Elisa Tsai

Email: eltsai@umich.edu Homepage: eltsai.github.io

INTERESTS

Web security; machine learning for security; machine learning efficiency.

My research focuses on building pragmatic machine learning systems for security [1]. I also design algorithms for data efficiency and inference efficiency for vision and large language models [2], [O1], [O2].

EDUCATION

University of Michigan, Ann Arbor 2020 - present
Ph.D. Candidate, Computer Science
Advisor: Prof. Atul Prakash

Univeristy of Science and Technology of China (USTC) 2020
B.S., Computer Science and Technology

WORK EXPERIENCE

Machine Learning Engineer Intern May 2025 - Aug 2025
I developed agentic LLM systems for the risk team. **Stripe**

Open-Source Contributor June 2023 - Aug 2023
I contributed to Riotpot, implementing multiple protocol emulations to increase the honeypot's relevance to real-world attacks. **Google Summer of Code - Honeynet RiotPot**

SELECTED PUBLICATIONS

1. Harmful Terms and Where to Find Them: Measuring and Modeling Unfavorable Financial Terms and Conditions in Shopping Websites at Scale
   Elisa Tsai, Neal Mangaokar, Boyuan Zheng, Haizhong Zheng, Atul Prakash
   WWW (The Web Conference) 2025 (Oral)
2. Label-Free Coreset Selection with Proxy Training Dynamics
   Haizhong Zheng (co-lead), Elisa Tsai (co-lead), Yifu Lu, Jiachen Sun, Brian R. Bartoldson, Bhavya Kailkhura, Atul Prakash
   ICLR (The International Conference on Learning Representations) 2025
3. Class-Proportional Coreset Selection for Difficulty-Separable Data
   Elisa Tsai, Haizhong Zheng, Atul Prakash
   ICCV Workshop on Curated Data for Efficient Learning (CDEL) 2025
4. Terms of Deception: Exposing Obscured Financial Obligations in Online Agreements with Deep Learning
   Elisa Tsai, Anoop Singhal, Atul Prakash
   DLSP (Deep Learning Security and Privacy Workshop) 2024
5. Detecting Social Engineering Scams While Preserving User Privacy in the Digital Era (Proposal Position Paper)
   Atul Prakash, Shivani Kumar, Elisa Tsai
   ConPro (Workshop on Technology and Consumer Protection) 2024

ONGING WORK

1. **LLM human preference data efficiency**: Investigating strategies to optimize data selection for fine-tuning large language models (LLMs) on human preference datasets, with a focus on maximizing performance while minimizing data usage.
2. **LLM inference efficiency**: Developing a parameter-efficient, lightweight adapter to improve LLM inference efficiency through dynamic, efficiency-aware training.

GRANT PROPOSALS

I actively contributed to the proposal design, proposal writing, and presentation for the following grants:

Data Efficiency of LLMs Fine-tuning with RLHF — Cisco, 2023
$150*K per year* — PI: Atul Prakash

Intelligent Assistants for Detecting Social Engineering Scams — OpenAI, 2023
$100*K* — PI: Atul Prakash

TEACHING

EECS 588 Computer & Network Security , **Grad Student Instructor** *Winter 2024, UMich*

EECS 281 Data Structures and Algorithms , **Grad Student Instructor** *Fall 2023, UMich*

EECS 598 Secure and Trustworthy ML , **Grad Student Instructor** *Winter 2023, UMich*

SERVICE

- SECRIT (SECurity Reading Is Terrific) Reading Group Host 2021 – 2024
- CSEG (CSE Graduate Students) Outreach Chair 2022 – 2023
- CSEG (CSE Graduate Students) Social Co-Chair 2022 – 2023