Mengwei Yang

Email: mengwey@uci.edu · Tel: +1(213)716-2673 · Irvine, USA

Website: https://mengwey.github.io.

Education

University of California Irvine 2021.09 – Present

Ph.D., Electrical Engineering GPA: 4/4 Advisor: Prof. Athina Markopoulou

University of Southern California 2019.08 – 2021.05

M.S., Electrical Engineering GPA: 3.71/4.0

Northeastern University (China) 2015.09 – 2019.06

B.E., Electronic Information Engineering GPA: 93/100

Skills

- **LLMs and Agents:** Hands-on experience building agentic applications with LLMs (Large Language Models), including RAG (Retrieval-Augmented Generation), tool-use orchestration, memory management, MCP-based agent collaboration, and OpenAI Agents SDK
- **Voice and Multimodal AI:** Expertise in voice cloning, multilingual speech translation, and lip-synced video dubbing; integration of ASR (Automatic Speech Recognition), TTS (Text-to-Speech), and generative models into interactive agent pipelines
- **Privacy and Security:** Differential privacy, secure aggregation, and defenses against model leakage attacks
- **Programming Languages:** Python, C/C++, Verilog
- **Machine Learning Frameworks:** PyTorch, TensorFlow, TensorFlow Federated
- **Federated Learning (FL):** Secure FL frameworks, personalized FL, model pruning, knowledge distillation, and Shapley value–based client valuation

Work Experience

Syntiant Corp. Irvine, California

Engineering Intern *2022.07 – 2022.09*

- Developed a confidence-aware, multi-teacher knowledge distillation framework for keyword spotting task, leveraging the student-teacher architecture. Used pre-trained transformer models to enhance student model performance effectively.

Ericsson Inc. Santa Clara, California

Data Science Intern *2020.06 – 2021.01*

- Built a secure federated XGBoost framework with an innovative secure quantile sketch and practical secure aggregation. Implemented pairwise masking of model parameters to protect against gradient leakage attacks during aggregation, strengthening client data privacy.

UCIrvine EECS Department. Irvine, California

Teaching Assistant *Winter and Spring 2024, Winter and Spring 2025*

- Taught the lab sessions of EECS 31L: Introduction to Digital Logic Lab, guiding students through Verilog module design and debugging.

Research Experience

ReTalk Agent: Multilingual Video Dubbing 2025.8 – 2025.10

- Developed an AI agent pipeline that transforms input videos into multilingual versions by leveraging LLM-based language processing while preserving the speaker's unique voice identity.
- Implemented lip-synchronization and voice cloning techniques to ensure natural alignment of facial movements with translated speech, producing high-quality, culturally adaptive video outputs.

LLM-Powered Voice Assistant with Voice Cloning 2025.6 – 2025.8

- Designed and implemented an interactive voice assistant leveraging large language models (LLMs) as the core reasoning engine, enabling natural, context-aware conversations with users.
- Integrated voice cloning technology to generate a personalized synthetic voice, allowing seamless spoken interaction and enhancing user engagement in real-time dialogue.

Valuing Solo and Synergy in Federated Learning 2024.10 – 2025.7

- Proposed *DuoShapley*, an efficient Shapley value approximation that adaptively balances individual (Solo) and collaborative (Leave-One-Out) user contributions to improve valuation efficiency across heterogeneous federated learning settings.
- Demonstrated that DuoShapley achieves over 200× speedup and higher robustness with existence of noisy users, enabling scalable and reliable client selection for large-scale federated learning.

Maverick-Aware Shapley Valuation for Client Selection in FL 2023.10 – 2024.9

- Designed a Maverick-aware Shapley valuation framework to quantify client contributions under data heterogeneity, addressing the systematic undervaluation of clients with rare or underrepresented classes.
- Developed *FedMS*, a Shapley-guided client selection mechanism that adaptively prioritizes high-value clients, improving global performance and robustness against adversaries and free-riders across diverse FL scenarios.

PriPrune: Quantifying and Preserving Privacy in Pruned FL 2022.11 – 2023.8

- Conducted research on privacy guarantees for model pruning in FL, deriving information-theoretic upper bounds on information leakage in pruned FL models.
- Developed *PriPrune* – a privacy-aware model pruning algorithm featuring personalized, per-client defense masks and adaptive pruning rates, balancing both privacy and model performance effectively.

Information Leakage In Personalized Federated Learning 2022.4 – 2022.6

- Executed the gradient leakage attack (DLG attack) in personalized FL, demonstrating how varying personalization levels impact vulnerability to DLG attacks.
- Proposed *PerFed-LDP*, a per-example level differential privacy method for personalized FL, analyzing its privacy-utility tradeoff and quantifying DLG attack performance under differentially private settings.

Privacy by Projection: Federated Population Density Estimation by Projecting on Random Features 2021.12 – 2022.7

- Designed a federated kernel density estimation (KDE) framework to estimate population density while ensuring user data remains local.
- Developed a federated Random Fourier Feature (RFF) KDE approach that leverages a random feature representation, irreversibly projecting user information onto spatially delocalized basis functions, thus enhancing privacy without compromising estimation accuracy.

Location Leakage in Federated Signal Maps 2021.11 – 2022.4

- Executed the gradient leakage (DLG) attack on federated signal map prediction tasks, successfully reconstructing the average location from users' private spatio-temporal datasets during federated training.
- Proposed a defense strategy that strategically selects local batches to obfuscate the true average location, misleading DLG attacks. Conducted an analysis on the tradeoff between utility in federated signal mapping and privacy protection for clients' location data.

SaferQ: Obfuscating Search Queries via Generative Adversarial Privacy 2020.3 – 2020.5

- Developed SaferQ, an extension of the Generative Adversarial Privacy (GAP) framework for sequence generation, designed to obfuscate search queries while balancing privacy and utility. Deployed between browsers and query log databases, SaferQ anonymizes query logs according to specified privacy-utility trade-off criteria.

Secure Federated XGBoost Framework 2018.9 – 2019.8

- Designed a secure federated XGBoost framework that incorporates anonymized data aggregation to balance privacy and model performance.

Publications

(* denotes equal contributions)

1. **Mengwei Yang**, Baturalp Buyukates, Yanning Shen, Athina Markopoulou, Valuing Solo and Synergy in Federated Learning, *Under submission.*

2. **Mengwei Yang**, Baturalp Buyukates, Athina Markopoulou, Rewarding the Rare: Maverick-Aware Shapley Valuation in Federated Learning, *Under submission.*

3. Tianyue Chu*, **Mengwei Yang***, Nikolaos Laoutaris, Athina Markopoulou, PriPrune: Quantifying and Preserving Privacy in Pruned Federated Learning, *accepted in ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ToMPECS), October 2024.*

4. **Mengwei Yang**, Ismat Jarin, Baturalp Buyukates, Salman Avestimehr, Athina Markopoulou, Maverick-Aware Shapley Valuation for Client Selection in Federated Learning, *accepted and to appear in ISIT workshop on Informational Theoretic methods for Trustworthy ML (IT-TML), 2024.*

5. Tianyue Chu, **Mengwei Yang**, Nikolaos Laoutaris, Athina Markopoulou, Information-Theoretical Bounds on Privacy Leakage in Pruned Federated Learning, *accepted and to appear in ISIT workshop on Informational Theoretic methods for Trustworthy ML (IT-TML), 2024.*

6. Evita Bakopoulou*, **Mengwei Yang***, Jiang Zhang, Konstantinos Psounis, Athina Markopoulou, Location Leakage in Federated Signal Maps, *in IEEE Transactions on Mobile Computing (TMC), June 2024.*

7. Zixiao Zong, **Mengwei Yang**, Justin Ley, Carter T. Butts, Athina Markopoulou, Privacy by Projection: Federated Population Density Estimation by Projecting on Random Features, *to appear in Proceedings on Privacy Enhancing Technologies (PoPETs), 2023.*

8. **Mengwei Yang**, Linqi Song, Jie Xu, Congduan Li, Guozhen Tan. The tradeoff between privacy and accuracy in anomaly detection using federated XGBoost, *accepted and to appear in IJCAI Workshop: International Workshop on Federated Learning for User Privacy and Data Confidentiality (FL-IJCAI'19), 2019.*

References

Available upon request.