GANGHUA WANG

Email: wang9019@umn.edu

Website: https://gwang.umn.edu/

Address: 492 Ford Hall, 224 Church Street SE, Minneapolis, MN 55455

Research Interests: Machine learning safety, deep learning theory and applications, machine learning theory, model fairness

EDUCATION

University of Minnesota, Twin Cities *Aug. 2019 - present*
Ph.D. in Statistics *Minneapolis, MN*
Advised by Prof. Jie Ding and co-advised by Prof. Yuhong Yang

Peking University *Sept. 2015 - July 2019*
B.S. in Statistics, Minor in Economics *Beijing, China*

PUBLICATIONS

Published

* indicates equal contributions

[1] **Ganghua Wang***, Xun Xian*, Jayanth Srinivasa, Ashish Kundu, Xuan Bi, Mingyi Hong, and Jie Ding. "Demystifying Poisoning Backdoor Attacks from a Statistical Perspective". In: *Proc. ICLR* (2024). [pdf].

[2] Enmao Diao*, **Ganghua Wang***, Jie Ding, Yuhong Yang, and Vahid Tarokh. "Pruning deep neural networks from a sparsity perspective". In: *Proc. ICLR* (2023). [pdf].

[3] Gen Li, **Ganghua Wang**, and Jie Ding. "Provable Identifiability of Two-Layer ReLU Neural Networks via LASSO Regularization". In: *IEEE Trans. Inf. Theory* 69.9 (2023), pp. 5921–5935. DOI: 10.1109/TIT.2023.3274152.

[4] **Ganghua Wang**, Jie Ding, and Yuhong Yang. "Regression with Set-Valued Categorical Predictors". In: *Statistica Sinica* 33.4 (2023), pp. 2545–2560. DOI: 10.5705/ss.202021.0332.

[5] Xun Xian, **Ganghua Wang**, Jayanth Srinivasa, Ashish Kundu, Xuan Bi, Mingyi Hong, and Jie Ding. "A Unified Framework for Inference-Stage Backdoor Defenses". In: *Proc. NeurIPS* (2023). [pdf].

[6] Xun Xian*, **Ganghua Wang***, Jayanth Srinivasa, Ashish Kundu, Xuan Bi, Mingyi Hong, and Jie Ding. "Understanding backdoor attacks through the adaptability hypothesis". In: *Proc. ICML* (2023). [pdf].

Under review

[7] **Ganghua Wang** and Jie Ding. "Subset Privacy: Draw from an Obfuscated Urn". In: *arXiv preprint* (2024). [pdf].

[8] **Ganghua Wang**, Ali Payani, Myungjin Lee, and Ramana Kompella. "Federated learning with group bias mitigation: beyond the local fairness". In: *arXiv preprint* (2024). [pdf].

[9] Xun Xian, **Ganghua Wang**, Xuan Bi, Jayanth Srinivasa, Ashish Kundu, Mingyi Hong, and Jie Ding. "RAW: A Robust and Agile Plug-and-Play Watermark Framework with Provable Guarantees". In: *Proc. CVPR* (2024).

[10] Wenjing Yang*, **Ganghua Wang***, Jie Ding, and Yuhong Yang. "A Theoretical Understanding of Neural Network Compression from Sparse Linear Approximation". In: *IEEE Trans. Pattern Anal. Mach. Intell., Under Major Revision* (2024). [pdf].

Manuscript

[11] **Ganghua Wang**, Zhiyuan Tang, and Jie Ding. "Classification with set-valued labels". In: *Preparation* (2024).

[12] **Ganghua Wang**, Yuhong Yang, and Jie Ding. "Model Privacy: A Framework to Understand Model Stealing Attack and Defense". In: *Preparation* (2024).

HONORS AND AWARDS

Awarded by University of Minnesota, Twin Cities

Doctoral Dissertation Fellowship Finalist	*2023*
School of Statistics Travel Award	*2023*
Summer Research Fellowship	*2020*
School of Statistics First Year Scholarship	*2019*

Awarded by Peking University

The Academic Excellence Scholarship (3 times)	*2016 - 2018*
Fang Zheng Scholarship	*2017*
Wu Si Scholarship	*2016, 2018*
Freshman Scholarship	*2015*

Awarded by Cisco Systems, Inc.

Cisco Research Graduate Award	*2022*
Cisco Research Fellow	*2022, 2023*

TEACHING

Teaching Assistant, University of Minnesota, Twin Cities

STAT 4102 Theory of Statistics II	*Fall 2020*
STAT 3021 Introduction to Probability and Statistics	*Spring 2020*
STAT 3011 Introduction to Statistical Analysis	*Fall 2019*

SERVICES

Member of

Graduate Student Liaison Committee,
School of Statistics, University of Minnesota *June 2023 - present*

Mentor of K-12 Outreach Program

AEOP High School Apprenticeship K-12 Outreach Program to enhance diversity, equity, and inclusion (sponsored by Army Research Office) *2022*

Reviewer of

ICASSP, AISTATS, AAAI, ICML, Neural Processing Letter

SOFTWARE

Python Packages

1. `SubsetPrivacy`: Implement the subset privacy mechanisms proposed in [7] and statistical inference methods for set-valued observations [4].
2. `ModelPrivacy`: Implement the common and proposed model stealing defense/attack strategies on benchmark datasets [12].
3. `FedGFT`: Implement a bias mitigation federated learning algorithm proposed in [8]. It has been incorporated into Flame, an open-source project for federated learning systems, which is developed by Cisco Systems, Inc.

4. `CBD`: Conformal backdoor defense that detects the backdoor inputs in the inference stage with provable guarantees on false positive rate [5].

5. `RAW`: A robust and agile watermarking technique for AI-generated images [9].

PATENTS

1. G. Wang, M. Lee, A. Payani, R. Kompella, "Group Bias Mitigation in Federated Learning Systems," 07/28/2023, US patent #18/227,535

TECHNICAL STRENGTHS

Computer Languages Python, MATLAB, R, SQL

PRESENTATIONS AND TALKS

Conference on Neural Information Processing Systems, New Orleans, LA	Dec. 2023
School of Statistics, University of Minnesota, Minneapolis, MN	Nov. 2023
Joint Statistical Meeting, Toronto, Canada	Aug. 2023
International Conference on Econometrics and Statistics, Tokyo, Japan	Aug. 2023
International Conference on Machine Learning, Honolulu, HI	July 2023
Joint Conference on Statistics and Data Science, Beijing, China	July 2023
Center of Statistical Research, Southwestern University of Finance and Economics, Chengdu, China	June 2023
Center for Statistical Science, Peking University, Beijing, China	June 2023
International Conference on Learning Representations, Kigali, Rwanda	May 2023
School of Statistics Student Seminar, Minneapolis, MN	Mar. 2023