Charlie Hou

Github: houcharlie · Email: hou.charlie2@gmail.com · Google Scholar · Website

Professional summary

I am a 5th year Ph.D. student at CMU working on **large language models (LLMs), learning-to-rank, security, and privacy**. Some research/work highlights: (1) I used **LLMs** to generate **privacy-preserving synthetic data** from users that is **high quality** enough to **replace on-device training** (2) I developed a way for **rankers (in recommendations, relevance)** to learn from unlabeled data.

Education

Carnegie Mellon University · Pittsburgh, PA……..……Aug 2019-May 2024(expected)
Ph.D. Candidate, Electrical and Computer Engineering, Advisor: Giulia Fanti
GPA: 3.92/4.0

Princeton University · Princeton, NJ……..……Sep 2015-Jun 2019
BSE, Major: Operations Research and Engineering, Minor: Applied Math
GPA: 3.843/4.0 (Graduated Magna cum laude)

Work Experience

Research/work experience in: Large language models (LLMs), differential privacy (DP), learning-to-rank (LTR), on-device/federated learning, security, reinforcement learning

---

Meta · AI research scientist intern · Redmond, WA……...Jun 2023-Sep 2023
Supervisor: Daniel Lazar

- Used **LLMs** to generate **differentially private synthetic data** that is high quality enough that it **(1) replaces the need for federated learning (2) allows customizing LLMs for user datasets.** Paper in preparation/legal review.

Meta · AI research scientist intern · Redmond, WA……...Sep 2022-Jan 2023
Supervisor: Daniel Lazar

- Developed FreD, a **differentially private** way to select pretraining datasets for federated learning training using **large language models (LLMs)**. FreD is a differentially private way to select the best pretraining dataset for **federated learning** and is capable of choosing between datasets even under high privacy constraints (eps=0.6).

- Produced a paper (Arxiv). This was presented at the ICLR 2023 TrustML workshop.

Amazon · Applied Science Intern · Palo Alto, CA……Jun 2022-Sep 2022
Supervisor: Sujay Sanghavi

- Developed new pretraining strategies for rankers which **improved the robustness** of Amazon shopping search engine rankers **by over 20% on retrieval metrics**
- Produced a paper (Arxiv) which was presented in the **ICML 2023 MFPL workshop as an oral presentation**. Full paper currently under submission at conference.

Amazon · Applied Science Intern · Seattle, WA……..Jun 2021-Sep 2021
Supervisor: Greg Herman

- Developed an **epsilon-greedy bandit algorithm** for product selection that outperforms Amazon's economics-based model in median profit within 30 selection periods, and outperforms the previous model in median profit by 4% after 100 selection periods.
- Introduced a **novel simulation-based offline evaluation framework** for **RL algorithms** on product selection. Amazon was previously using backtests, which use past sales data. These could not produce counterfactuals for experimental product selections.

Uber · Research Intern · San Francisco, CA……Jun 2019-May 2020
Supervisor: Ersin Yumer

- Developed a radar simulation model that predicts radar detections via combining classical physical simulation and the U-net **computer vision** architecture
- U.S. Patent submitted: "Radar Simulation".

Publications

---

1. **Charlie Hou,** Kiran Koshy Thekumparampil, Michael Shavlovsky, Giulia Fanti, Yesh Dattatreya, Sujay Sanghavi. "Pretrained deep models outperform GBDTs in Learning-To-Rank under label scarcity". In: **ICML 2023 workshop for preference-based learning as an oral presentation.** Under submission for full conference. Arxiv
2. **Charlie Hou,** Hongyuan Zhan, Akshat Shrivastava, Sid Wang, Sasha Livshits, Giulia Fanti, Daniel Lazar. "Privately Customizing Prefinetuning to Better Match User Data in Federated Learning". In: **ICLR 2023 workshop for TrustML**. Arxiv
3. **Charlie Hou**, Kiran K. Thekumparampil, Giulia Fanti, Sewoong Oh. "FedChain: Chained Algorithms for Near-Optimal Communication Cost in Federated Learning". In: **ICLR 2022**, also appeared at **ICML-FL 2021 workshop as an oral presentation.** Arxiv, Github
4. **Charlie Hou**, Kiran K. Thekumparampil, Giulia Fanti, Sewoong Oh. "Efficient Algorithms for Saddle Point Optimization". Arxiv

5. **Charlie Hou***, Mingxun Zhou*, Yan Ji, Phil Daian, Florian Tramer, Giulia Fanti, Ari Juels. (* represents equal contribution) “SquirRL: Automating Attack Analysis on Blockchain Incentive Mechanisms with Deep Reinforcement Learning”. In: **NDSS 2021**. Arxiv, Github

Technical Skills

Tensorflow, Python, NumPy, Tensorflow-Federated, Pytorch, Pandas, AWS Batch/ECR, Docker

Relevant Coursework

Machine learning: convex optimization, learning theory, online learning, topics in machine learning systems, sketching & streaming algorithms for big data, differential privacy
Computer science: CS theory toolkit, algorithms & data structures, functional programming, discrete mathematics, linear programming
Mathematics: real analysis, high dimensional probability, probability and stochastic processes, mathematical statistics, mathematical finance

Professional Service

---

1. Student Volunteer at Symposium on Theory for Computing (STOC) 2020
2. Sub-reviewer for ICML 2022 (for Kiran K. Thekumparampil, reviewed 5 papers)
3. Reviewer for ICLR 2023
4. Reviewer for NeurIPS 2023

Other activities

Tiger Chef Champion (2018), CMU running club member