

Hanhan Zhou

800 22nd St NW, Washington, DC 20052
Github: github.com/hanhanAnderson/

Email: hanhan@gwu.edu
Mobile: +1-202-290-4126
LinkedIn: [linkedin.com/in/hanhanzhou/](https://www.linkedin.com/in/hanhanzhou/)

SUMMARY: Current Ph.D. student with full-stack programming knowledge and a wide range of working and project experience in machine learning, research areas include **Reinforcement Learning, Distributed Learning, Generative AI (Foundation Models and Large Language Models)**.

EDUCATION

The George Washington University

Ph.D. Student in Department of Electrical and Computer Engineering

Research Interests: Reinforcement Learning, Distributed Computing, Generative Models, Recommendation Systems, Cyber Security

Washington, DC
Sept 2019 — May 2024

The George Washington University

M.S. - Electrical and Computer Engineering

Core Courses: Design & Analysis of Algorithm, Computer System Architecture, Embedded Systems, Telecommunications Security

Washington, DC
Sept 2017 — May 2019

Zhejiang Sci-tech University

B.S. - Electronic and Information Engineering

Hangzhou, China
Sept 2013 — May 2017

SKILLS

- **Languages:** Object Oriented Design with **Python, JAVA, Go, Kotlin, JavaScript, MATLAB**
- **Frameworks:** PyTorch, Scikit, Keras, TensorFlow, Hugging Face, LangChain, Keras, Spring, Android Studio, Streamlit
- **Tools&DB:** Kubernetes, Docker, GIT, Jenkins, JIRA, Maven, Gradle, MongoDB, ChromaDB, Pinecone, MySQL

SELECTED PUBLICATIONS

1. Yongsheng Mei*, **Hanhan Zhou***, Tian Lan, “Projection-Optimal Monotonic Value Function Factorization in Multi-Agent Reinforcement Learning”: Accepted at the Int. Conf. Auton. Agents & MultiAgent Syst (**AAMAS**) 2024.
2. Huiqun Li, **Hanhan Zhou**, Yifei Zou, Dongxiao Yu, and Tian Lan, “ConcaveQ: Non-Monotonic Value Function Factorization via Concave Representations in Deep Multi-Agent Reinforcement Learning”: Accepted at the **AAAI** 2024.
3. **Hanhan Zhou**, Tian Lan, Guru Venkataramani, Wenbo Ding, “Every Parameter Matters: Ensuring the Convergence of Federated Learning with Dynamic Heterogeneous Models Reduction”: Accepted at **NeurIPS** 2023.
4. **Hanhan Zhou**, Tian Lan, and Vaneet Aggarwal, “Value Functions Factorization with Latent State Information Sharing in Decentralized Multi-Agent Policy Gradients”: IEEE Transactions on Emerging Topics in Comput. Intelligence, 2023.
5. Yongsheng Mei, **Hanhan Zhou**, Tian Lan, Guru Venkataramani, and Peng Wei, “MAC-PO: Multi-agent Experience Replay via Regret Minimization”: Accepted at the Int. Conf. Auton. Agents & MultiAgent Syst (**AAMAS**) 2023.
6. **Hanhan Zhou**, Tian Lan, and Vaneet Aggarwal, “Double Policy Estimation for Importance Sampling in Sequence Modeling-Based Reinforcement Learning”: Accepted at Foundation Models for Decision Making -NeurIPS, 2023.
7. Chang-Lin Chen, **Hanhan Zhou**, Jiayu Chen et al. “Two-tiered Online Optimization of Region-wide Datacenter Resource Allocation via Deep Reinforcement Learning”: submitted to IEEE Transactions on Network and Service Management.

8. **Hanhan Zhou**, Tian Lan, and Vaneet Aggarwal, “PAC: Assisted Value Factorisation with Counterfactual Predictions in Multi-Agent Reinforcement Learning”: Accepted at **NeurIPS** 2022.
9. **Hanhan Zhou**, Tian Lan, Guru Venkataramani, Wenbo Ding, “Heterogeneous Models with Online Pruning in Federated Learning”: Accepted at FL -NeurIPS, 2022.
10. **Hanhan Zhou**, Tian Lan, Guru Venkataramani, “PT-VTON: an Image-Based Virtual Try-On Network with Progressive Pose Attention Transfer”: arXiv preprint, 2021
11. **Hanhan Zhou**, Tian Lan, Guru Venkataramani, “Hunting Garbage Collection Related Concurrency Bugs through Critical Condition Restoration”: In Proceedings of the 2020 ACM CCS Workshop on Forming an Ecosystem Around Software Transformation (FEAST), 2020

ACADEMIC & INDUSTRY EXPERIENCE

Graduate Research Assistant

Jun 2019 — Present
Washington, DC

The George Washington University

Research Topic 1: Multi-Agent Reinforcement Learning

- Analyzed the drawbacks of popular state-of-the-art Multi-Agent Reinforcement Learning algorithms and proposed several frameworks to address them by optimizing the learning algorithm or introducing more effective designs.
- Evaluated proposed methods on several baselines including the StarCraft II challenge and demonstrated a substantial performance improvement by over 30%, several works are published in NeurIPS, AAI and AAMAS.

Research Topic 2: Federated and Distributed Learning

- Proposed a unified framework for optimization algorithms on Federated Learning and conducted convergence analysis with heterogeneous local clients through model pruning and quantization.
- Validated the proposed theory and provided actionable insights for designing and developing models of local clients of heterogeneous federated learning for foundation and large language models. Several works have been accepted in FL-NeurIPS 2022 and NeurIPS 2023, etc.

Research Topic 3: Offline Reinforcement Learning and Optimization

- Investigated deficiencies in Sequence-Modeled offline reinforcement learning, such as Decision Transformers, focusing on their high variance in complex environments and susceptibility to long-term dependencies.
- Assessed the proposed algorithm against various SOTA methods, showcasing up to 20% performance improvement and variance reduction, with results presented at foundation models for decision making workshop at NeurIPS.

Research Topic 4: Cyber and Software Security

- Conducted software security testing using AFL-based fuzzing tools on communication protocols like SSL and proposed a machine-learning-based privacy-preserving app for Android using the Xposed framework.
- Studied the occurrence of concurrent bugs in JavaScript inside the WebKit engine, and proposed a framework that generates critical conditions to promote the reproduction of concurrency-related bugs within limited execution overhead, the work is accepted at ACM-CCS-FEAST 2020.
- Examined the current network intrusion detection problem and proposed the real-time network intrusion detection modeled as a sequence modeling sequence of arriving packets using causal decision transformer, the work is accepted at Artificial Intelligence for Cyber Security workshop at AAI 2024.

App Development Intern - Full Stack

Jun 2018 - August 2018
New York, NY

AdviceCoach LLC

- Designed and developed a mobile-friendly website using React.js and Redux, allowing users to edit and create playbooks on their mobile phones' browser before installing the app.

- Implemented REST APIs, which allows the frontend to interact with the backend server, for customizing the playbook, sending messages to other customers and receiving feedback, etc.
- Tested the iOS app with the latest version, debugged, and improved user experience based on users' feedback.

Software Engineer Intern - Embedded Systems & Design

Lanxum Gushenxing NetSec Inc.

Dec 2016 - Mar 2017

Hangzhou, China

- Worked in a group of six software engineers responsible for designing and testing industrial RTU (Remote Terminal Unit) and ZigBee Modules.
- Implemented customized models of industrial firewalls like anti-sniffing, and enhanced whitelist (using snort etc.) based on clients' requirements.
- Designed preset safe configuration rules page on a Node.js based monitoring portal Website.
- Performed safety check for clients' network robustness and security with GE Achilles

SELECTED DEVELOPMENT PROJECTS

Weekend: A Event Search and Ticket Recommendation with fine-tuned LLM-based ChatBot

- Developed an interactive web page for users to search events and purchase tickets based on their geo-location with MongoDB stored data from TicketmasterAPI and user information.
- Implemented Elastic Stack (Elasticsearch, Logstash, Kibana) for advanced user profiling and insightful data analysis.
- Designed an LLM agent using APICChain as Retrieval Augmented Generation for real-time event recommendations.
- Instruction-Finetuned a Flan-T5 on AWS Sagemaker for Recommendation task with Tuning framework for Aligning LLMs with Recommendations(TALLRec) which increased 10% performance on MovieLens(100k) Benchmark.

Circa: A Geo-indexed Based Social Network WebApp with Stable-Diffusion Image Generation

- Built a geo-based social network webapp using React.JS and React Router with Material UI Design, with a local Stable Diffusion based image generation module using Gradio library.
- Developed a scalable web service in Golang to handle requests and deployed to Google App Engine flex.
- Utilized ElasticSearch (on Google Compute Engine) to provide geolocation-based search functions and Google Dataflow to implement a daily dump of posts to BigQuery table for offline analysis.
- Integrated a remote Stable-Diffusion-XL-Refiner using HuggingFace API with the Golang service.

AliDada: An online Shopping WebApp with DRL-based Recommendation System

- Developed an online shopping web app with Spring MVC for item search and Spring Web Flow for order processing.
- Implemented recommendation system by training a reinforcement learning agent under soft-actor-critic algorithm which improved 15% click-through rate on JESTER dataset.
- Employed Hibernate to provide better support for database operations.
- Integrated an ad bidding system based on second-highest-bid to deliver ads on the website.

AWARDS

NeurIPS Scholar Award- 2022, 2023

GW University Fellowship - 2019 — 2023

Runner Up Prize at GW New Venture Competition - 2021

Lin Weng Graduate Scholarship - 2021

Facebook Research Scholarship - 2019

GW SEAS Graduate Ambassador Fellowship - 2019

PRESENTATIONS

ACM CCS FEAST 2019

Nov 2019, London, United Kindom

ACM CCS FEAST 2020

Nov 2020, Online

International Workshop on Federated Learning @NeurIPS

Dec 2022, New Orleans, LA

Foundation Models for Decision Making @NeurIPS

Dec 2023, New Orleans, LA

SERVICES

- Journal Reviewer

- IEEE Transactions on Artificial Intelligence (TAI)
- IEEE/ACM Transactions on Networking (TNET)
- IEEE Transactions on Communications (TCOM)

- Conference Reviewer

- IEEE International Conference on Computer Communications (INFOCOM)
- Advances in Neural Information Processing Systems (NeurIPS)
- International Conference on Learning Representations(ICLR)
- New Frontiers in Graph Learning (GLFrontiers)
- Foundation Models for Decision Making (FMDM), NeurIPS Workshop
- Temporal Graph Learning Workshop, NeurIPS Workshop

REFERENCES

Prof. Tian Lan

Department of Electrical and Computer Engineering, the George Washington University

E-mail: tlan@gwu.edu

Professor

Washington, DC

<https://www2.seas.gwu.edu/~tlan/>

Prof. Guru Venkataramani

Department of Electrical and Computer Engineering, the George Washington University

E-mail: guruv@gwu.edu

Professor

Washington, DC

<https://www2.seas.gwu.edu/~guruv/>

Prof. Vaneet Aggarwal

School of Industrial Engineering, Purdue University

E-mail: vaneet@purdue.edu

Professor

West Lafayette, IN

<https://web.ics.purdue.edu/~vaneet/>