Bhaskara-Yashwant Bitra

Binghamton, NY | (813) 593-8899 | yashubitra@gmail.com | linkedin.com/in/yashwant-bitra | github.com/yashzord | Portfolio

Education

State University of New York at Binghamton — Binghamton, NY
Master of Science in Computer Science, Artificial Intelligence Track — May 2025
Bachelor of Science in Computer Science — May 2024
GPA: 3.82 / 4.00 | Dean's List: Fall 2022 – Spring 2024

Technical Skills

AI/ML & GenAI: Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), Agentic AI, LangChain, LangGraph, Model Context Protocol (MCP), Prompt Engineering, Hugging Face, Ollama, vLLM, Groq
Machine Learning: TensorFlow, Keras, PyTorch, Scikit-learn, XGBoost, Random Forest, SMOTE, Autoencoders, Feature Engineering, Cross-validation, Hyperparameter Tuning
MLOps & Observability: LLMOps, Langfuse, MLflow, DeepEval, promptfoo, Guardrails AI, CI/CD Pipelines, GitHub Actions, Experiment Tracking
Cloud & Infrastructure: AWS (Bedrock, SageMaker, EC2, S3, RDS), Docker, Docker Compose, FastAPI, Flask, REST APIs, Git, Linux/Unix
Security: Penetration Testing, Intrusion Detection Systems (IDS), OWASP Top 10 (Web & LLM), Zero Trust Architecture, DevSecOps, Kali Linux, Wireshark, Burp Suite, mitmproxy, aircrack-ng, LUKS Encryption, Network Anomaly Detection
Compliance: CMMC (L1/L2/L3), HIPAA, FedRAMP, HITRUST-CSF, DORA, CRI Profile, NIST
Frontend: React.js, Next.js, TypeScript, JavaScript, Tailwind CSS, D3.js, Chart.js, Streamlit
Databases: PostgreSQL, MongoDB, MySQL, SQLite, ChromaDB, AWS RDS, JSONB
Languages: Python, TypeScript, JavaScript, C++, C, SQL
Certifications: Google Data Analytics, Google Cybersecurity, Relational Database Systems

Professional Experience

State University of New York at Binghamton, School of Computing — Binghamton, NY
Teaching Assistant & Research Assistant — August 2024 – Present

- Mentored 40+ students in React.js, TypeScript, MongoDB, REST API design, secure database integration, and ML pipeline development across 10+ capstone projects
- Automated TypeScript project grading via custom GradeScope scripts, cutting review time by 15 min/submission across 40+ students
- Researched LLM architectures, RAG systems, and agentic AI frameworks; applied findings to independent production AI system development

CoreCard Software, Inc. — Atlanta, GA
Software Developer / Machine Learning Engineer Intern — May 2023 – August 2023

- Shipped a real-time fraud detection REST API (XGBoost, Random Forest, SMOTE) achieving 85% accuracy and cutting fraud 25% across 1M+ transactions in production; applied anomaly detection to live threat mitigation
- Applied SMOTE oversampling to resolve severe class imbalance; performed feature engineering on raw transaction data improving minority-class recall significantly
- Delivered a 3D interactive analytics dashboard (D3.js + Plotly + Flask) that saved 10 min/session per client and directly informed a $500K+ credit card strategy decision
- Partnered with analysts and business stakeholders to define requirements, validate model outputs, and iterate across delivery cycles

Project Experience

ST Intelligent Assistant: Agentic AI for Enterprise MFT — Jan 2026 – Present
Python, FastAPI, LangGraph, MCP, Groq, Ollama, vLLM, ChromaDB, Langfuse, DeepEval, promptfoo, Streamlit

- Built and benchmarked dual agent architectures (custom ReAct + LangGraph) for production agentic AI, evaluating tool selection accuracy, latency, and hallucination rates
- Implemented RAG-based semantic tool routing (ChromaDB + nomic-embed-text embeddings) dynamically selecting from 49 business tools, zero manual configuration, measurable recall improvement over keyword routing
- Architected a model-agnostic inference layer enabling A/B testing across backends (Groq, Ollama, vLLM); wired Langfuse for full LLM observability (traces, latency, cost)
- Built CI/CD evaluation pipelines (promptfoo + DeepEval) tracking tool accuracy and hallucination rates across model updates

GRC Compliance Data Ingestion Platform Jan 2026 – Present

Python, PostgreSQL, AWS RDS, JSONB, psycopg3, uv, Docker, REST APIs

- Engineered ETL pipelines for 6 regulatory compliance frameworks (CMMC L1/L2/L3, HIPAA, FedRAMP, HITRUST-CSF, DORA, CRI Profile) into a unified AWS RDS PostgreSQL production database
- Normalized heterogeneous data sources (Excel, PDF, HTML, live REST APIs via eCFR) into a single JSONB schema supporting hierarchical compliance trees up to 4 levels deep
- Designed content-aware versioning, auto-increments version only on data change, ensuring idempotent pipeline runs and preventing duplicate writes
- Maintained multi-tenant data isolation across 20+ frameworks via source-scoped queries; collaborated with team scaling the platform to additional compliance standards

Mobile Network Intrusion Detection System Jan 2025 – May 2025

Python, mitmproxy, TensorFlow/Keras, Kali Linux, LUKS, Streamlit, PostgreSQL

- Intercepted and processed 25,000+ HTTP/HTTPS flows from a live iOS device via mitmproxy; stored in LUKS-encrypted volumes for data-at-rest security compliance
- Trained a deep autoencoder on 21,993 labeled flows for unsupervised anomaly detection, reduced false positives from 1,100 to 103 (90% reduction) at the 95th percentile threshold
- Shipped a Streamlit real-time traffic monitoring dashboard replacing manual log triage with visual flow investigation

Evil Twin Attack: Full Wireless Attack & Defense Framework 2024

Kali Linux, aircrack-ng, dnsmasq, Wireshark, Python (pywifi, tkinter), hostapd

- Executed complete penetration test: monitor mode, deauth via aireplay-ng, rogue AP with airbase-ng impersonating target SSID, captive portal credential harvesting, confirmed via live testing
- Built real-time BSSID anomaly detection script flagging rogue APs by MAC mismatch; wrote pywifi prevention script to auto-disconnect from untrusted networks
- Deployed honeypot APs for attacker fingerprinting; documented full red-team and blue-team lifecycle across five technical reports

Social Media Intelligence Platform: Toxicity Classification Aug 2024 – Dec 2024

Python, MongoDB, Flask, Hugging Face, LangChain, Scikit-learn

- Ingested 100K+ posts across Reddit, 4chan, and YouTube into a structured MongoDB document store; built normalized schema for cross-platform content analysis
- Integrated Hugging Face NLP toxicity classifiers achieving 90% accuracy; built RAG pipeline (LangChain + MongoDB) enabling natural language querying over 100K+ documents

AI Assistant: OWASP Top 10 LLM Security Hardening Jan 2026 – Present

Python, FastAPI, LangGraph, Guardrails AI, Langfuse

- Hardened a production AI assistant against OWASP Top 10 for LLM Applications: prompt injection, insecure output handling, sensitive data exposure, model denial-of-service
- Implemented Guardrails AI output validation and PII scrubbing; integrated Langfuse for full security audit trail observability

Leadership

Theta Tau Engineering Fraternity Binghamton, NY

Vice President December 2023 – May 2024

- Led organizational planning, facilitated leadership meetings, tracked action items, and managed competing stakeholder priorities across concurrent chapter initiatives