

Wenxuan Guo

☎ (908) 842-4659

✉ w9guo@ucsd.edu

🌐 linkedin.com/in/wenxuan-guo-a4ba72265

🐙 github.com/wendyguo2002

Education

University of California, San Diego
M.S. in Computer Science

Sep 2025 – Present | San Diego, CA

University of Michigan
B.S. in Computer Science & Economics, GPA: 3.71 / 4.00

Aug 2022 – May 2025 | Ann Arbor, MI

Experience

University of California, San Diego – ML Research Assistant Synthetic Animal Motion Generation & Evaluation for Radar Simulation *Jan 2026 – Present*

- Built an end-to-end pipeline to generate realistic animal motion for **synthetic mmWave radar training data**, addressing limited real-world motion datasets.
- Developed a **video-to-3D motion reconstruction pipeline** using segmentation, skeletonization, 2D keypoint detection, and **SMAL/SMALify** fitting to recover 3D quadruped pose sequences.
- Implemented motion processing with **Python, PyTorch, and OpenCV**, representing motion as $T \times J \times 3$ **joint trajectories** for dataset construction and analysis.
- Integrated reconstructed motion into a **Unity motion matching animation system** to generate controllable quadruped locomotion and export animation clips for radar simulation.
- Designed a **Fréchet Distance-based evaluation pipeline** comparing Unity-generated and real dog motion using normalized joint position and velocity features.

University of Michigan – HoloLens Radiation Visualization Project

May 2025 – Aug 2025

- Developed a **feed-forward neural network** for real-time classification of gamma-ray waveforms, engineering 15+ temporal features (rise time, peak ratios, pulse width) for detector signal analysis.
- Implemented Bayesian uncertainty calibration, achieving 91% precision in signal-noise discrimination versus 76% baseline threshold methods, yielding an 18% SNR improvement.
- Optimized **CNN** design from five to three convolutional layers and compiled critical inference loops with **Cython**, meeting sub-millisecond latency for continuous data streams.
- Integrated **Python** inference with Unity through TCP, deploying to HoloLens 2 with 3D overlays; boosted valid-signal detection from 89% to 94% and cut operator hotspot detection time by 40%.

University of Michigan – Discrete Math PrairieLearn System

May 2024 – Jan 2025

- Built and merged an ordering-feedback engine for Proof Blocks into the **PrairieLearn** platform, providing deterministic tag-based feedback for 800+ students through a Python backend and Mustache.js frontend.
- Deployed and configured the **PrairieLearn system** locally and in production, proposing modular APIs to support future LLM-based feedback extensions and improve long-term maintainability.
- Collaborated with core developers on refactoring architecture and testing; streamlined feedback workflows that improved structured learning outcomes and reduced manual grading overhead.
- Led end-to-end design, testing, and rollout; authored written and video onboarding guides, enabling full departmental adoption for Discrete Math and related proof-based courses.

Digital China – Software & Machine Learning Engineering Intern

Jun 2023 – Aug 2023

- Developed a baseline-management web platform using **Java (Spring Boot)** and **Vue.js**, automating document version control and multi-team workflow approvals across 200+ daily users.
- Designed relational schemas and data-access layers in **MySQL + MyBatis Plus**, supporting transactional integrity and efficient multi-threaded CRUD operations under concurrent workloads.
- Integrated a **Python-based document classification module** with TF-IDF and logistic regression models, enabling automatic tagging and search recommendation within the platform.
- Built scheduled retraining pipelines using **Airflow** and monitored inference accuracy and latency metrics, improving keyword classification precision from 82% to 91% and cutting manual labeling effort by 45%.

Independent Projects

Synthetic Data Distillation with Nemotron-4

Dec 2025 – Jan 2026

- Designed a **Synthetic Data Generation (SDG)** pipeline using **Nemotron-4 340B** to generate high-fidelity financial instruction-response pairs, addressing label scarcity in domain-specific NLP.
- Architected a **teacher-student distillation** framework to compress Nemotron-level reasoning into a **sub-5B Small Language Model**, enabling accurate ticker extraction and sentiment inference from noisy news.
- Eliminated hallucination in financial outputs by enforcing strict **JSON schema constraints**, achieving deterministic structured predictions suitable for downstream trading and analytics systems.

AI Portfolio Knowledge Base via RAG

Dec 2025 - Jan 2026

- Architected a **Retrieval-Augmented Generation (RAG)** system using **LlamaIndex** to transform static professional documents into an interactive, queryable knowledge base.
- Designed a semantic ingestion pipeline for structured Markdown, applying **task-aware chunking** to preserve context for complex behavioral and technical queries.
- Optimized retrieval quality with **ChromaDB** embeddings and prompt tuning, enabling low-latency, structured response synthesis aligned with interview evaluation frameworks.

Technical Skills

Programming: Python, C++, Java, JavaScript, SQL, C#

Machine Learning: PyTorch, TensorFlow, Scikit-learn, Transformers, Hugging Face

Tools: Unity, HoloLens, MRTK, AWS, PrairieLearn, Git, Jira