

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, and privacy**. In the past, I have also worked on **security**.

Education

Carnegie Mellon University · Pittsburgh, PA.....Aug 2019-Aug 2024
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

- Developed PrE-Text: a method to train LLMs of any size with private federated data.
- Paper accepted as an **Oral Presentation at ICML 2024**.
- Paper accepted as an **Honorable Best Paper Award + Oral Presentation at ICLR 2024 PrivML workshop**.

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)**.
- Paper accepted 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**

- Paper accepted as an **Oral Presentation at ICML 2023 MFPL workshop**. 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 in warehouses.
- Introduced a **novel simulation-based offline evaluation framework** for **RL algorithms** on product selection.

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**, Akshat Shrivastava, Hongyuan Zhan, Rylan Conway, Trang Le, Adithya Sagar, Giulia Fanti, Daniel Lazar. “PrE-Text: Training Language Models on Private Federated Data in the Age of LLMs”. **Oral Presentation at ICML 2024, Oral Presentation at ICLR 2024 PrivML workshop**.
2. Shuqi Ke, **Charlie Hou**, Giulia Fanti, Sewoong Oh. “On the Convergence of Differentially-Private Fine-tuning: To Linearly Probe or to Fully Fine-tune?”. Under submission.
3. **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”. **Oral Presentation at ICML 2023 MFPL workshop**. Under submission.
4. **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”. **ICLR 2023 TrustML workshop**.
5. **Charlie Hou**, Kiran K. Thekumparampil, Giulia Fanti, Sewoong Oh. “FedChain: Chained Algorithms for Near-Optimal Communication Cost in Federated Learning”. **ICLR 2022, Oral Presentation at ICML 2021 FL workshop**.
6. **Charlie Hou**, Kiran K. Thekumparampil, Giulia Fanti, Sewoong Oh. “Efficient Algorithms for Saddle Point Optimization”.
7. **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”. **NDSS 2021**.

Awards

1. Honorable Mention for Best Paper at ICLR 2024 PrivML Workshop
2. Google Collabs 2022 Research Award (\$100k grant)
3. Tiger Chef Champion 2018

Professional Service

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

Other activities

[Tiger Chef](#) Champion (2018), CMU running club member