Su Zhang

suzhang94@gmail.com · (812)349-8473 · Homepage · LinkedIn

Research Interests

Machine Learning, Reinforcement Learning, Human-in-the-loop RL, Agent-based Systems, LLM Applications

Education

Washington State University — Pullman, WA
Ph.D. in Computer Science, GPA: 3.9/4.0 — July 2025
Advisor: Dr. Matthew E. Taylor

Indiana University — Bloomington, IN
M.S. in Computer Science, GPA: 3.97/4.0 — May 2017

Wuhan University — Wuhan, China
B.S. in Computer Science & B.S. in Economics, GPA: 3.48/4.0 — June 2015

Professional Experience

Nirva Labs — Pullman, WA
Machine Learning Engineer — Oct. 2025 - Present

- Architected end-to-end audio ML pipeline integrating multiple models for speaker diarization, voiceprint matching, and LLM-driven content extraction
- Built LLM-powered analytics features with observability tracing, caching layers, and configurable thresholds via RESTful APIs
- Developed adaptive scoring and classification pipelines for multi-modal signal processing with real-time visualization on iOS
- Implemented 15+ native iOS screens (SwiftUI) with offline-first local persistence, including onboarding flows, real-time chat UI, and data visualization dashboards
- Contributed to production Python backend (FastAPI) with vector search and relational database integration, deployed on AWS with infrastructure-as-code
- Designed evaluation framework with custom quality metrics for LLM-generated conversational outputs

Washington State University — Pullman, WA
Staff Assistant – Enterprise Systems / BI — Sep. 2021 - Jun. 2025

- Developed, modified, and validated SQL logic, ETL processes, and data tables in OBIEE/OAC for university-wide enrollment and demographic reporting
- Created and standardized dashboards, data canvases, and visualizations across multiple organizational levels for the Graduate School Fact Book

Selected Publications

[1] **Zhang, S.**, Das, S., Ganapathi Subramanian, S., & Taylor, M. (2023). Two-Level Actor-Critic Using Multiple Teachers. *TMLR & AAMAS 2023.*

[2] **Zhang, S.**, & Taylor, M. (2018). Enhanced Learning from Multiple Demonstrations with a Two-level Structured Approach. *ALA Workshop at AAMAS 2018.*

[3] Zhang, Y., **Zhang, S.**, & Leake, D. (2017). Maintenance for Case Streams: A Streaming Approach to Competence-Based Deletion. *ICCBR*.

Selected Research

Multiagent Coordination via Option Value Decomposition 2023 – 2024

- Proposed multi-agent option-critic framework (CTDE); addressed sparse reward with optimistic heuristic integration

Two-Level Actor-Critic Using Multiple Teachers Mar. 2021 - Nov. 2023

- Developed hierarchical RL framework for learning from heterogeneous-quality demonstrations via teacher selection and low-level policy optimization; published in TMLR and AAMAS

Efficient Exploration with Probability Map May 2019 - Jan. 2021

- Leveraged prior probability distributions as exploration heuristics to accelerate RL agent training in robotic search tasks

Technical Skills

Languages:	Python, C/C++, Java, Swift, SQL, Lua
ML & AI:	PyTorch, TensorFlow, OpenAI Gym, LLM APIs (Gemini, OpenAI), RLHF, RAG, Qdrant, LangSmith
MLOps / DevOps:	Docker, FastAPI, CI/CD, Pulumi IaC
Cloud & Data:	AWS (ECS, S3, RDS), PostgreSQL, Redis, OBIEE/OAC
Mobile:	SwiftUI, GRDB, XcodeGen, BLE, LiveKit

Service

Program Committee: Human-aligned RL Workshop at ICRA 2024
Reviewer: TMLR, MRS 2023, AAMAS 2024