

TIAN LI

CURRICULUM VITAE

Carnegie Mellon University
Computer Science Department

tianli@cmu.edu
cs.cmu.edu/~litian

RESEARCH INTERESTS

My research centers around optimization, trustworthy machine learning, and learning in heterogeneous environments, with applications to federated learning.

Specifically, I am interested in designing, analyzing, and evaluating principled distributed learning algorithms, taking into account real-world constraints (e.g., communication and heterogeneity) to address issues related to accuracy, scalability, trustworthiness (fairness, robustness, and privacy), and their interplays.

EDUCATION

Carnegie Mellon University	PA, USA
Ph.D. in Computer Science	2018 - present
M.S. in Computer Science	2018 - 2020
<i>Advisor: Virginia Smith</i>	
Peking University	Beijing, China
B.S. in Computer Science (summa cum laude)	2014 - 2018
B.A. in Economics	2015 - 2018

EXPERIENCES

Carnegie Mellon University	PA, USA
<i>Research Assistant</i>	2018 - present
Google Research	WA, USA
<i>Research Intern</i>	Summer 2022
Peking University	Beijing, China
<i>Undergraduate Researcher</i>	2016 - 2018
Microsoft Research Asia	Beijing, China
<i>Research Intern</i>	Winter 2017
ETH Zurich	Zurich, Switzerland
<i>Visiting Student</i>	Summer 2017

AWARDS & HONORS

Rising Stars in EECS Workshop, Invited Participant (2022)
Rising Stars in Data Science, UChicago (2022)
Oral Presentation (top 5%) at NeurIPS 2022 OPT-ML Workshop (2022)

Rising Stars in Machine Learning, UMD (2021)
Best Paper Award at ICLR Workshop on Security and Safety in ML Systems (2021)
Outstanding Reviewer Award (top 8%), NeurIPS (2021)
Top 10% Reviewers, ICML (2021)
Several awards and scholarships, Peking University (2014 - 2018)
Student Summer Research Fellowship, ETHZ (2017)
Selected into CS Elite Class (Turing Class), Peking University (2016)

PUBLICATIONS (* indicates equal contribution)

Manuscripts

T. Li*, A. Beirami*, M. Sanjabi, and V. Smith. On Tilted Losses in Machine Learning: Theory and Applications.

Y. J. Cho, D. Jhunjhunwala, T. Li, V. Smith, and G. Joshi. To Federated or Not To Federate: Incentivizing Client Participation in Federated Learning.

J. Wang, Z. Charles, Z. Xu, G. Joshi, H. B. McMahan, et al. A Field Guide to Federated Optimization.

Journal Articles

T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated Learning: Challenges, Methods, and Future Directions. In *IEEE Signal Processing Magazine (SPM), Special Issue on Streaming, Distributed Machine Learning, 2020. (Most Popular SPM Article of 2020: Link)*

Conference Publications

T. Li, M. Zaheer, K. Liu, S. Reddi, B. McMahan, and V. Smith. Differentially Private Adaptive Optimization with Delayed Preconditioners. In *International Conference on Learning Representations (ICLR), 2023. (Oral Presentation (top 5%) at NeurIPS 2022 OPT-ML Workshop)*

T. Li, M. Zaheer, S. Reddi, and V. Smith. Private Adaptive Optimization with Side Information. In *International Conference on Machine Learning (ICML), 2022.*

R. Balakrishnan*, T. Li*, T. Zhou*, N. Himayat, V. Smith, and J. Bilmes. Diverse Client Selection for Federated Learning via Submodular Maximization. In *International Conference on Learning Representations (ICLR), 2022.*

M. Khodak, R. Tu, T. Li, L. Li, M-F. Balcan, V. Smith, and A. Talwalkar. Federated Hyperparameter Optimization: Challenges, Baselines, and Connections with Weight-Sharing. In *Neural Information Processing Systems (NeurIPS), 2021.*

T. Li*, A. Beirami*, M. Sanjabi, and V. Smith. Tilted Empirical Risk Minimization. In *International Conference on Learning Representations (ICLR), 2021.*

T. Li, S. Hu, A. Beirami, and V. Smith. Ditto: Fair and Robust Federated Learning Through Personalization. In *International Conference on Machine Learning (ICML), 2021. (Best Paper Award at ICLR 2021 Secure ML Workshop)*

D. Dennis, **T. Li**, and V. Smith. Heterogeneity for the Win: One-Shot Federated Clustering. In *International Conference on Machine Learning (ICML)*, 2021.

L. A. Melgar, D. Dao, S. Gan, N. M. Gürel, N. Hollenstein, J. Jiang, B. Karlas, T. Lemmin, **T. Li**, Y. Li, X. Rao, J. Rausch, C. Renggli, L. Rimanic, M. Weber, S. Zhang, Z. Zhao, K. Schawinski, W. Wu, and C. Zhang. Ease.ML: A Lifecycle Management System for MLDev and MLOps. In *Conference on Innovative Data Systems Research (CIDR)*, 2021.

T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. Federated Optimization in Heterogeneous Networks. In *Conference on Machine Learning and Systems (MLSys)*, 2020.

T. Li, M. Sanjabi, A. Beirami, and V. Smith. Fair Resource Allocation in Federated Learning. In *International Conference on Learning Representations (ICLR)*, 2020.

T. Yu, **T. Li**, Y. Sun, S. Nanda, V. Smith, V. Sekar, and S. Seshan. Learning Context-Aware Policies from Smart Homes via Federated Multitask Learning. In *Conference on Internet of Things Design and Implementation (IoTDI)*, 2020.

T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. FedDANE: A Federated Newton-Type Method. In *Asilomar Conference on Signals, Systems, and Computers (Asilomar)*, 2019. (Invited Paper)

T. Li, J. Zhong, J. Liu, W. Wu, and C. Zhang. Ease.ml: Towards Multi-Tenant Resource Sharing for Machine Learning Workloads. In *Very Large Data Bases Conferences (VLDB)*, 2018.

Workshop Papers

S. Wu, **T. Li**, Z. Charles, Y. Xiao, Z. Liu, Z. Xu, and V. Smith. Motley: Benchmarking Heterogeneity and Personalization in Federated Learning. In *Workshop on Federated Learning: Recent Advances and New Challenges, NeurIPS 2022*.

S. Caldas, S. K. Duddu, P. Wu, **T. Li**, J. Konecny, H. B. McMahan, V. Smith, and A. Talwalkar. LEAF: A Benchmark for Federated Settings. In *Workshop on Federated Learning for Data Privacy and Confidentiality, NeurIPS 2019*.

Z. Wang*, **T. Li***, Y. Shao, and B. Cui. CUTE: Query Knowledge Graphs by Tabular Examples. In *Web and Information Management Conference (WAIM)*, 2018. (Demo)

C. Zhang, W. Wu, and **T. Li**. An Overreaction to the Broken Machine Learning Abstraction: The Ease.ml Vision. In *Human-In-the-Loop Data Analytics Workshop, SIGMOD 2017*.

TALKS

Differentially Private Adaptive Optimization with Delayed Preconditioners
- *NeurIPS OPT-ML Workshop, LA, Dec. 2022*.

Scalable and Trustworthy Learning in Heterogeneous Networks
- *UChicago Rising Stars in Data Science Workshop, IL, Nov. 2022*.
- *USC Symposium on Frontiers of Machine Learning and Artificial Intelligence, CA, Nov. 2022*.

Differential Privacy Meets Adaptive Optimization
- *UCSD HDSI Seminar, virtual, Nov. 2022*

Trustworthy Learning in Heterogeneous Networks

- *Andalusian Research Institute in DaSCI Seminar, virtual, Nov. 2022.*
- *Qualcomm AI Research DistributedML Seminar, virtual, Oct. 2022.*

On Out-Of-Distribution Generalization in Personalized Federated Learning

- *Google Research, WA, Aug. 2022.*

Motley: Benchmarking Heterogeneity and Personalization in Federated Learning

- *Intel-NSF Workshop on Machine Learning for Wireless Systems, virtual, Oct. 2022.*
- *Google Research, WA, July 2022.*

Personalized Federated Learning: Interplays with Competing Constraints and Beyond

- *International Conference on Continuous Optimization, PA, July 2022.*

On Heterogeneity in Federated Settings

- *UMD Rising Stars in Machine Learning Speaker Series, virtual, Nov. 2021.*
- *CMU Catalyst Group meeting, virtual, Apr. 2021.*

Tilted Empirical Risk Minimization

- *Tsinghua University AI TIME forum, virtual, June 2021.*

Fair and Robust Federated Learning Through Personalization

- *Stanford Software Lunch, virtual, Apr. 2022.*
- *TrustML Young Scientists Seminar series, virtual, Feb. 2022.*
- *ICLR Secure ML Workshop, virtual, May 2021.*

Learning in Heterogeneous Networks: Optimization and Fairness

- *CONIX Student Research Seminar, virtual, Aug. 2020.*
- *Federated Learning One World Seminar, virtual, Aug. 2020.*

Federated Optimization in Heterogeneous Networks

- *MLSys Conference, TX, Mar. 2020.*
- *On-device Intelligence Workshop, TX, Mar. 2020.*
- *Carnegie Bosch Institute Research Projects Workshop, PA, Mar. 2019.*

SERVICES

Co-Organizer

MLSys 2023 Workshop on Federated Learning Systems

ACL 2022 Workshop on Federated Learning for Natural Language Processing

Program Committee

International Conference on Very Large Data Bases (VLDB) 2023

Reviewer

Conferences

Regularly Reviewing: International Conference on Machine Learning (ICML) (*Top 10% reviewers in 2021*), Neural Information Processing Systems (NeurIPS) (*Outstanding Reviewer Award 2021*),

International Conference on Learning Representations (ICLR), Transactions on Machine Learning Research (TMLR)

Workshops

AAAI 2022 Federated Learning Workshop, NeurIPS 2021 Federated Learning Workshop, ICML 2021 Federated Learning Workshop, ICML 2021 Workshop on Information-Theoretic Methods for Responsible ML, ICLR 2021 Responsible AI Workshop (Area Chair), NeurIPS 2020 Federated Learning Workshop

Journals

SIAM Journal on Mathematics of Data Science (SIMODS), IEEE Journal on Selected Areas in Communications (JSAC) Series on Machine Learning for Communications and Networks, IEEE Journal on Selected Areas in Communications (JSAC) Special Issue on Distributed Learning over Wireless Edge Networks, IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Nature Communications, Journal of Machine Learning Research (JMLR)

Others

CMU Computer Science Department Mentorship Program mentor, 2022
CMU Computer Science Department Faculty Hiring Committee student member, 2021 - 2022
Reviewer for Grant Proposal 'Robust Federated Learning for IoT Services', CES 23, French National Research Agency, 2021
CMU School of Computer Science Graduate Application Support Program mentor, 2020
CMU Computer Science Department Ph.D. Orientation Committee member, 2019
EuroSys Shadow Program Committee member, 2018

TEACHING

Guest Lecturer at CSE 598: Machine Learning Security, Privacy, and Fairness
Arizona State University, Fall 2022
"Learning in Heterogeneous Networks"

Guest Lecturer at 15-884: Special Topic: Machine Learning Systems
Carnegie Mellon University, Spring 2021
"Federated Learning"

Teaching Assistant for 15-884: Special Topic: Machine Learning Systems
Carnegie Mellon University, Spring 2021
Instructor: Tianqi Chen

Head Teaching Assistant for 10-405 / 10-605: Machine Learning with Large Datasets
Carnegie Mellon University, Spring 2020
Instructors: Virginia Smith and Heather Miller