Yuechun Gu (谷岳錞)
Ph.D. Candidate in Computer Science

- **Phone: +1 (414) 285-8673, +86 18512837931**
- **Email: ygu2@umbc.edu**
- **LinkedIn: linkedin.com/in/ethanyuechun**
- **Homepage: yuechun-ethan-gu.net**

Profile

I am a Ph.D. candidate in Computer Science at the University of Maryland, Baltimore County (UMBC), working under the guidance of Dr. Keke Chen. My research focuses on privacy and confidentiality issues in machine learning, with a special interest in the practical application of differentially private ML systems and their implications in fields like physics and biomedical science.

Education

- **Ph.D. in Computer Science**
  University of Maryland, Baltimore County, Baltimore, MD
  2024 – Present
- **Ph.D. in Computer Science**
  Marquette University, Milwaukee, WI
  2021 – 2024
- **Bachelor of Science in Mathematics**
  University of Electronic Science and Technology of China
  2016 – 2020

Professional Experience

- **Research Assistant**
  UMBC, Baltimore, MD
  2024 – Present
  Conducting research on privacy-preserving machine learning, focusing on machine unlearning methodologies.
- **Research Assistant**
  Marquette University, Milwaukee, WI
  2021 – 2024
  Worked on inference attacks, privacy in machine learning, and dataset encoding. Published work in top-tier conferences and journals.
- **Systems Specialist**
  EF Education First, Tianjin, China
  2020 – 2021
  Developed and maintained a Salesforce-based data warehouse, generating continuous profit for EF Tianjin.
- **Research Assistant (Remote)**
  UCLA, Los Angeles, CA
  2018 – 2020

Collaborated on applying regression algorithms to economic analyses, which resulted in a publication at a top conference.

Research and Publications

Ongoing Projects

1. **Towards Membership Inference Attack Against Recommender Systems and CTR Prediction**
   Designing a likelihood-ratio-based attack against traditional recommender systems and novel CTR prediction systems, analyzing the implications of membership inference on these systems.
2. **Efficient Privacy Risk Estimation System**
   Optimizing the FT-PrivacyScore system for large-scale models (ViT, LLM), exploring efficient methods to estimate the privacy risk of participants.
3. **Auditing machine unlearning in Recommender Systems**
   We are trying to use privacy estimation tools to show the privacy protection given by the machine unlearning in recommender systems.
4. **Privacy-aware wound image classification system through dataset encoding (Collaborated with UW Milwaukee)**
   We are designing a private wound image classification system that maintains the classification performance and protects patients' privacy.

First Author Projects

1. **Calibrating Practical Privacy Risks for Differentially Private Machine Learning.**
   Accepted by IEEE Big Data (CCF C)
   Demonstrated that removing sensitive features significantly improves the trade-off between utility and privacy in differentially private models.
2. **Auditing Privacy Protection of Machine Unlearning.**
   Under review at ICLR
   Developed an efficient augmentation-based attack to evaluate the privacy risks of samples after applying novel machine unlearning methods.
3. **Gu, Y., & Chen, K. (2023). Adaptive Domain Inference Attack.**
   arXiv preprint arXiv:2312.15088, Under review at KDD 2025 (CCF-A)
   Proposed a tree-like architecture for crafting efficient domain inference attacks in scenarios where attackers have no prior domain knowledge.
4. **FT-PrivacyScore: Personalized Privacy Scoring Service for Machine Learning Participation.**
   ACM CCS 2024 (CCF-A)
   Developed a system that quantifies privacy risk in model fine-tuning tasks, allowing data contributors to assess privacy risks before participating.
5. **Gu, Y., Sharma, S., & Chen, K. (2023, November). Image Disguising for Scalable GPU-accelerated Confidential Deep Learning.**
   Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (pp. 3679-3681). **(CCF-A)**

Designed an image-disguising system for secure cloud-based model training, allowing users to encrypt data using RMT or AES.

6. **Gu, Y., & Chen, K. (2023, June). GAN-based Domain Inference Attack.** ***Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 12, pp. 14214-14222).*** **(CCF-A)**
   Developed a GAN-based method to infer likely domains of target models without prior domain knowledge, enhancing model inversion attack strategies.
7. **Gu, Y., Yan, D., Yan, S., & Jiang, Z. (2020, October). Price Forecast with High-frequency Finance Data: An Autoregressive Recurrent Neural Network Model with Technical Indicators.**
   Proceedings of the 29th ACM International Conference on Information & Knowledge Management (pp. 2485-2492). **(CCF-B)**
   Demonstrated that autoregressive recurrent networks with technical indicators outperformed traditional LSTM and GARCH models in financial data forecasting.

Second Author Projects

1. **Chen, K., Gu, Y., & Sharma, S. (2023). DisguisedNets: Secure Image Outsourcing for Confidential Model Training in Clouds.**
   ACM Transactions on Internet Technology, 23(3), 1-26. **(JCR Q1)**
   Presented a secure image disguising approach that enables users to outsource images for confidential, GPU-accelerated cloud model training.

Services and Involvement

- **Committee Member:** IEEE CogMI 2024
- **Reviewer:** KDD 2024, AAAI 2023,2024,2025, ASML 2024
