# Nam Hoang

(334) 492-4343 | namhd.work@gmail.com | linkedin.com/in/nam-hd | github.com/nam-ruto

## Education

**Troy University** Troy, Alabama
*Bachelor of Computer Science* *(Expected) May 2026*

- **GPA**: 3.60/4.0
- **Chancellor's List**: Recognized for academic excellence with GPA score 4.0/4.0 (3 semesters)
- **First Place at Troy Hackathon 2024**: Developed an AI Chatbot system during a 24-hour Hackathon
- **Second Place at Troy Hackathon 2025**: Developed an anomaly behavior detection system
- **Troy University International Trojan Scholarship**: Merit-based scholarship for top 5% of incoming students

## Technical Skills

**Languages**: Python, R, C/C++, Java, JavaScript, TypeScript, Latex, HTML/CSS
**Database**: PostgreSQL, MySQL, SQL Server
**Clouds**: Microsoft Azure, Google Cloud Platform
**LLM Platforms**: Hugging Face, Ollama, WebLLM
**AI Frameworks**: PyTorch, MXNet, TensorFlow, Keras
**Web Frameworks**: Flask, FastAPI, Nuxt, React, React Native, Streamlit

## Working Experience

**Software Engineer Intern** June 2024 - Present
*The George Washington Institute at ISODS* *Remote*

- **Technologies**: *RAG, Python, JavaScript, Nuxt, FastAPI, LangChain, Ollama, Web-LLM, ChromaDB*
- Built an end-to-end document processing pipeline to clean, chunk, and embed **11,700+** medical documents across 6 clinical domains for AI-driven applications
- Developed a full-stack web platform **(Nuxt + FastAPI)** enabling users to interact with domain-specific AI agents (Law, Medical, etc.) through a unified interface
- Integrated local LLM inference using Ollama and WebLLM to enable fully offline, privacy-preserving inference, eliminating external API dependency and cost
- Collaborated with cross-functional teams to define data schemas, retrieval workflows, and system boundaries for scalable AI services

**Software Engineer Intern (Mobile)** August 2025 - December 2025
*Troy University* *Troy, AL*

- **Technologies**: *JavaScript, TypeScript, React Native, PostgreSQL, Supabase, Stripe*
- Built a full-stack **React Native** mobile application for campus transportation, serving **2,000+** users and reducing manual booking workflows by 70%
- Designed and implemented scalable backend data models and APIs using PostgreSQL and Supabase to manage bookings, routes, users, and payments
- Integrated **Stripe payments** to support secure, real-time transactions with end-to-end booking consistency and zero manual reconciliation
- Demonstration: **Trojan Bus: Bus-Ticket Booking System**

**Data Engineer Intern** May 2025 - August 2025
*Cerrowire* *Hartselle, AL*

- **Technologies**: ETL, Python, SQL Server, Azure VM, Window Task Scheduler
- Designed and implemented an **ETL pipeline** to ingest daily copper settlement prices from the commodity market and load the data into SQL Server for company-wide reporting and analysis
- Automated pipeline execution on **Azure VM** using Windows Task Scheduler, ensuring reliable, hands-free data updates and reducing manual intervention by **100%**

**Data Engineer Intern** May 2024 - August 2024
*FPT Software* *Vietnam*

- **Technologies**: ETL, Python, C#, SQL Server, MySQL, PostgreSQL, Azure, Blob Storage
- Developed a Python framework capable of retrieving data from multiple database systems and cloud platforms, enhancing data accessibility and integration
- Implemented optimized database queries and efficient data handling techniques, achieving a **30%** reduction in ETL execution time and significantly improving system responsiveness

## Research Experience

**Research Intern - Developing a Multi-Agent System for Medical Queries** June 2024 - Present
*The George Washington Institute at ISODS* *Remote*

- Co-authored a paper on **DoctorAI**, a multi-agent RAG-based healthcare assistant for medical question answering under resource constraints
- Built and preprocessed a **11,700+** page medical dataset (6 domains: diabetes, hypertension, heart failure, etc.) using Docling + Vintern-1B-v2 OCR
- Developed semantic retrieval pipeline with Sentence Transformer embeddings and ChromaDB for efficient context retrieval, achieving **90%** context precision and reduced query latency by **30%**
- Designed and tested a **multi-agent RAG architecture** (Diagnostic, Doctor, Summarize Agents) with ReAct prompting, improving diagnostic reasoning and grounding responses in verified medical sources

## Conference Publications

- Cuong Do, Quy Quach, Dang Quach, **Nam Hoang**, Khuong Vu, Tin Huynh, Huy Do. **(2025). Efficient Multi-Agent Collaboration for Medical Question Answering with Quantized LLMs under Resource Constraints**. *ICITE 2025 - International Conference On Information Technology Education, Ho Chi Minh City Vietnam, November 2025*

## Projects

**Anomaly Behavior Detection System (Hackathon Project)** | *YOLO + CLIP + Gemini* November 2025

- Developed a real-time intelligent surveillance system that detects weapons and anomalous human behaviors using YOLOv8 and CLIP models
- Integrated Gemini (VLM) for contextual reasoning and automatic incident report generation from visual data
- Built a Streamlit dashboard for real-time video analysis, alert visualization, and threat-level summaries

**Trojan AI Chatbot (Hackathon Project)** | *Python, FastAPI, React, ChromaDB, RAG, LLMs* November 2024

- Developed an AI-powered chatbot using the Retrieval-Augmented Generation (RAG) technique to assist students with academic program inquiries, degree information, and campus resources
- Built a full-stack solution with a **React** frontend and **FastAPI** backend, integrating the **all-MiniLM-L6-v2** model for precise data retrieval and embedding

**ETL Process For Online Retail Data** | *Python, PostgreSQL, DBT, BigQuery, Airflow* January 2024

- Built an **ETL** process that extracts raw data from PostgreSQL, implements **DBT** tool for data cleaning, and loads processed data into Google BigQuery for scalable storage and analysis
- Orchestrated the entire ETL process using **Apache Airflow** DAG. Scheduled the pipeline to run daily, ensuring up-to-date and accurate data for analysis

**Restaurant Management System** | *Java, JavaFX, PostgreSQL, NetBeans, SceneBuilder* November 2023

- Developed a full-stack restaurant management application in Java, leveraging JavaFX for a responsive, desktop-based UI and SceneBuilder for rapid UI prototyping
- Architected a PostgreSQL relational database to support dynamic operations, including order processing, reservations, menu customization, billing, and inventory control
- Integrated business logic with a modular, object-oriented design to ensure maintainability and scalability for future feature extensions
- Optimized database queries and transactions to improve performance and ensure reliable data persistence under concurrent user operations
- Streamlined development workflow using NetBeans IDE and incorporated clean UI/UX principles to deliver an intuitive user experience for both staff and administrators