label Noise: Disparate Impacts and Treatments](#)
Yang Liu
ICML (Long talk) – *International Conference on Machine Learning*, 2021
23. [Clusterability as an Alternative to Anchor Points When Learning with Noisy Labels](#)
Zhaowei Zhu[△], Yiwen Song[△] and Yang Liu
ICML – *International Conference on Machine Learning*, 2021
24. [When Optimizing \$f\$ -divergence is Robust with Label Noise](#)
Jiaheng Wei[△] and Yang Liu
ICLR – *International Conference on Learning Representations*, 2021
25. [Learning with Instance-Dependent Label Noise: A Sample Sieve Approach](#)
Hao Cheng[△], Zhaowei Zhu[△], Xingyu Li[△], Yifei Gong, Xing Sun and Yang Liu
ICLR – *International Conference on Learning Representations*, 2021
26. [Linear Models are Robust Optimal Under Strategic Behavior](#)
Wei Tang, Chien-Ju Ho and Yang Liu
AISTATS – *International Conference on Artificial Intelligence and Statistics*, 2021
27. [Sample Elicitation](#)
Jiaheng Wei[△], Zuyue Fu*, Yang Liu*, Xingyu Li[△], Zhuoran Yang and Zhaoran Wang
AISTATS – *International Conference on Artificial Intelligence and Statistics*, 2021
28. [Federated Bandit: A Gossiping Approach](#)
Zhaowei Zhu[△], Jingxuan Zhu, Ji Liu and Yang Liu
SIGMETRICS – *ACM International Conference on Measurement and Modeling of Computer Systems*, 2021
29. [A Second-Order Approach to Learning with Instance-Dependent Label Noise](#)

- Zhaowei Zhu[△], Tongliang Liu and Yang Liu
CVPR (Oral selection) – *Conference on Computer Vision and Pattern Recognition*, 2021
30. [Fair Classification with Group-Dependent Label Noise](#)
Jialu Wang[△], Yang Liu and Caleb Levy[△]
ACM FAccT – *ACM Conference on Fairness, Accountability, and Transparency*, 2021
31. [Are Gender-Neutral Queries Really Gender-Neutral? Mitigating Gender Bias in Image Search](#)
Jialu Wang[△], Yang Liu and Xin Eric Wang
EMNLP (Oral selection) – *Empirical Method in Natural Language Processing*, 2021
- 2020
32. [Learning Strategy-Aware Linear Classifiers](#)
 α - β Yiling Chen, Yang Liu and Chara Podimata
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2020
33. [Optimal Query Complexity of Secure Stochastic Convex Optimization](#)
Wei Tang, Chien-Ju Ho and Yang Liu
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2020
34. [How Do Fair Decisions Fare in Long-Term Qualification?](#)
Xueru Zhang*, Ruibo Tu*, Yang Liu, Mingyan Liu, Hedvig Kjellström, Kun Zhang and Cheng Zhang
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2020
35. [Peer Loss Functions: Learning from Noisy Labels without Knowing Noise Rates](#)
Yang Liu and Hongyi Guo[△]
ICML (Long talk) – *International Conference on Machine Learning*, 2020
36. [Online Learning Using Only Peer Prediction](#)
Yang Liu and David Helmbold
AISTATS – *International Conference on Artificial Intelligence and Statistics*, 2020
37. [Reinforcement Learning with Perturbed Rewards](#)
Jingkang Wang, Yang Liu and Bo Li
AAAI – *AAAI Conference on Artificial Intelligence*, 2020
38. [How do fairness definitions fare? Testing public attitudes towards three algorithmic defn. of fairness in loan allocations](#)
Nripsuta Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David Parkes and Yang Liu
AIJ – *Artificial Intelligence Journal*, 2020
39. [Surrogate Scoring Rules](#)
Yang Liu*, Juntao Wang* and Yiling Chen
EC – *ACM Conference on Economics and Computation*, 2020
40. [Research Replication Prediction Using Weakly Supervised Learning](#)
Tianyi Luo[△], Xingyu Li[△], Hainan Wang[△] and Yang Liu
EMNLP – *Findings at Empirical Method in Natural Language Processing*, 2020
41. [Are replication rates the same across academic fields? Community forecasts from the DARPA SCORE programme](#)
Michael Gordon*, Domenico Viganola*, Michael Bishop, Yiling Chen, Anna Dreber, Brandon Goldfeder, Felix Holzmeister, Magnus Johannesson, Yang Liu, Charles Twardy, Juntao Wang and Thomas Pfeiffer
RSOS – *Royal Society Open Science*, 2020
- 2019
42. [Fairness without Harm: Decoupled Classifiers with Preference Guarantees](#)
Berk Ustun, Yang Liu and David Parkes
ICML – *International Conference on Machine Learning*, 2019
43. [Actionable Recourse in Linear Classification](#)

- Berk Ustun, Alexander Spangher and **Yang Liu**
ACM FAccT – *ACM Conference on Fairness, Accountability, and Transparency*, 2019
44. [Bayesian Fairness](#)
 α - β Christos Dimitrakakis, **Yang Liu**, David Parkes and Goran Radanovic
AAAI – *AAAI Conference on Artificial Intelligence*, 2019
45. [Randomized Wagering Mechanism](#)
 α - β Yiling Chen, **Yang Liu**, and Juntao Wang
AAAI – *AAAI Conference on Artificial Intelligence*, 2019
- 2018
46. [Inference Aided Reinforcement Learning for Incentive Mechanism Design in Crowdsourcing](#)
Zehong Hu, Yitao Liang, Jie Zhang, Zhao Li and **Yang Liu**
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2018
47. [Incentivizing High Quality User Contributions: New Arm Generation in Bandit Learning](#)
Yang Liu and Chien-Ju Ho
AAAI – *AAAI Conference on Artificial Intelligence*, 2018
48. [Active Information Acquisition for Linear Optimization](#)
Shuran Zheng, Bo Waggoner, **Yang Liu** and Yiling Chen
UAI – *Uncertainty in Artificial Intelligence*, 2018
- 2017
49. [Machine Learning aided Peer Prediction](#)
Yang Liu and Yiling Chen
EC – *ACM Conference on Economics and Computation*, 2017
50. [Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices](#)
Yang Liu and Yiling Chen
AAAI – *AAAI Conference on Artificial Intelligence*, 2017
51. [Crowd Learning: Improving Online Decision Making Using Crowdsourced Data](#)
Yang Liu and Mingyan Liu
IJCAI – *International Joint Conference on Artificial Intelligence*, 2017
52. [Calibrated Fairness in Bandit](#)
Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalya Mandal and David Parkes
FAT-ML – *Fairness, Accountability, and Transparency in Machine Learning*, 2017
53. [Fair Optimal Stopping Policy for Matching with Mediator](#)
Yang Liu
UAI – *Uncertainty in Artificial Intelligence*, 2017
- 2016 AND BEFORE
54. [A Bandit Framework for Strategic Regression](#)
Yang Liu and Yiling Chen
NeurIPS – *Annual Conference on Neural Information Processing Systems*, 2016
55. [Finding One's Best Crowd: Online Prediction By Exploiting Source Similarity](#)
Yang Liu and Mingyan Liu
AAAI – *AAAI Conference on Artificial Intelligence*, 2016
56. [Learning to Incentivize: Eliciting Effort via Output Agreement](#)
Yang Liu and Yiling Chen
IJCAI – *International Joint Conference on Artificial Intelligence*, 2016

57.	An Online Learning Approach to Improving the Quality of Crowd-Sourcing Yang Liu and Mingyan Liu SIGMETRICS – ACM International Conference on Measurement and Modeling of Computer Systems, 2015	
58.	Sufficient Conditions on the Optimality of Myopic Sensing in Opportunistic Channel Access: A Unifying Framework Yang Liu, Mingyan Liu and Sahand Haji Ali Ahmad TIT – IEEE Transaction on Information Theory, 2014	
59.	Detecting Hidden Propagation Structure and Its Application to Analyzing Phishing Yang Liu and Mingyan Liu DSAA – International Conference on Data Science and Advanced Analytics, 2014	
OTHER APPLICATIONS USING MACHINE LEARNING		
60.	From Patching Delays to Infection Symptoms: Using Risk Profiles for an Early Discovery of Vulne. Exploited in the Wild Chaowei Xiao, Armin Sarabi, Yang Liu, Bo Li, Mingyan Liu and Tudor Dumitras USENIX Security – USENIX Security Symposium, 2018	
61.	Cloudy with a Chance of Breach: Forecasting Cyber Security Incidents Yang Liu, Armin Sarabi, Jing Zhang, Parinaz Ardabili, Manish Karir, Michael Bailey and Mingyan Liu USENIX Security – USENIX Security Symposium, 2015	
PATENTS	Rating Network Security Posture and Comparing Network Maliciousness Mingyan Liu, Michael Bailey, Manish Karir, Yang Liu, Jing Zhang, US Patent 14/801,016	
SELECTED MEDIA COVERAGE	WIRED report on our science reproducibility study via building a prediction market WIRED report on our actionable recourse paper (FAccT 19) FICO acquisition of Quadmetrics ¹ WSJ report on our incident prediction work (USENIX Security 15)	2019 2018 2016 2016
ADVISING (CSE PhD)	Yatong Chen, 3rd year, Anti-Racism Research Fellow Hao Cheng, 2nd year, Chancellor's fellowship Zonglin Di, 1st year, Chancellor's fellowship Reilly Raab, 4th year, ARCS Scholar Jialu Wang, 4th year Jiaheng Wei, 4th year, CROSS Fellow Jayanth Yetukuri, 4th year Zhaowei Zhu, 4th year Tianyi Luo, PhD. Joined Amazon as an Applied Scientist	2019 – PRESENT 2021 – PRESENT 2022 – PRESENT 2019 – PRESENT 2019 – PRESENT 2019 – PRESENT 2021 – PRESENT 2019 – PRESENT 2019 – 2022
WORKSHOP COMPETITION TUTORIALS	Tutorial on Machine Learning in Strategic Settings, WINE Workshop on Weakly Supervised Learning, ACML Tutorial on Learning and Mining from Noisy Labels, CIKM Learning with Noisy Labels Competition, IJCAI-ECAI Workshop on Incentives in Machine Learning, ICML Data Science Day (Workshop), UC Santa Cruz Workshop on Game-Theoretic Mechanisms for Data and Information, IJCAI-ECAI/ICML/AAMAS	2022 2022 2022 2022 2020 2019 2018
INVITED TALKS	Fairness in Machine Learning when Agents Respond, University of Washington ECE Colloquium Agency Bias in Machine Learning, University of Southern California ML Seminar Fairness in Machine Learning when Agents Respond, Brandeis University CS Colloquium Learning from Noisy Labels without Knowing Noise Rate, University of Michigan, Ann Arbor Learning from Noisy Labels without Knowing Noise Rate, IDEAL workshop@Northwestern Learning from Noisy Labels without Knowing Noise Rate, UCSD ML Seminar	11/2022 11/2022 11/2022 11/2022 10/2022 10/2022

	The Matthew Effect When Learning from Weak Supervisions, RIKEN	08/2022
	Consequential Machine Learning, ByteDance AI	06/2022
	Learning with Noisy Labels, Google Research	04/2022
	Learning with Noisy Labels, Amazon Science	03/2022
	The Matthew Effect When Learning from Weak Supervisions, Oracle Lab	02/2022
	Revisiting Dynamics in Strategic ML, StraML@NeurIPS	12/2021
	Learning from Weak Supervisions, Marvell AI	11/2021
	Learning from Noisy Labels: Some Lessons and Challenges, Weakly Supervised Learning @IJCAI	08/2021
	Consequential Machine Learning, Institute for Foundations of Data Science (IFDS) Colloquium	03/2021
	Consequential Machine Learning, Washington University in St. Louis CS Colloquium	12/2020
	Fair Classification When Imperfect Training Labels Present, UW, Statistics Seminar	10/2020
	Peer Loss Functions, Google Research, Seattle	07/2020
	Multi-Armed Bandit with Delayed Impact of Actions, UCLA	05/2020
	Peer Loss Functions, Google Brain	01/2020
	Delayed Impacts of Actions in Multi-Armed Bandit, INFORMS Annual Meeting	10/2019
	Fairness in Machine Learning, Theoretical Computer Science Seminar, UC Santa Cruz	04/2019
	Fairness in Machine Learning, Data Science Day, UC Santa Cruz	05/2019
	Verifying the Value of Information without Verification, Rutgers IOE Seminar	10/2018
	Incentivizing Societal Contributions for and via Machine Learning	
	★ Shanghai Jiao Tong University	05/2018
	★ Microsoft Research NYC	04/2018
	★ Purdue, UC Santa Cruz, Vanderbilt, U Pitt	03/2018
	★ UIC, Penn State, U of Arizona	02/2018
	Surrogate Scoring Rules and Dominant Truth Serum, UPenn Theory Seminar	10/2017
	Machine Learning aided Peer Prediction, BayesianCrowd	07/2017
	Machine Learning aided Incentive Design, Dartmouth College CS Colloquium	05/2017
	Learning and Harnessing the Power of Human Computation, USC ISI	02/2017
	A Bandit Framework for Strategic Regression, UIUC CSL seminar	02/2017
	Bandit in Crowdsourcing, Models and Algorithms for Crowds and Networks	09/2016
	Improving Prediction By Exploiting Source Similarity, ITA	02/2015
	Anatomy of Network-Level Malicious Activities, statistics seminar, UMich	05/2014
	Diversity Gains in Multiuser Wireless Network, Qualcomm research	08/2013
TEACHING	CSE 242 Machine Learning, UC Santa Cruz	F2021, F2020
	CSE 142 Machine Learning, UC Santa Cruz	F2020, S2020
	CSE 290 Computing for Society, UC Santa Cruz	S2021, F2019
	CSE 248 Foundations of Data Science, UC Santa Cruz	F2022, S2019
	CSE 140 Artificial Intelligence, UC Santa Cruz	W2019
CONFERENCE PROGRAM COMMITTEE	Neural Information Processing Systems (NeurIPS), Area Chair	2021 - 2022
	ACM Fairness, Accountability and Transparency (FAccT), Area Chair	2022 - 2023
	Causal Learning and Reasoning (CLear), Area Chair	2022
	Uncertainty in Artificial Intelligence (UAI), Area Chair	2022 - 2023
	AAAI Conference on Artificial Intelligence (AAAI), Senior PC	2021 - 2022
	International Joint Conference on Artificial Intelligence (IJCAI), Senior PC	2021 - 2023
	International Conference on Machine Learning (ICML), PC	2020 - 2022
	ACM Economics and Computation (EC), PC	2021 - 2022
	International Conference on Artificial Intelligence and Statistics (AISTATS), PC	2021 - 2022
	International Conference on Autonomous Agents and Multiagent Systems (AAMAS), PC	2019 - 2021
	Neural Information Processing Systems (NeurIPS), PC	2020
	AAAI Conference on Artificial Intelligence (AAAI), PC	2019 - 2020
	International Joint Conference on Artificial Intelligence (IJCAI), PC	2019 - 2020

SELECTED	Journal of Machine Learning Research	
JOURNAL	Journal of Artificial Intelligence Research	
REVIEWER	Management Science	
	ACM Transactions on Economics and Computation	
	IEEE Transaction on Signal Processing	
	IEEE Transaction on Information Forensics & Security	
	IEEE Transaction on Information Theory	
	ACM/IEEE Transaction on Networking	
	Neurocomputing	
	IEEE Transaction on Communications	
	IEEE Journal on Selected Areas in Communications	
	IEEE Transaction on Mobile Computing	
	IEEE Transactions on Cognitive Communications and Networking	
UNIVERSITY	ML Faculty search committee of CSE department	2021, 2022
SERVICES	HCI faculty search committee of Computational Media department	2020
	Faculty search committee of CSE department	2020
	2019-2020 Graduate admission committee of CSE department	2019
	Co-organized Data Science Day at UC Santa Cruz	2019
	Co-organized the Joint AI/ML Poster Presentation Session at CSE department	2019
	Co-organized the Theoretical Computer Science Seminar at CSE department	2019
OTHER	Visiting researcher, Microsoft Research, NYC	04/2018
ACTIVITIES	Research intern, Qualcomm Research, San Diego, California	SUMMER 2013
REFERENCES	Available upon request.	