Wenqi Wei

Tenure-Track Assistant Professor
Computer and Information Sciences department
Address: 113 West 60th street, Room 610H, New York, NY 10023
Email: wenqiwei@fordham.edu
Homepage: https://wenqiwei789.github.io/Homepage/

Education

- **Georgia Institute of Technology** Atlanta, GA
  *PhD in Computer Science (PhD advisor: **Ling Liu**)* *Aug. 2017 to May 2022*
  - **Thesis**: Adversarial Resilient and Privacy Preserving Deep learning
  - **PhD Minor**: Quantitative and Computational Finance with a focus on AI-augmented financial data system
- **Huazhong University of Science and Technology** Wuhan, China
  Bachelor of Engineering in Electronics and Information Engineering *Sept. 2013 to June. 2017*
  Signal processing track, summa cum laude

Experience

- **Fordham University** New York, NY
  Computer and Information Sciences Department
  Tenure-Track Assistant Professor (January 2023 - now)
  - **Current research areas of Interest**: AI and ML algorithm for big data systems and services; security, privacy, and fairness enhanced ML and AI systems; efficient AI; data mining.
- **IBM Almaden Research Center** San Jose, CA
  Applied Intelligence Department
  Research Staff Member (May 2022 - January 2023)
  - **Project**: Trustworthy foundation models for financial services, Vehicle-IOT collaboration for Smart City, OpenShift for AI-driven Ransomware detection on Cloud. One patent under USPTO review.
- **Georgia Institute of Technology** Atlanta, GA
  Distributed Data Intensive Systems Lab *advisor: Ling Liu*
  Graduate Research Assistant (Aug 2017 - May 2022)
- **IBM Almaden Research Center** San Jose, CA
  Applied Intelligence Department
  Research Intern (June 2020 - August 2020, May 2021 - August 2021)
  - **Project**: Accelerating ransomware detection with graph learning (2020), Graph Neural Networks ensemble learning (2021). 2 patents published.
- **IBM Thomas J. Watson Research Center** Yorktown Heights, NY
  Enterprise Solutions Department
  Research Intern (May 2019 - August 2019)
  - **Project**: Graph representation learning for Bitcoin transaction data mining. Paper published on IEEE TETC.
- **Samsung Research America** Mountain View, CA
  Data Intelligence Group, AI center
  Research Intern (May 2018 - August 2018)
  - **Project**: Computation-efficient deep learning with differential privacy.
- **Huazhong University of Science and Technology** Wuhan, China
  Signal Processing and Information Networking in Communication Lab *advisor: Pan Zhou*
  Undergraduate Research Assistant (Sept 2015 - June 2017)
  - **Privacy-Preserving Networking**: Research on game-theoretic mechanism design with differential privacy for large-scale spectrum sharing while protecting user data privacy.

Journal Publications

[15] **Wenqi Wei** and Ling Liu, “Trustworthy Distributed AI Systems: Robustness, Privacy, and Governance”, accepted by ACM Computing Surveys (ACM CSUR), 2024.

[14] **Wenqi Wei**, Ka-Ho Chow, Tiansheng Huang, Sihao Hu, Yanzhao Wu, and Ling Liu, “Demystifying Data Poisoning Attacks in Distributed Learning as a Service”, accepted by IEEE Transactions Services Computing (IEEE TSC), 2024.

[13] Yanzhao Wu, Ka-Ho Chow, **Wenqi Wei**, and Ling Liu. ”Hierarchical Pruning of Deep Ensembles with Focal Diversity.” accepted by ACM Transactions on Intelligent Systems and Technology (ACM TIST), 2024.

[12] **Wenqi Wei**, Ling Liu, Jingya Zhou, Ka-Ho Chow, and Yanzhao Wu. ”Securing Distributed SGD against Gradient Leakage Threats.” IEEE Transactions on Parallel and Distributed Systems (IEEE TPDS), 34, no. 7, (2023): 2040-2054.

[11] Xigang Sun, Jingya Zhou, Ling Liu, and **Wenqi Wei**. ”Explicit time embedding based cascade attention network for information popularity prediction.” Information Processing & Management, 60, no. 3 (2023): 103278.

[10] Huanhuan Xu, Jingya Zhou, **Wenqi Wei**, and Baolei Cheng. ”Multiuser computation offloading for long-term sequential tasks in mobile edge computing environments.” Tsinghua Science and Technology 28, no. 1 (2023): 93-104.

[9] **Wenqi Wei** and Ling Liu. ”Gradient leakage attack resilient deep learning.” IEEE Transactions on Information Forensics and Security (IEEE TIFS), 17 (2022): 303-316.

[8] Jingya Zhou, Ling Liu, **Wenqi Wei**, and Jianxi Fan. ”Network representation learning: from preprocessing, feature extraction to node embedding.” ACM Computing Surveys (ACM CSUR), 55, no. 2 (2022): 1-35.

[7] Mehmet Emre Gursoy, Ling Liu, Ka-Ho Chow, Stacey Truex, and **Wenqi Wei**. ”An adversarial approach to protocol analysis and selection in local differential privacy.” IEEE Transactions on Information Forensics and Security (IEEE TIFS), 17 (2022): 1785-1799.

[6] **Wenqi Wei** and Ling Liu. ”Robust deep learning ensemble against deception.” IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 18, no. 4 (2021): 1513-1527.

[5] Mehmet Emre Gursoy, Acar Tamersoy, Stacey Truex, **Wenqi Wei**, and Ling Liu. ”Secure and utility-aware data collection with condensed local differential privacy.” IEEE Transactions on Dependable and Secure Computing (IEEE TDSC), 18, no. 5 (2021): 2365-2378.

[4] **Wenqi Wei**, Qi Zhang, and Ling Liu. ”Bitcoin transaction forecasting with deep network representation learning.” IEEE Transactions on Emerging Topics in Computing, 9, no. 3 (2021): 1359-1371.

[3] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, Lei Yu, and **Wenqi Wei**. ”Demystifying membership inference attacks in machine learning as a service.” IEEE Transactions on Services Computing (IEEE TSC), 14, no. 6 (2021): 2073-2089.

[2] Yanzhao Wu, Ling Liu, Calton Pu, Wenqi Cao, Semih Sahin, **Wenqi Wei**, and Qi Zhang. ”A comparative measurement study of deep learning as a service framework.” IEEE Transactions on Services Computing (IEEE TSC), 15, no. 1 (2021): 551-566.

[1] Pan* Zhou, **Wenqi Wei***, Kaigui Bian, Dapeng Oliver Wu, Yuchong Hu, and Qian Wang. ”Private and truthful aggregative game for large-scale spectrum sharing.” IEEE Journal on Selected Areas in Communications (IEEE JSAC), 35, no. 2 (2017): 463-477. (* equal contribution)

Conference Publications

[27] Sihao Hu, Tiansheng Huang, Ka-Ho Chow, **Wenqi Wei**, Yanzhao Wu and Ling Liu, ”ZipZap: Efficient Training of Language Models for Ethereum Fraud Detection”, the ACM Web Conference (theWebConf), Singapore, May 2024.

[26] Fatih Ilhan, Ka-Ho Chow, Sihao Hu, Tiansheng Huang, Selim Tekin, **Wenqi Wei**, Yanzhao Wu, Myungjin Lee, Ramana Kompella, Hugo Latapie, Gaowen Liu and Ling Liu, ”Adaptive Deep Neural Network Inference Optimization with EENet”, Winter Conference on Applications of Computer Vision (WACV), 2024.

[25] **Wenqi Wei** and Ling Liu, ”Gradient Coupling Effect of Poisoning Attacks in Federated Learning”, Hawaii International Conference on System Sciences (HICSS), 2024.

[24] **Wenqi Wei**, Mu Qiao, and Divyesh Jadav. ”GNN-Ensemble: Towards Random Decision Graph Neural Networks”, IEEE International Conference on Big Data (Big Data), 2023.

[23] Ka-Ho Chow, Ling Liu, **Wenqi Wei**, Fatih Ilhan, Yanzhao Wu. ”STDLens: Securing Federated Learning Against Model Hijacking Attacks.”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023.

[22] **Wenqi Wei**, Ka-Ho Chow, Fatih Ilhan, Yanzhao Wu, Ling Liu, ”Model Cloaking against Gradient Leakage”, IEEE International Conference on Data Mining (ICDM), 2023.

[21] Yanzhao Wu, Ka-Ho Chow, **Wenqi Wei**, Ling Liu, ”Exploring Model Learning Heterogeneity for Boosting Ensemble Robustness.”, IEEE International Conference on Data Mining (ICDM), 2023.

[20] Xirong Cao, Xiang Li, Divyesh Jadav, Yanzhao Wu, Zhehui Chen, Chen Zeng and **Wenqi Wei**, ”Invisible Watermarking for Audio Generation Diffusion Models”, IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications, 2023. (mentored student)

[19] Hongpeng Jin, **Wenqi Wei**, Xuyu Wang, Wenbin Zhang and Yanzhao Wu, ”Rethinking Learning Rate Tuning in the Era of Large Language Models.”, IEEE International Conference on Cognitive Machine Intelligence, 2023.

[18] Gaolei Li, Yuanyuan Zhao, **Wenqi Wei**, and Yuchen Liu, ”Few-shot Multi-domain Knowledge Rearming for Context-aware Defence against Advanced Persistent Threats,” IEEE International Conference on Smart Applications, Communications and Networking, pp. 1-8. 2023,

[17] **Wenqi Wei**, Mu Qiao, Eric Butler, and Divyesh Jadav. ”Graph Representation Learning based Vulnerable Target Identification in Ransomware Attacks.” IEEE International Conference on Big Data (Big Data), pp. 2423-2430. 2022.

[16] **Wenqi Wei**, Ling Liu, Yanzhao Wu, Gong Su, and Arun Iyengar. ”Gradient-leakage resilient federated learning.” IEEE International Conference on Distributed Computing Systems (ICDCS), pp. 797-807. 2021.

[15] Yanzhao Wu, Ling Liu, Zhongwei Xie, Ka-Ho Chow, and **Wenqi Wei**. ”Boosting ensemble accuracy by revisiting ensemble diversity metrics.” IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 16469-16477. 2021.

[14] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, **Wenqi Wei**, and Ka Ho Chow. ”The tsc-pfed architecture for privacy-preserving fl.” IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, pp. 207-216. 2021.

[13] **Wenqi Wei**, Ling Liu, Margaret Loper, Ka-Ho Chow, Mehmet Emre Gursoy, Stacey Truex, and Yanzhao Wu. ”A framework for evaluating client privacy leakages in federated learning.” European Symposium on Research in Computer Security (ESORICS), pp. 545-566. 2020.

[12] **Wenqi Wei**, Ling Liu, Margaret Loper, Ka-Ho Chow, Emre Gursoy, Stacey Truex, and Yanzhao Wu. "Cross-layer strategic ensemble defense against adversarial examples." International Conference on Computing, Networking and Communications (ICNC), pp. 456-460. 2020.

[11] Ka-Ho Chow, Ling Liu, Mehmet Emre Gursoy, Stacey Truex, **Wenqi Wei**, and Yanzhao Wu. "Understanding object detection through an adversarial lens." European Symposium on Research in Computer Security (ESORICS), pp. 460-481. 2020.

[10] Stacey Truex, Ling Liu, Ka-Ho Chow, Mehmet Emre Gursoy, and **Wenqi Wei**. "LDP-Fed: Federated learning with local differential privacy." ACM International Workshop on Edge Systems, Analytics and Networking, pp. 61-66. 2020. (Best paper)

[9] **Wenqi Wei**, Ling Liu, Margaret Loper, Ka-Ho Chow, Mehmet Emre Gursoy, Stacey Truex, and Yanzhao Wu. "Adversarial Deception in Deep Learning: Analysis and Mitigation." IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, pp. 236-245. 2020.

[8] Ka-Ho Chow, Ling Liu, Margaret Loper, Juhyun Bae, Mehmet Emre Gursoy, Stacey Truex, **Wenqi Wei**, and Yanzhao Wu. "Adversarial objectness gradient attacks in real-time object detection systems." IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, pp. 263-272. 2020.

[7] Yanzhao Wu, Ling Liu, Zhongwei Xie, Juhyun Bae, Ka-Ho Chow, and **Wenqi Wei**. "Promoting high diversity ensemble learning with ensemblebench." IEEE International Conference on Cognitive Machine Intelligence, pp. 208-217. 2020.

[6] Ka-Ho Chow, **Wenqi Wei**, Yanzhao Wu, and Ling Liu. "Denoising and verification cross-layer ensemble against black-box adversarial attacks." IEEE International Conference on Big Data (Big Data), pp. 1282-1291. 2019.

[5] Yanzhao Wu, Ling Liu, Juhyun Bae, Ka-Ho Chow, Arun Iyengar, Calton Pu, **Wenqi Wei**, Lei Yu, and Qi Zhang. "Demystifying learning rate policies for high accuracy training of deep neural networks." IEEE International conference on Big Data (Big Data), pp. 1971-1980. IEEE, 2019.

[4] Stacey Truex, Ling Liu, Mehmet Emre Gursoy, **Wenqi Wei**, and Lei Yu. "Effects of differential privacy and data skewness on membership inference vulnerability." IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, pp. 82-91. IEEE, 2019.

[3] Ling Liu, **Wenqi Wei**, Ka-Ho Chow, Margaret Loper, Emre Gursoy, Stacey Truex, and Yanzhao Wu. "Deep neural network ensembles against deception: Ensemble diversity, accuracy and robustness." IEEE international conference on mobile ad hoc and sensor systems (MASS), pp. 274-282. IEEE, 2019.

[2] Mehmet Emre Gursoy, Ling Liu, Stacey Truex, Lei Yu, and **Wenqi Wei**. "Utility-aware synthesis of differentially private and attack-resilient location traces." ACM SIGSAC Conference on Computer and Communications Security (CCS), pp. 196-211. 2018.

[1] Ling Liu, Yanzhao Wu, Wenqi **Wenqi Wei** Cao, Semih Sahin, and Qi Zhang. "Benchmarking deep learning frameworks: Design considerations, metrics and beyond." IEEE International Conference on Distributed Computing Systems (ICDCS), pp. 1258-1269. IEEE, 2018.

Patent Filed

[2] Mu Qiao, **Wenqi Wei**, and Divyesh Jadav. "Graph Neural Network Ensemble Learning." U.S. Patent Application 17/562,080, 20230206029A1, 2023.

[1] Mu Qiao, **Wenqi Wei**, Eric Butler, and Divyesh Jadav. "Machine learning based vulnerable target identification in ransomware attack." U.S. Patent Application 17/113,464, US20220179964A1, 2022.

Teaching Experience

- Instructor: CISC5325 Databases, Fordham University 24 Spring
- Instructor: CISC4080 Computer Algorithms, Fordham University 23 Spring, 23 Fall
- Instructor: CISC5835 Algorithms for Data Science, Fordham University 23, 24 Spring, 23 Fall
- TA: CS6220 Big Data Systems, Georgia Tech 19, 20, 21 Fall
- TA: CS6675/CS4675 Advanced Internet Computing, Georgia Tech 19, 22 Spring

Services (More than 100 times PC/Reviewer experience)

- **Conference reviewer/PC**: ICDE18, ICDM (20,21), TheWebConf(21,23), ICLR-DPML21, KDD(21,22,23,24), MM21, Middleware21, NeurIPS-AI4Science21, ICML-AI4Science22, ML4H (20,21,22,23), AAAI(22,23,24), CVPR(22,23,24), ECCV(22,24), TPS(22,23), SDM(22,24), NeurIPS(22,23), ICWSM(23,24), IJCAI(23,24), ICCV23, VTC23, ISI23, HICSS24, WACV24, ICLR24, SACMAT24

- **Senior PC**: AAAI23-Safe and Robust Artificial Intelligence track

- **Conference chairing**: Publicity chair @ CIC/TPS/CogMI(22,23), Session chair @ (AAAI23,CIC/TPS/CogMI23), Student Mentoring Workshop Chair @ CIC/TPS/CogMI23, Tutorial co-chair @ IEEE BigData23

- **Journal reviewer**: IEEE TIFS, IEEE TMC, IEEE TNNLS, IEEE ToN, IEEE TNSE, IEEE TSC, IEEE CL, IEEE IoTJ, IEEE TBD, IEEE TKDE, IEEE TMC, ACM TOIT, ACM TIST, Elsevier JISA, Elsevier CHB, Elsevier IP&M, Springer SCIS, Springer ML, SCN

- **Distinguished Review Board**: ACM TWEB

- **NSF Panelist**: SafeLearning, SaTC

- **Internal Service**: Admission Committee for MS for Data Science, Fordham HPC Research Initiative review committee (2023-2026), Program committee of Fordham Cybersecurity program, session chair @ Fordham-IBM Workshop23, Fordham-IBM Fellow