Yiming Feng

909-348-4612 | yimingoscarfeng@gmail.com | linkedin.com/in/yiming-feng | github.com/quiet98k | quiet98k.com

Education

Duke University — Durham, NC
Computer Science — Master of Science — *August 2025 – Present*
GPA: 4.0

University of California, Davis — Davis, CA
Computer Science — Bachelor of Science — *July 2022 – June 2025*
GPA: 3.9
Honors Graduation

Experience

AI/ML Intern — May 2026 – Present
Qorvo — *Greensboro, NC*

- Designed and developed an agentic GenAI workflow (As-Is Documentation Agent) that analyzes SAP source code and technical artifacts to automatically generate, summarize, and iteratively refine solution design documents and process documentation.
- Built RECEFW Agent, a Retrieval-Augmented Generation (RAG) chatbot integrating multiple enterprise data sources, enabling context-aware question answering and knowledge retrieval for internal users.

Graduate Teaching Assistant — September 2025 – Present
Duke University CS Department — *Durham, NC*

- Evaluate assignments promptly and provide detailed, actionable feedback that strengthens student understanding and performance.
- Hold weekly office hours to answer questions and debug coding issues, improving comprehension of core course material.
- Lead pre-exam review sessions with targeted practice problems to reinforce key concepts and improve exam readiness.

System Architect Intern — July 2023 – August 2023
Guangdong Yixun Technology — *Guangzhou, China*

- Integrated Vue.js with Spring Boot to streamline the user experience and improve system responsiveness.
- Raised YOLOv3 image recognition accuracy by developing and curating specialized training datasets.
- Built monitoring pipelines with Prometheus and Grafana to surface system health and performance metrics.

Projects

As-Is Documentation Agent | *Agentic AI, LLMs, Databricks, Openai SDK* — May 2026 – Present

- Built an agentic GenAI workflow that analyzes SAP source code, configuration, and technical artifacts to automatically generate, refine, and validate solution design and process documentation.
- Designed a multi-stage LLM orchestration pipeline with structured outputs, prompt engineering, iterative validation, and traceability generation to improve documentation accuracy and consistency.

RECEFW Agent | *RAG, LLMs, Vector Search, Databricks* — May 2026 – Present

- Developed an enterprise Retrieval-Augmented Generation (RAG) chatbot integrating multiple internal knowledge sources for context-aware question answering and technical knowledge retrieval.
- Implemented document ingestion, embedding, indexing, and semantic retrieval pipelines using vector databases and hybrid search techniques to improve retrieval quality and reduce hallucinations.

StyleDou | *Reinforcement Learning, Opponent Modeling, Deep Monte Carlo* — January 2026 – May 2026

- Built a style-conditioned DouDizhu RL agent on top of DouZero/DouZero+, adding opponent-style LSTM encoders and auxiliary next-action prediction for imperfect-information gameplay.
- Trained distributed self-play agents for $\sim 10^9$ frames on a 3–4 GPU setup; achieved $\sim 98\%$ win rate vs. Random, $\sim 83\%$ vs. RLCard, and improved average win rate by ~ 4 percentage points over a matched-compute baseline.
- GitHub: github.com/quiet98k/StyleDou

SafeBench-DRA | *LLM Agents, Deep Research Agent, Benchmarking, Agent Security* January 2026 – May 2026

- Built an offline Chinese–English benchmark for evaluating security failures in deep-research agents using deterministic retrieval, adversarial tool-response overlays, and trace-level scoring.
- Created a 37-task bilingual benchmark for evaluating prompt injection, goal drift, and taint attacks in deep-research agents.
- GitHub: github.com/Orange23334/ECE590_Genai_Safebench_DRA

Hand Gestures Classifier | *GNN, ResNet-34, Image Classification, PyTorch* July 2025 – December 2025

- Developed a multimodal deep learning system for classifying 18 hand gestures using Graph Neural Networks, ResNet-34, and late fusion.
- Achieved 99.72% accuracy on the HaGRID dataset.
- GitHub: github.com/quiet98k/Hand-Gestures-Classifier

LOL-Win-Prediction | *Adversarial Robustness, PGD, TRADES, Deep Learning* April 2025 – June 2025

- Predicted win probability in League of Legends matches using a deep neural network trained on the first 15 minutes of in-game data.
- Improved local robustness with PGD attacks, noise injection, Jacobian regularization, and TRADES.
- GitHub: github.com/quiet98k/LOL-Win-Prediction

Proactive-AI | *Chrome Extension, AI Tools, OCR* July 2025 – December 2025

- Built a Chrome extension that detects user hovers and selections to surface the most relevant AI tools in place.
- Enabled real-time code explanation, article summarization, equation graphing, table extraction, and OCR on images without leaving the webpage.
- Improved user productivity by integrating multi-functional AI features directly into the browsing workflow.
- GitHub: github.com/Arvid-pku/Proactive-AI

Chinese Handwritten Digit Classifier | *CNN, VGG16, LeNet* September 2024 – December 2024

- Developed deep learning models with VGG16 and LeNet to classify handwritten Chinese numerals.
- Achieved approximately 97% accuracy on the Chinese MNIST dataset.
- GitHub: github.com/quiet98k/Chinese-Digit-Classifier

Blood-On-Mahjong | *TypeScript, WebSocket, Google OAuth, Kubernetes, CI-CD* July 2025 – December 2025

- Built a real-time multiplayer Sichuan Mahjong web application with WebSocket-based gameplay and Google OAuth authentication.
- Designed deployment for scalable Kubernetes-based hosting.
- GitHub: github.com/quiet98k/Blood-On-Mahjong

WOAA-trading | *Full-Stack, WebSocket, Backtesting* May 2025 – December 2025

- Developed a full-stack stock trading simulation platform with real-time market data streaming and historical backtesting.
- Implemented portfolio management features that let users practice and refine trading strategies without financial risk.
- GitHub: github.com/quiet98k/woaa-trading

Technical Skills

Programming Languages: Python, C++, C, JavaScript / TypeScript, Java
Machine Learning & Deep Learning: PyTorch, Scikit-learn, TensorFlow, Supervised & Unsupervised Learning, Regression, Tree-based Models, Neural Networks, CNNs, Transfer Learning, Model Evaluation & Hyperparameter Tuning
Computer Vision: Image Classification, Object Detection (YOLO), Data Augmentation, OpenCV, MediaPipe
Generative AI: Large Language Models (LLMs), Agentic AI, Deep Research Agent, Prompt Engineering, Inference Pipelines, Agent Security
Reinforcement Learning: Deep Reinforcement Learning, Multi-Agent RL, Deep Monte Carlo (DMC), Self-Play Training, Policy Evaluation
MLOps & Systems: Docker, Kubernetes, CI/CD (GitHub Actions), Model Deployment, Linux, Git
Software Development: React, Nuxt, Node.js, FastAPI, REST, WebSocket, MongoDB, SQL, Redis