

Yosef Shammout

(734)-709-4358 | yosefshammout123@gmail.com | linkedin.com/in/yosef-shammout | github.com/ysham123

Software Engineer — AI Research & High-Performance Systems

Specializing in GNNs, Knowledge Graph Learning, and resilient Python infrastructure. Expertise in bridging research-grade ML with production-scale data engineering.

EDUCATION

Wayne State University

Detroit, MI

B.S. Computer Science; GPA 3.7/4.0

Expected Dec 2026

- **Relevant Coursework:** Software Engineering, Algorithm Design & Analysis, Operating Systems, Multivariable Calculus, Linear Algebra, Probability & Statistics, Physics II (E&M).

TECHNICAL SKILLS

Languages: Python (Expert), SQL, TypeScript/JavaScript, Cypher (Neo4j), C++

ML & Research: PyTorch (Expert), GNNs (RGCN), Knowledge Graphs, Selective Reranking, Model Calibration (ECE), Logit Adjustment, Scikit-learn, NumPy, Pandas, Grad-CAM

Systems & DevOps: Docker, FastAPI, Airflow, CI/CD (GitHub Actions), Linux (Bash), PostgreSQL, Snowflake, BigQuery, AWS, Vercel

RESEARCH EXPERIENCE

Wayne State University (Advisor: Dr. Lihui Liu)

Detroit, MI

AI Research Assistant — Knowledge Graph Learning

Oct 2025 – Present

- Developed the **CSR-HARD** pipeline using a Relational Graph Convolutional Network (**RGCN**) to generate graph-aware logits for long-tail medication prediction, managing patient-disjoint evaluation splits.
- Engineered a **Selective Rerank** module to fuse graph logits with DXPOP popularity priors, optimizing accuracy for rare queries via a lightweight gating mechanism; achieved MRR of **0.211 ± 0.030**.
- Improved model calibration via **logit adjustment**, reducing Expected Calibration Error (ECE@top1) from **0.0997 to 0.0708** while conducting seed sweeps and overlap/OOV audits.

INDUSTRY EXPERIENCE

Ford Motor Company

Dearborn, MI

Software Engineering Intern — Enterprise Technology

May 2025 – Jul 2025

- Built 4 production Airflow pipelines for automated data ingestion into Snowflake and BigQuery; utilized Docker for containerized, idempotent workflows to ensure high-fidelity datasets for downstream ML.
- Implemented Python-based **schema validation gates** in CI/CD (GitHub Actions), defining **25+** source-target contracts to block breaking changes and improve overall release health.

SamSearch (Government-Contracted AI Startup)

Vancouver, BC

Software Engineer Intern

May 2024 – Aug 2024

- Deployed a deep learning-assisted chatbot backend using **FastAPI**; achieved a **40% reduction in p95 latency** through asynchronous processing and optimized Astra DB (NoSQL) query patterns.
- Architected modular API endpoints to handle secure data retrieval for government-contracted requirements, improving system engagement by **30%** through optimized model performance metrics.

Code Emergency (Health-Tech Startup)

Detroit, MI

Founder & Lead Engineer

Aug 2024 – Feb 2025

- Architected and led end-to-end execution of a privacy-focused health platform (Django/React); utilized AWS and Vercel for a pilot that reduced median medical info retrieval time by **20%** in emergency scenarios.
- Developed and secured a QR-based data platform, focusing on HIPAA-compliant infrastructure and high-availability design for critical healthcare data access.

SELECTED PROJECTS & AWARDS

Myndra: Developed a multi-agent orchestration framework using LLM planning and graph memory; implemented a Dockerized evaluation workflow with rigorous ablation testing.

VisionAssist: Engineered a real-time CV application (MobileNetV2 + LSTM) with Grad-CAM visualization; optimized for low-latency feedback using Python-based WebRTC and TTS integration.

Michigan DevFest AI Hackathon (IBM Watsonx, Google): Finalist and Invited Presenter (2025).