Nguyen Van Quang

AI Engineer

Email: qcontact.12@gmail.com
Mobile: (+84) 888 987 400
Linkedin: linkedin.com/in/quaang
Github: github.com/quangliz

Education

- **Hanoi University of Civil Engineering** — **Ha Noi, Viet Nam**
  Major: **BSc in Computer Science** | **CPA: 3.21** — Sep 2022 - Sep 2026 (Expected)
  - Awarded Scholarships for Good Academic Achievements — 2023, 2024

Skills

- Programming Languages: **Python, TypeScript, SQL**
- ML/AI: **Regression/Boosting, Neural Networks, Transformer, LLMs, AI Agents, Fine-tuning**
- Database: **PostgreSQL, Vector databases**
- Data Engineering: **ETL, Orchestration, Feature Engineering, Data Preprocessing**
- Tools: **Git, Docker, Linux, AWS, Coding Agents**
- Languages: **Vietnamese** (native), **English** (B2)

Experience

- **AI Engineer Intern** | VinSmart Future — May 2026 - Present
  - Participating in the process of building a high-quality data pipeline for the AI system in the tourism domain.
- **AI Engineer Intern** | 8SENECA — Nov 2025 - Mar 2026
  [Python, FastAPI, Langchain/Langgraph, OpenAI, HuggingFace, Docker, Postgres, Neo4j, ...]
  - Developed a robust AI-powered system for automating the process of creating and managing software requirements (SDLC) on Sparx Enterprise Architecture platform.
  - Implemented a custom LLM inference pipeline based on client's requirement.
  - Fine-tuned BERT model for classification and embedding model for text vectorization within domain.
  - Participated in the deployment process.

Projects

- **ThePawsome - AI-native Pet E-commerce Platform** | https://github.com/quangliz/petshop
  [Next.js, FastAPI, PostgreSQL + pgvector, LangGraph, OpenAI API, Docker, AWS EC2/RDS]
  - Built a full-stack Vietnamese pet e-commerce platform with product discovery, checkout, payment, order management, customer support, forum, and role-based admin workflows.
  - Designed and implemented Catbot, an agentic RAG assistant with planner-enforced tool calling, hybrid retrieval, pet-care knowledge grounding, pet-profile personalization, deterministic product-safety checks, streamed responses, verified product-card rendering, and human handoff.
  - Created an AI evaluation pipeline with 50 LangSmith-traced cases using LLM-as-judge scoring, measuring answer quality, required-tool usage, tool order, safety preflight, valid product references, handoff behavior, and latency; achieved 96% pass rate, 97.0 average final score, 0% eval error rate, 100% required-tool/tool-order/safety/handoff accuracy, and 11.16s p50 / 22s p99 agent latency.

Certificates/Certifications

- DeepLearning.ai - **Machine Learning Specialization Certificate** — Jun 2025
- British Council - **IELTS 6.0 Overall** — Jul 2025