

AMEET TALWALKAR

Machine Learning Department
Carnegie Mellon University
talwalkar@cmu.edu

Education

- 2006 – 2010 M.S./Ph.D in Computer Science, Courant Institute, New York University
Thesis: *Matrix Approximation for Large-scale Learning* (Advisor: Mehryar Mohri)
- 1998 – 2002 B.S. in Computer Science, summa cum laude, Yale University

Experience

- 2023 – Associate Professor (tenured) at Carnegie Mellon University, Machine Learning Department
- 2022 – Venture Partner, Amplify Partners
- 2022 – 2023 Associate Professor (untenured) at Carnegie Mellon University, Machine Learning Department
- 2021 – 2022 Senior Director, Hewlett Packard Enterprises
- 2018 – 2022 Assistant Professor at Carnegie Mellon University, Machine Learning Department
- 2017 – 2021 Co-founder and Chief Scientist at Determined AI (acquired by Hewlett Packard Enterprises)
- 2016 – 2017 Affiliated Assistant Professor at University of California, Los Angeles, Statistics Department
- 2014 – 2017 Assistant Professor at University of California, Los Angeles, Computer Science Department
- 2014 – 2017 Technical Advisor at Databricks
- 2014 – 2015 Visiting Assistant Professor at University of California Berkeley, EECS Department
- 2010 – 2014 Postdoctoral Scholar at UC Berkeley, EECS Department (Advisor: Michael I. Jordan)
- 2007 – 2009 Intern at Google New York
- 2004 – 2006 Researcher in Paul Greengard's Neuroscience Laboratory at The Rockefeller University
- 2003 – 2004 Software Developer at Wireless Generation
- 2002 – 2003 Consultant at Oliver, Wyman & Company

Awards and Honors

- 2023 Best Paper Award at Conference on Human Computation and Crowdsourcing (HCOMP)
- 2021 NSF CAREER Award
- 2020 Facebook Faculty Research Award
- 2019 JP Morgan Faculty Research Award
- 2019 Carnegie Bosch Institute Research Award
- 2018 Okawa Research Grant
- 2018 Google Faculty Research Award
- 2018 Machine Learning Research Award (MLRA) through Amazon Web Services (AWS)
- 2016 edX Prize Finalist for Exceptional Contributions in Online Teaching and Learning

2016	Microsoft Azure Research Award
2015	Google Faculty Research Award
2015	Bloomberg Research Grant Award
2015	Amazon Web Services Research Award
2011	NSF Office of Cyberinfrastructure (OCI) Postdoctoral Fellowship
2011	Genentech Innovation Postdoctoral Fellowship
2011	Janet Fabri Prize for best doctoral dissertation in NYU’s Computer Science Department
2009	Best Student Paper at New York Academy of Sciences’ Symposium on Machine Learning
2008	Henning Biermann Award for exceptional service to NYU’s Computer Science Department
2002	Yale Computer Science Undergraduate Prize
2001	Member of Yale’s Phi Beta Kappa and Tau Beta Pi Honor Societies
2000	Yale Science and Engineering Association Junior High Scholarship Award
1998	Robert C. Byrd Scholarship
1998	Westinghouse Science Talent Search Finalist

Publications (* alphabetical, † equal contribution)

BOOK AND BOOK CHAPTERS

- M. Mohri*, A. Rostamizadeh*, A. Talwalkar*. *Foundations of Machine Learning, 2nd Edition*. The MIT Press, 2018.
- M. Mohri*, A. Rostamizadeh*, A. Talwalkar*. *Foundations of Machine Learning*. The MIT Press, 2012.
- S. Kumar*, M. Mohri*, A. Talwalkar*. “Ensemble Nyström.” *Ensemble Machine Learning: Methods and Applications*. Editors: C. Zhang and Y. Ma. Springer, 2012.
- A. Talwalkar, S. Kumar, M. Mohri, H. Rowley. “Large-Scale Manifold Learning.” *Manifold Learning Theory and Applications*. Editors: Y. Ma and Y. Fu. CRC Press, 2011.

PEER-REVIEWED JOURNAL PUBLICATIONS

- V. Chent†, U. Bhatt†, H. Heidari, A. Weller, A. Talwalkar. *Perspectives on Incorporating Expert Feedback into Model Updates*. Patterns, 2023.
- A. Chiu, E. Molloy, Z. Tan, A. Talwalkar, S. Sankararaman. *Inferring Population Structure in Biobank-scale Genomic Data*. The American Journal of Human Genetics, 2022.
- V. Chent†, J. Lit†, J. Kim, G. Plumb, A. Talwalkar. *Interpretable Machine Learning: Moving From Mythos to Diagnostics*. Communications of the ACM (CACM), 2022.
- T. Li, A. Sahu, A. Talwalkar, V. Smith. *Federated Learning: Challenges, Methods, and Future Directions*. IEEE Signal Processing Magazine, Special Issue on Distributed, Streaming Machine Learning, 2020.
- L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar. *Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization*. Journal of Machine Learning Research (JMLR), 2018.
- X. Meng, J. Bradley, B. Yuvaz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M. Franklin, R. Zadeh, M. Zaharia, A. Talwalkar. *MLlib: Machine Learning in Apache Spark*. Journal of Machine Learning Research (JMLR), 2016.
- L. Mackey†, A. Talwalkar†, M.I. Jordan. *Distributed Matrix Completion and Robust Factorization*. Journal of Machine Learning Research (JMLR), 2015.

- A. Talwalkar†, J. Liptrap†, J. Newcomb, C. Hartl, J. Terhorst, K. Curtis, M. Bresler, Y. S. Song, M. I. Jordan, D. Patterson. *SMaSH: A Benchmarking Toolkit for Variant Calling*. Bioinformatics, 2014.
- N. Gong, A. Talwalkar, L. Mackey, L. Huang, R. Shin, E. Stefanov, E. Shi, D. Song. *Joint Link Prediction and Attribute Inference using a Social-Attribute Network*. ACM Transactions on Intelligent Systems and Technology (TIST), 2014.
- A. Talwalkar, S. Kumar, M. Mohri, H. Rowley. *Large-scale SVD and Manifold Learning*. Journal of Machine Learning Research (JMLR), 2013.
- A. Kleiner, A. Talwalkar, P. Sarkar, M.I. Jordan, *A Scalable Bootstrap for Massive Data*. Journal of the Royal Statistical Society, Series B (JRSS-B), 2013.
- S. Kumar*, M. Mohri*, A. Talwalkar*. *Sampling methods for the Nystrom Method*. Journal of Machine Learning Research (JMLR), 2012.
- P. Radivojac, W. Clark, T. Oron, A. Schnoes, T. Wittkop, A. Sokolov, K. Graim, C. Funk, K. Verspoor, A. Ben-Hur, G. Pandey, J. Yunes, A. Talwalkar, et. al. *A Large-scale Evaluation of Computational Protein Function Prediction*. Nature Methods, 2012.

PEER-REVIEWED CONFERENCE PUBLICATIONS

- N. Johnson, A. Cabrera, G. Plumb, A. Talwalkar. *Where Does My Model Underperform? A Human Evaluation of Slice Discovery Algorithms*. Conference on Human Computation and Crowdsourcing (HCOMP), 2023.
- M. Barker, E. Kallina, D. Ashok, K. Collins, A. Casovan, A. Weller, A. Talwalkar, V. Chen, U. Bhatt. *FeedbackLogs: Recording and Incorporating Stakeholder Feedback into Machine Learning Pipelines*. Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO), 2023.
- J. Kim, V. Chen, D. Pruthi, N. Shah, A. Talwalkar. *Assisting Human Decisions in Document Matching*. Transactions on Machine Learning Research (TMLR), 2023.
- G. Plumb, N. Johnson, A. Cabrera, A. Talwalkar. *Towards a More Rigorous Science of Blindspot Discovery in Image Models*. Transactions on Machine Learning Research (TMLR), 2023.
- J. Shen, L. Li, L. Dery, C. Staten, M. Khodak, G. Neubig, A. Talwalkar. *Cross-Modal Fine-Tuning: Align then Refine*. International Conference on Machine Learning (ICML), 2023.
- A. Cabrera, E. Fu, D. Bertucci, K. Hostein, A. Talwalkar, J. Hong, A. Perer. *Zeno: An Interactive Framework for Behavioral Evaluation of Machine Learning*. Conference on Human Factors in Computing Systems (CHI), 2023.
- K. Kuo, P. Thaker, M. Khodak, J. Ngyuen, D. Jiang, A. Talwalkar, V. Smith. *On Noisy Evaluation in Federated Hyperparameter Tuning*. Conference on Machine Learning and Systems (MLSys), 2023.
- L. Dery, P. Michel, M. Khodak, G. Neubig, A. Talwalkar. *AANG: Automating Auxiliary Learning*. International Conference on Learning Representations (ICLR), 2023.
- R. Tu†, N. Roberts†, M. Khodak, J. Shen, F. Sala, A. Talwalkar. *NAS-Bench-360: Benchmarking Diverse Tasks for Neural Architecture Search*. Neural Information Processing Systems (NeurIPS) Datasets and Benchmarking Track, 2022.
- V. Chen, N. Johnson, N. Topin, G. Plumb, A. Talwalkar. *Use-Case-Grounded Simulations for Explanation Evaluation*. Neural Information Processing Systems (NeurIPS), 2022.
- J. Shen†, M. Khodak†, A. Talwalkar. *Efficient Architecture Search for Diverse Tasks*. Neural Information Processing Systems (NeurIPS), 2022.
- M. Khodak*, M. F. Balcan*, A. Talwalkar*, S. Vassilvitskii*. *Learning Predictions for Algorithms with Predictions*. Neural Information Processing Systems (NeurIPS), 2022.
- K. Harris, V. Chen, J. Kim, A. Talwalkar, H. Heidari, Z. Wu. *Bayesian Persuasion for Algorithmic Recourse*. Neural Information Processing Systems (NeurIPS), 2022.

- M. F. Balcan, M. Khodak, D. Sharma, A. Talwalkar. *Provably Tuning the ElasticNet Across Instances*. Neural Information Processing Systems (NeurIPS), 2022.
- G. Plumb, M. Ribeiro, A. Talwalkar. *Finding and Fixing Spurious Patterns with Explanations*. Transactions on Machine Learning Research (TMLR), 2022.
- J. Kim, G. Plumb, A. Talwalkar. *Sanity Simulations for Saliency Methods*. International Conference on Machine Learning (ICML), 2022.
- L. Dery, P. Michel, A. Talwalkar, G. Neubig. *Should We Be Pre-training? An Argument for End-task Aware Training as an Alternative*. International Conference on Learning Representations (ICLR), 2022.
- M. Khodak, R. Tu, T. Li, L. Li, M.F. Balcan, V. Smith, A. Talwalkar. *Federated Hyperparameter Tuning Challenges, Baselines, and Connections to Weight-Sharing*. Neural Information Processing Systems (NeurIPS), 2021.
- N. Roberts†, M. Khodak†, T. Dao, L. Li, Christopher Ré, A. Talwalkar. *Rethinking Neural Operations for Diverse Tasks*. Neural Information Processing Systems (NeurIPS), 2021.
- M.F. Balcan*, M. Khodak*, D. Sharma*, A. Talwalkar*. *Learning-to-learn non-convex piecewise-Lipschitz functions*. Neural Information Processing Systems (NeurIPS), 2021.
- M. Al-Shedivat, L. Li, E. Xing, A. Talwalkar. *On Data Efficiency of Meta-learning*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2021.
- J. Cohen, S. Kaur, Y. Li, Z. Kolter, A. Talwalkar. *Gradient Descent on Neural Networks Typically Occurs at the Edge of Stability*. International Conference on Learning Representations (ICLR), 2021.
- J. Li†, V. Nagarajan†, G. Plumb, A. Talwalkar. *A Learning Theoretic Perspective on Local Explainability*. International Conference on Learning Representations (ICLR), 2021.
- L. Li†, M. Khodak†, M.F. Balcan, A. Talwalkar. *Geometry-Aware Gradient Algorithms for Neural Architecture Search*. International Conference on Learning Representations (ICLR), 2021.
- G. Plumb, M. Al-Shedivat, A. Cabrera, A. Perer, E. Xing, A. Talwalkar. *Regularizing Black-box Models for Improved Interpretability*. Neural Information Processing Systems (NeurIPS), 2020.
- J. Kim, J. Chen, A. Talwalkar. *FACT: A Diagnostic for Group Fairness Trade-offs*. International Conference on Machine Learning (ICML), 2020.
- G. Plumb, J. Terhorst, S. Sankararaman, A. Talwalkar. *Explaining Groups of Points in Low-Dimensional Representations*. International Conference on Machine Learning (ICML), 2020.
- Z. Tan, S. Yeom, M. Fredrikson, A. Talwalkar. *Learning Fair Representations for Kernel Models*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2020.
- T. Li†, A. Sahu†, M. Zaheer, M. Sanjabi, A. Talwalkar, V. Smith. *Federated Optimization for Heterogeneous Networks*. Conference on Machine Learning and Systems (MLSys), 2020.
- L. Li, K. Jamieson, A. Rostamizadeh, E. Gonina, J. Ben-tzur, M. Hardt, B. Recht, A. Talwalkar. *A System for Massively Parallel Hyperparameter Tuning*. Conference on Machine Learning and Systems (MLSys), 2020.
- J. Li, M. Khodak, S. Caldas, A. Talwalkar. *Differentially Private Meta-Learning*. International Conference on Learning Representations (ICLR), 2020.
- T. Li, A. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, V. Smith. *FedDANE: A Federated Newton-Type Method*. Asilomar Conference on Signals, Systems and Computers (Invited Paper), 2019.
- M. Khodak, M.F. Balcan, A. Talwalkar. *Adaptive Gradient-Based Meta-Learning Methods*. Neural Information Processing Systems (NeurIPS), 2019.
- L. Li, A. Talwalkar. *Random Search and Reproducibility for Neural Architecture Search*. Conference on Uncertainty in Artificial Intelligence (UAI), 2019.
- M. Khodak, M. Balcan, A. Talwalkar. *Provable Guarantees for Gradient-Based Meta-Learning*. International Conference on Machine Learning (ICML), 2019.

- G. Plumb, D. Molitor, A. Talwalkar. *Supervised Local Modeling for Interpretability*. Neural Information Processing Systems (NeurIPS), 2018.
- V. Smith, C. Chiang, M. Sanjabi, A. Talwalkar. *Federated Multi-task Learning*. Neural Information Processing Systems (NIPS), 2017.
- A. Amini, S. Kazemitabar, A. Bloniarz, A. Talwalkar. *Variable Importance Using Decision Trees*. Neural Information Processing Systems (NIPS), 2017.
- A. Das, I. Upadhyaya, X. Meng, A. Talwalkar. *Collaborative Filtering as a Case-Study for Model Parallelism on Bulk Synchronous Systems*. International Conference on Information and Knowledge Management (CIKM), 2017.
- L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, A. Talwalkar. *Hyperband: Bandit-Based Configuration Evaluation for Hyperparameter Optimization*. International Conference on Learning Representations (ICLR), 2017.
- H. Qi, E. Sparks, A. Talwalkar. *Paleo: A performance model for deep neural networks*. International Conference on Learning Representations (ICLR), 2017.
- F. Abuzaid, J. Bradley, F. Liang, A. Feng, L. Yang, M. Zaharia, A. Talwalkar. *Yggdrasil: An Optimized System for Training Deep Decision Trees at Scale*. Neural Information Processing Systems (NIPS), 2016.
- A. Bloniarz, C. Wu, B. Yu, A. Talwalkar. *Supervised neighborhoods for distributed nonparametric regression*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2016.
- K. Jamieson, A. Talwalkar. *Non-stochastic Best Arm Identification and Hyperparameter Optimization*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2016.
- E. Sparks, A. Talwalkar, D. Haas, M. Franklin, M. I. Jordan, T. Kraska. *Automating Model Search for Large Scale Machine Learning*. Symposium on Cloud Computing (SOCC), 2015.
- S. Agarwal, H. Milner, A. Kleiner, A. Talwalkar, B. Mozafari, M. I. Jordan, S. Madden, and I. Stoica. *Knowing When You're Wrong: Building Fast and Reliable Approximate Query Processing Systems*. Special Interest Group on Management of Data (SIGMOD), 2014.
- A. Bloniarz†, A. Talwalkar†, J. Terhorst†, M. I. Jordan, D. Patterson, B. Yu, Y. S. Song. *Changepoint Analysis for Efficient Variant Calling*. International Conference on Research in Computational Molecular Biology (RECOMB), 2014.
- E. Sparks, A. Talwalkar, V. Smith, J. Kottalam, X. Pan, J. Gonzalez, M. Franklin, M. I. Jordan, T. Kraska. *MLI: An API for Distributed Machine Learning*. International Conference on Data Mining (ICDM), 2013.
- A. Talwalkar†, L. Mackey†, Y. Mu, S. Chang, M. I. Jordan. *Distributed Low-rank Subspace Segmentation*. International Conference on Computer Vision (ICCV), 2013.
- A. Kleiner, A. Talwalkar, S. Agarwal, I. Stoica, M.I. Jordan. *A General Bootstrap Performance Diagnostic*. Conference on Knowledge, Discovery and Data Mining (KDD oral), 2013.
- T. Kraska†, A. Talwalkar†, J. Duchi, R. Griffith, M. Franklin, M.I. Jordan. *MLbase: A Distributed Machine Learning System*. Conference on Innovative Data Systems Research (CIDR), 2013.
- A. Kleiner, A. Talwalkar, P. Sarkar, M.I. Jordan, *The Big Data Bootstrap*. International Conference on Machine Learning (ICML oral), 2012.
- L. Mackey†, A. Talwalkar†, M.I. Jordan. *Divide-and-Conquer Matrix Factorization*. Neural Information Processing Systems (NIPS), 2011.
- M. Mohri*, A. Talwalkar*. *Can Matrix Coherence be Efficiently and Accurately Estimated?* International Conference on Artificial Intelligence and Statistics (AISTATS oral), 2011.
- A. Talwalkar, A. Rostamizadeh. *Matrix Coherence and the Nystrom Method*. Conference on Uncertainty in Artificial Intelligence (UAI oral), 2010.
- C. Cortes*, M. Mohri*, A. Talwalkar*. *On the Impact of Kernel Approximation on Learning Accuracy*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2010.

- S. Kumar*, M. Mohri*, A. Talwalkar*. *The Ensemble Nyström Method*. Neural Information Processing Systems (NIPS), 2010.
- S. Kumar*, M. Mohri*, A. Talwalkar*. *On Sampling-based Approximate Spectral Decomposition*. International Conference on Machine Learning (ICML), 2009.
- S. Kumar*, M. Mohri*, A. Talwalkar*. *Sampling Techniques for the Nyström Method*. International Conference on Artificial Intelligence and Statistics (AISTATS), 2009.
- C. Allauzen*, M. Mohri*, A. Talwalkar*. *Sequence Kernels for Predicting Protein Essentiality*. International Conference on Machine Learning (ICML), 2008.
- A. Talwalkar, S. Kumar, H. Rowley. *Large-Scale Manifold Learning*. International Conference on Vision and Pattern Recognition (CVPR oral), 2008.

PEER-REVIEWED WORKSHOP PUBLICATIONS

- S. Caldas, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith, A. Talwalkar. *LEAF: A Benchmark for Federated Settings*. Workshop on Federated Learning for Data Privacy and Confidentiality at NeurIPS, 2019.
- S. Caldas, J. Konečný, H. B. McMahan, A. Talwalkar. *Expanding the Reach of Federated Learning by Reducing Client Resource Requirements*. Workshop on Federated Learning for Data Privacy and Confidentiality at NeurIPS, 2019.
- N. Guha, A. Talwalkar, V. Smith. *One-shot Federated Learning*. Workshop on Machine Learning on Devices at NeurIPS, 2018.
- L. Li, E. Sparks, K. Jamieson, A. Talwalkar. *Exploiting Reuse in Pipeline-Aware Hyperparameter Tuning*. Workshop on Systems for Machine Learning at NeurIPS, 2018.
- N. Gong, A. Talwalkar, L. Mackey, L. Huang, E. Shin, E. Stefanov, E. Shi, D. Song. *Jointly Predicting Links and Inferring Attributes using a Social-Attribute Network*. ACM Workshop on Social Network Mining and Analysis (SNA-KDD), 2012.
- L. Boucher, R. Sekuler, A. Talwalkar, A. B. Sekuler. *Motion Perception is Influenced by Sound: Two- and Three-Dimensional Motion*. Association for Research in Vision and Ophthalmology, 1998.

TECHNICAL REPORTS

- V Chent†, M. Yang†, W. Cui, J. Kim, A. Talwalkar, J. Ma. *Best Practices for Interpretable Machine Learning in Computational Biology*. 2022.
- K. Amarasinghe, K. Rodolfa, S. Jesus, V. Chen, V. Balayan, P. Saleiro, P. Bizarro, A. Talwalkar, R. Ghani. *On the Importance of Application-Grounded Experimental Design for Evaluating Explainable ML Methods*. 2022.
- E. Jääsaari, M. Ma, A. Talwalkar, T. Chen. *SONAR: Joint Architecture and System Optimization Search*. 2022.
- J. Wang, Z. Charles, Z. Xu, G. Joshi, H. B. McMahan, et al. *A Field Guide to Federated Optimization*. 2021.
- A. Ratner, ... 60+ authors ..., A. Talwalkar. *MLSys: The New Frontier of Machine Learning Systems*. 2019.
- S. Kazemitabar, A. Amini, A. Talwalkar. *On the support recovery of marginal regression*. 2018.
- P. Chaudhari, C. Baldassi, R. Zecchina, S. Soatto, A. Talwalkar, A. Oberman. *Parle: parallelizing stochastic gradient descent*. 2017.
- K. Curtis, A. Talwalkar, M. Zaharia, A. Fox, D. Patterson. *Hybrid Genomic Processing Via Similar Regions*. 2015.

POPULAR MEDIA

- E. Sparks, A. Talwalkar. *AI Can Do Great Things — If it Doesn't Burn the Planet*. Eye at The Six Five Summit. 2022.
- A. Talwalkar. *Databrew by Databricks Podcast*. 2021.
- A. Talwalkar, C. Smith. *Determined AI Podcast Series*. 2020.

A. Talwalkar. *Democratizing Machine Learning*. The Data Exchange Podcast. 2020
A. Talwalkar. *AI in the 2020s Must Get Greener—and Here's How*. IEEE Spectrum. 2020.
E. Sparks, A. Talwalkar. *AutoML with Determined AI*. Eye on A.I. Podcast with Craig Smith. 2019.
L. Li, A. Talwalkar. *What is Neural Architecture Search?* O'Reilly Ideas AI Blog. 2018.
A. Talwalkar. *Toward the Jet Age of Machine Learning*. O'Reilly Ideas AI Blog. 2018.
A. Talwalkar. *How to Train and Deploy Deep Learning at Scale*. O'Reilly Data Show Podcast. 2018.

Teaching

Fall 2023, MLD 10-605/805 Machine Learning with Large Datasets, with Geoffrey Gordon
Fall 2022, MLD 10-605/805 Machine Learning with Large Datasets, with Henry Chai
Spring 2021, MLD 10-405/605 Machine Learning with Large Datasets, with Virginia Smith
Fall 2020, MLD 10-605/805 Machine Learning with Large Datasets, with Virginia Smith
Fall 2019, MLD 10-701 Introduction to Machine Learning (PhD), with Ziv Bar-Joseph
Spring 2019, MLD 10-718 Data Analysis Course
Fall 2018, MLD 10-718 Data Analysis Course, with Aarti Singh
Spring 2018, MLD 10-821 DAP Prep Course
Spring 2017, UCLA: CS269 Scalable and User-Friendly Machine Learning
Winter 2017, UCLA: CS260 Machine Learning Algorithms
Summer 2016, EdX: Data Science and Engineering with Spark XSeries
Spring 2016, UCLA: CS269 Scalable Machine Learning
Fall 2015, UCLA: CS260 Machine Learning Algorithms
Summer 2015, edX: CS190.1x Scalable Machine Learning

Advising

PHD STUDENTS

Junhong Shen, Summer 2021 – present
Valerie Chen, Summer 2020 – present
Lucio Dery (joint with Graham Neubig), Fall 2019 – present
Jeremy Cohen (joint with Zico Kolter), Fall 2019 – present
Joon Sik Kim, Fall 2019 – Spring 2023 (graduated)
Misha Khodak (joint with Nina Balcan), Fall 2018 – present
Gregory Plumb, Fall 2017 – Spring 2023 (graduated)
Liam Li, Fall 2015 – Spring 2020 (graduated)

POSTDOCTORAL ASSOCIATES

Eric Tan, Postdoctoral Associate, Jan 2018 - Oct 2018
Maziar Sanjabi, Postdoctoral Associate, Mar 2017 - Oct 2017

UNDERGRADUATE AND MS STUDENTS

Ritvik Gupta, Fall 2023 - present
Xuchen Gong, Fall 2023 - present
Zongzhe Xu, Fall 2023 - present
Arya Shah, Summer 2023
Samuel Guo, Spring 2022 – Summer 2023
Nathan Moss, Fall 2021 – Summer 2022
Wenbo Cui, Spring 2021 – Fall 2021
Renbo Tu, Fall 2020 – Fall 2021
Chloe Peng, Fall 2020
Haoping Bai, Spring 2020
Nicholas Roberts, Spring 2020 – Summer 2021
Ye Yuan, Fall 2019
Karthik Duddu, Spring 2019
Jeffrey Li, Fall 2018 – Summer 2020
Peter Wu, Summer 2018 – Fall 2018
Nolan Donaghue (UCLA), Fall 2016 – Fall 2019
Christopher Wu (UCLA), Fall 2015 – Spring 2017
Brooke Wenig (UCLA), Fall 2016 – Spring 2017
David Nola (UCLA), Fall 2016 – Spring 2017

PHD THESIS COMMITTEES

Alex Cabrera, Spring 2024 (expected)
Tian Li, Summer 2023
Shilpa George, Spring 2023
Chih-Kuan Yeh, Summer 2022
Nicholay Topin, Spring 2022
Vaishnavh Nagarajan, July 2021
Ellen Viterchik, July 2021
Maruan Al-Shedivat, May 2021
Alnur Ali, May 2019
Aaron Harlap, April 2019
Pratik Chaudhari (UCLA), June 2018
Nikolaos Karianakis (UCLA), May 2017
Konstantine Tsotsos (UCLA), June 2016

Invited Talks

Machine Learning for Diverse Tasks

- Microsoft Research New England (virtual), Apr. 2023.
- AI-EDGE Seminar Series (virtual), Feb. 2023.
- Aurora Innovation (virtual), Feb. 2023.
- UCLA ML Seminar Series, Los Angeles, CA, Mar. 2022.
- LinkedIn Invited TechTalk (virtual), Feb. 2022.
- ELLIS Automated Machine Learning Seminar (virtual), Nov. 2021.
- Stanford MLSys Seminar (virtual), Apr. 2021.

Federated Hyperparameter Tuning: Challenges, Baselines, and Weight Sharing

- WSDM Federated Learning Workshop (virtual), Feb. 2022.
- Amazon Alexa AI (virtual), Sep. 2021.
- Facebook AI (virtual), Sep. 2021.
- ICML Workshop on Federated Learning for User Privacy and Data Confidentiality (virtual), July 2021.

Personalized Federated Learning

- CoNEXT Workshop on DistributedML (virtual), Dec. 2020.
- NeurIPS Workshop on Federated Learning for Data Privacy, Vancouver, Canada, Dec. 2019.

Algorithmic Foundations of Neural Architecture Search

- Microsoft Research New England (virtual), July 2020.
- IQVIA (virtual), July 2020.
- Goldman Sachs (virtual), July 2020.
- Berkeley RiseLab (virtual), Sept. 2020.

Toward the Jet Age of Machine Learning

- UT Austin Wireless Networking and Communications Group, Austin, TX, Feb. 2020.
- Carnegie Mellon and Code & Supply Confluence Talk Series, Pittsburgh, PA, Feb. 2020.

Random Search and Reproducibility for Neural Architecture Search

- Argo AI, Pittsburgh, PA, Oct. 2019.
- Intel, Santa Clara, CA, Sept. 2019.
- Google, Pittsburgh, PA, Apr. 2019.
- O'Reilly AI Conference, New York City, NY, Apr. 2019.

Overcoming Deep Learning's Massive Barriers to Entry

- Morgan Stanley Disruptor Series, New York City, NY, Apr. 2019.

Massively Parallel Hyperparameter Optimization

- Deep Learning NYC Meetup, New York City, NY, Feb. 2019.
- IPAM New Architectures and Algorithms Workshop, Los Angeles, CA, Nov. 2018.

Scalable Deep Learning

- University of Minnesota, Minneapolis, MN, July 2018.
- O'Reilly AI Conference, New York City, NY, May 2018.
- Not Another Big Data Conference, Palo Alto, CA, Dec. 2017.
- UCSC, Santa Cruz, CA, Nov. 2017.
- O'Reilly AI Conference, San Francisco, CA, Sep. 2017.
- MIT, Boston, MA, Sep. 2017.

Scalable Machine Learning Pipelines

- CMU Machine Learning Department Seminar, Pittsburgh, PA, May 2017.
- Yale University CS Department Seminar, New Haven, CT, Mar. 2017.
- Brown University CS Department Seminar, Providence, RI, Mar. 2017.
- Netflix, Los Gatos, CA, Mar. 2017.

Paleo: A Performance Model for Deep Neural Networks

- Information Theory and Applications Workshop, San Diego, CA, Feb. 2017.
- NIPS Workshop on Machine Learning Systems, Barcelona, Spain, Dec. 2016.

A Novel Bandit-Based Approach to Hyperparameter Optimization

- NIPS Workshop on Optimizing the Optimizers, Barcelona, Spain, Dec. 2016.
- Google Research, New York City, NY, Sep. 2016.
- UCLA Statistics Seminar Series, Los Angeles, CA, Feb. 2016.
- Information Theory and Applications Workshop, San Diego, CA, Feb. 2016.

Scalable and Accurate Local Learning via Random Forests

- Joint Statistical Meetings, Chicago, IL, Aug. 2016.
- Bloomberg R&D Seminar, New York City, NY, June 2016.

Machine Learning in Apache Spark

- UT Austin Wireless Networking and Communications Group, Austin, TX, May 2015.
- Columbia Machine Learning Seminar, New York City, NY, Apr. 2015.
- USC CS Colloquium, Los Angeles, CA, Mar. 2015.
- Bloomberg R&D Seminar, New York City, NY, Mar. 2015.
- Information Theory and Applications Workshop, San Diego, CA, Feb. 2015.

MLlib: Apache Spark's Machine Learning Library

- AMPCamp5, Berkeley, CA, Nov. 2014.
- Spark User Meetup, Sunnyvale, CA, July 2014.
- Spark Summit 2014, San Francisco, CA, June 2014.
- NIPS Workshop on Machine Learning Open Source Software, Lake Tahoe, CA, Dec. 2013.

Machine Learning in the Wild

- Facebook AI Lab, Menlo Park, CA, Apr. 2014.
- UCSD CS Seminar, San Diego, CA, Apr. 2014.
- Wisconsin ECE Seminar, Madison, WI, Apr. 2014.
- Princeton CS Seminar, Princeton, NJ, Mar. 2014.
- UCLA CS Seminar, Los Angeles, CA, Mar. 2014.
- UCSD ECE Seminar, San Diego, CA, Feb. 2014.
- Harvard Biostatistics Seminar, Boston, MA, Feb. 2014.
- UCSC CS Seminar, Santa Cruz, CA, Feb. 2014.

Addressing Speed and Accuracy in Variant Calling

- Illumina, San Diego, CA, Jan. 2014.

MLbase: A User-friendly System for Distributed Machine Learning

- NIPS Workshop on Distributed ML and Matrix Computations, Montreal, Canada, Dec. 2014.
- The Machine Learning Conference (MLconf), San Francisco, CA, Nov. 2014.
- Stanford Spark Class, Palo Alto, CA, Aug. 2014.
- DIMACS Workshop on Systems and Analytics of Big Data, Piscataway, NJ, Mar. 2014.
- Information Theory and Applications Workshop, San Diego, CA, Feb. 2014.
- Strata Conference, Santa Clara, CA, Feb. 2014.
- Silicon Valley Machine Learning Meetup, Mountain View, CA, Jan. 2014.
- AMPCamp3, Berkeley, CA, Aug. 2013.
- Spark User Meetup, San Francisco, CA, Aug. 2013.
- SAP, Palo Alto, CA, July 2013.
- IBM Research Almaden, San Jose, CA, Apr. 2013.

Changepoint Analysis for Efficient Variant Calling

- UCSC Center for Biomolecular Science and Engineering Seminar, Santa Cruz, CA, Nov. 2013.

MLI: An API for Distributed Machine Learning

- Bay Area Machine Learning Symposium, Menlo Park, CA, Aug. 2013.

SMaSH: A Benchmarking Suite for Variant Calling

- UCSC Center for Biomolecular Science and Engineering Seminar, Santa Cruz, CA, May 2013.

Scalable and User-Friendly Machine Learning

- Netflix, Los Gatos, CA, Sep. 2013.
- Microsoft Silicon Valley (CISL Group), Mountain View, CA, Aug. 2013.
- SF Bay Area Machine Learning Meetup, San Francisco, CA, May 2013.
- AT&T, San Francisco, Apr. 2013.
- Brown University CS Department Seminar, Providence, RI, Mar. 2013.

Divide-and-Conquer Learning for Big Data

- Boston University ECE Seminar, Boston, MA, Feb. 2013.

Divide-and-Conquer Matrix Factorization

- The Machine Learning Conference (MLconf), San Francisco, CA, Nov. 2013.
- Information Theory and Applications Workshop, San Diego, CA, Feb. 2013.
- INRIA, Paris, France, Dec. 2012.
- NIPS Workshop on Sparse and Low-rank Approximation, Sierra Nevada, Spain, Dec. 2011.

Matrix Approximation for Large-scale Learning

- NYU CS Seminar, New York City, NY, May 2010.
- Statistical Artificial Intelligence Lab Seminar, Berkeley, CA, Mar. 2010.
- Google Research, New York City, NY Feb. 2010.

The Ensemble Nyström Method

- New York Academy of Sciences, New York City, NY, Nov. 2009.

Improved Bounds for the Nyström Method

- New York Academy of Sciences, New York City, NY, Oct. 2008.

Sequence Kernels for Predicting Protein Essentiality

- Google Research, New York City, NY, Aug. 2008.

Professional Activities

CONFERENCE / SEMINAR / WORKSHOP ORGANIZATION

- | | |
|-------------|--|
| 2022 | Co-organizer, AutoML Decathlon Competition at Neural Information Processing Systems |
| 2021 | Publicity Chair, International Conference on Machine Learning |
| 2020 | Co-organizer, Workshop on <i>Scalability, Privacy, and Security in Federated Learning (SpicyFL)</i> at Neural Information Processing Systems (virtual) |
| 2019 – 2023 | Board President, Machine Learning and Systems (MLSys) Conference |
| 2019 | Co-organizer, Workshop on <i>Adaptive and Multitask Learning: Algorithms & Systems</i> at International Conference on Machine Learning, Long Beach, California |
| 2018 | General Chair, Machine Learning and Systems (MLSys18) Conference |
| 2017 | Inaugural Program Chair, Machine Learning and Systems (MLSys17) Conference |
| 2015 – 2016 | UCLA Computer Science Weekly Seminar Series |
| 2014 – 2016 | Information Theory and Applications Workshop Session Organizer |
| 2015 | Co-organizer, 3-day course on <i>Distributed Analytics and Machine Learning with Apache Spark</i> at Berkeley Institute for Data Science, UC Berkeley |
| 2014 | Co-organizer, Workshop on <i>Distributed Machine Learning and Matrix Computations</i> at Neural Information Processing Systems, Montreal, Canada |
| 2014 | Organizer, AMPCamp 5, UC Berkeley |

- 2013 Organizer, Systems and Machine Learning (SysML) Seminar, AMPLab, UC Berkeley
- 2011 Co-organizer, Workshop on *Sparse Representation and Low-rank Approximation* at Neural Information Processing Systems, Sierra Nevada, Spain
- 2010 Co-organizer, Workshop on *Low-rank Methods for Large-scale Machine Learning* at Neural Information Processing Systems, Whistler, Canada
- 2009 Co-organizer, Machine Learning Seminar, Courant Institute, NYU

OTHER

- 2018 – CMU Machine Learning Blog Co-Creator / Faculty Advisor
- 2018 – 2019 CMU AI+ Advisor
- 2017 NSF Panelist
- 2010 NSF Panelist
- 2006 – 2008 Courant Institute Graduate Student Representative