Yuntao Du

ytdu@purdue.edu
www.zealscott.com

Research Interests

Data Privacy; AI Security/Privacy; Trustworthy Machine Learning.

Education

2023 - 2028 (expected) **Purdue University**, USA
Ph.D. in Computer Science (Advisor: Ninghui Li)

2020 - 2023 **Zhejiang University**, China
M.E. in Computer Technology. *Rank 1/49.*

2016 - 2020 **East China Normal University**, China
B.E. in Data Science, *GPA 3.61/4 (rank 1/23)*

Selected Honors and Awards

2026 **NDSS Fellowship**, Internet Society (1 of 24 worldwide)
2023 **Ross Fellowship**, Purdue University (1 of 10)
2023 **Herbold Scholarship**, Purdue University (1 of 7)
2023 **Presidential Doctoral Excellence Award**, Purdue University (1 of 150)
2023 **Excellent Masters Dissertation**, China (1 of 43)
2023 **Provincial Outstanding Graduates**, Zhejiang, China
2022 **National Scholarship**, China (0.1%)
2021 **National Scholarship**, China (0.1%)

Working Experience

2024.05 – 08 **Alibaba US**, Seattle, USA
Research Intern (Mentor: Dr. Bolin Ding)

2022 – 2023 **Microsoft Research Aisa**, Beijing, China
Research Intern, Award of Excellent Internship (top 10%) (Mentor: Dr. Xing Xie)

Publications

First & Co-First Author: CCS×1, USENIX Security×2, NDSS×1, VLDB×1, WWW×1, SIGIR×1, TKDE×1, TCDE×1.

Mentored Student Publications: USENIX Security×1, SIGIR×1, ICDE×2.

(equal contribution, † mentored student)*

Manuscript **Yuntao Du***, Tanishq Praveen Pauskar†*, Hao Wang, Jing Su, Ninghui Li. Privacy Leakage from a Thousand Words: Sub-Pixel Location Recovery from Dot Maps.

Manuscript **Yuntao Du**, Zitao Li, Bolin Ding, Yaliang Li, Hanshen Xiao, Jingren Zhou, Ninghui Li. Automated Profile Inference with Language Model Agents.

USENIX 2026 **Yuntao Du**, Yuetian Chen, Hanshen Xiao, Bruno Ribeiro, Ninghui Li. Imitative Membership Inference Attack. In USENIX Security Symposium (USENIX Security), 2026.

USENIX 2026 Meng Tong[†*], **Yuntao Du**[*], Kejiang Chen, Weiming Zhang, Ninghui Li. Membership Inference Attacks on Tokenizers of Large Language Models. In USENIX Security Symposium (USENIX Security), 2026.

USENIX 2026 Yuetian Chen, **Yuntao Du**, Kaiyuan Zhang, Ashish Kundu, Charles Fleming, Bruno Ribeiro, Ninghui Li. Window-based Membership Inference Attacks Against Fine-tuned Large Language Models. In USENIX Security Symposium (USENIX Security), 2026.

ICLR 2026 Yuetian Chen, Kaiyuan Zhang, **Yuntao Du**, Edoardo Stoppa, Charles Fleming, Ashish Kundu, Bruno Ribeiro, Ninghui Li. Membership Inference Attacks Against Fine-tuned Diffusion-Based Language Models. In International Conference on Learning Representations (ICLR), 2026.

NDSS 2026 **Yuntao Du**, Jiacheng Li, Yuetian Chen, Kaiyuan Zhang, Zhizhen Yuan, Hanshen Xiao, Bruno Ribeiro, Ninghui Li. Cascading and Proxy Membership Inference Attacks. In Network and Distributed System Security Symposium (NDSS), 2026.

IEEE Data Eng. Bulletin **Yuntao Du**, Zitao Li, Ninghui Li, Bolin Ding. Beyond Data Privacy: New Privacy Risks for Large Language Models. In IEEE Data Engineering Bulletin. 2025.

CCS 2025 **Yuntao Du**, Ninghui Li. Systematic Assessment of Tabular Data Synthesis Algorithms. In ACM Conference on Computer and Communications Security (CCS), 2025.

USENIX 2025 Kaiyuan Zhang, Siyuan Cheng, Hanxi Guo, Yuetian Chen, Zian Su, Shengwei An, **Yuntao Du**, Charles Fleming, Ashish Kundu, Xiangyu Zhang, Ninghui Li. SOFT: Selective Data Obfuscation for Protecting LLM Fine-tuning against Membership Inference Attacks. In USENIX Security Symposium (USENIX Security), 2025.

Before Ph.D.

ICDE 2024 Yujia Hu[†], **Yuntao Du**, Zhikun Zhang, Ziquan Fang, Lu Chen, Kai Zheng, Yunjun Gao. Real-Time Trajectory Synthesis with Local Differential Privacy. In IEEE International Conference on Data Engineering (ICDE), 2024.

VLDB 2023 **Yuntao Du**, Yujia Hu, Zhikun Zhang, Ziquan Fang, Lu Chen, Baihua Zheng, Yunjun Gao. LDPTrace: Locally Differentially Private Trajectory Synthesis. In International Conference on Very Large Data Bases (VLDB), 2023.

WWW 2023 **Yuntao Du**, Jianxun Lian, Jing Yao, Xiting Wang, Mingqi Wu, Lu Chen, Yunjun Gao, Xing Xie. Towards Explainable Collaborative Filtering with Taste Clusters Learning. In ACM Web Conference (WWW), 2023.

SIGIR 2023 Xinjun Zhu†, **Yuntao Du**, Yuren Mao, Lu Chen, Yujia Hu, Yunjun Gao. Knowledge-refined Denoising Network for Robust Recommendation. In International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 2023.

TKDE 2023 Minjun Zhao, Lu Chen, Keyu Yang, **Yuntao Du**, Yunjun Gao. Finding Materialized Models for Model Reuse. In IEEE Transactions on Knowledge and Data Engineering (TKDE), 2023.

ICDE 2023 Zhihao Zeng†, **Yuntao Du**, Ziquan Fang, Lu Chen, Shiliang Pu, Guodong Chen, Hui Wang, Yunjun Gao. FLBooster: A Unified and Efficient Platform for Federated Learning Acceleration. In IEEE International Conference on Data Engineering (ICDE), 2024.

KDD 2022 Ziquan Fang, **Yuntao Du**, Xinjun Zhu, Danlei Hu, Lu Chen, Yunjun Gao, Christian S. Jensen. Spatio-Temporal Trajectory Similarity Learning in Road Networks. In ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), 2022.

SIGIR 2022 Yunjun Gao, **Yuntao Du**, Yujia Hu, Lu Chen, Xinjun Zhu, Baihua Zheng. Self-Guided Learning to Denoise for Robust Recommendation. In International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 2022.

SIGIR 2022 **Yuntao Du**, Xinjun Zhu, Lu Chen, Baihua Zheng, Yunjun Gao. HAKG: Hierarchy-Aware Knowledge Gated Network for Recommendation. In International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR), 2022.

TKDE 2022 **Yuntao Du**, Xinjun Zhu, Lu Chen, Ziquan Fang, Yunjun Gao. MetaKG: Meta-learning on Knowledge Graph for Cold-start Recommendation. In IEEE Transactions on Knowledge and Data Engineering (TKDE), 2022.

ICDE 2021 Ziquan Fang, **Yuntao Du**, Xinjun Zhu, Lu Chen, Ziquan Fang, Yunjun Gao. E^2DTC: An End to End Deep Trajectory Clustering Framework via Self-Training. In IEEE International Conference on Data Engineering (ICDE), 2021.

VLDB 2021 Ziquan Fang, Lu Pan, Lu Chen, **Yuntao Du**, Yunjun Gao. MDTP: A Multi-source Deep Traffic Prediction Framework over Spatio-Temporal Trajectory Data. In International Conference on Very Large Data Bases (VLDB), 2021.

Patents

U.S. Patent Mingqi Wu, Jianxun Lian, **Yuntao Du**, Jing Yao, Xiting Wang, Bei Lu, Xing Xie. Explainable Cluster-based Collaborative Filtering. U.S. Patent Application No. 18/191,681.

Invited Talks

2026 Feb **NDSS 2026**
Cascading and Proxy Membership Inference Attacks

2025 Oct **ACM CCS 2025**
Systematic Assessment of Tabular Data Synthesis Algorithms

2025 Jun **Research Trend**
Automated Profile Inference with Language Model Agents

2025 Mar **Graduate Research Symposium, Purdue University**
Automated Profile Inference with Language Model Agents

2023 Nov **AI Time**
Towards Explainable Collaborative Filtering with Taste Clusters Learning

2023 Jan **Microsoft Research Aisa**
Explainable Clustering and Cluster-based Collaborative Filtering

Teaching & Mentoring

Teaching

Teaching Assistant, CS556: Data Security and Privacy, Purdue University. 2025

Mentoring

Kimaya Deshpande, B.S. Computer Science, Purdue University. 2026 - now
Annabelle Nakamura DeFosse, B.S Computer Engineering, Purdue University. 2026 - now
Pauskar Tanishq, B.S. Computer Science, Purdue University. 2025 - now
Tong Meng, M.S. Computer Science, University of Science and Technology of China. 2025 - now
Hillary Yang, high school student, Carmel High School. 2024
Yujia Hu, M.S. Computer Science, Zhejiang University. 2022 - 2024
Zhihao Zeng, M.S. Computer Science, Zhejiang University. 2022 - 2023
Xinjun Zhu, M.S. Computer Science, Zhejiang University. 2022 - 2023

Academic Service

Conference Reviewing

PC, International Conference on Learning Representations (ICLR) 2025 - 2026
PC, International Conference on Artificial Intelligence and Statistics (AISTATS) 2025 - 2026
PC, ACM Web Conference (WWW) 2026
PC, ACM Conference on Data and Application Security and Privacy (CODASPY) 2026
PC, International Conference on Web Search and Data Mining (WSDM) 2026
PC, International Joint Conference on Neural Networks (IJCNN) 2025
PC, Table Representation Learning Workshop (TRL) 2024 - 2025
PC, ACM Conference on Information and Knowledge Management (CIKM) 2024 - 2025
PC, ACM Conference on Research and Development in Information Retrieval (SIGIR) 2023 - 2026
PC, AAAI Conference on Artificial Intelligence (AAAI) 2023 - 2026
PC, ACM Conference on Information Retrieval in the Asia Pacific (SIGIR-AP) 2023 - 2024

Journal Reviewing

Reviewer, ACM Computing Surveys (CSUR)

Reviewer, IEEE Transactions on Dependable and Secure Computing (TDSC)
Reviewer, ACM Transactions on Privacy and Security (TOPS)
Reviewer, International Journal on Very Large Data Bases (VLDBJ)
Reviewer, ACM Transactions on Recommender Systems (TORS)
Reviewer, IEEE Transactions on Knowledge and Data Engineering (TKDE)
Reviewer, IEEE Transactions on Big Data (TBD)