Shishir Lohar

U.S. Citizen | (484)-784-8419 | shishirlohar@gmail.com | linkedin.com/in/shishirlohar | github.com/winterblacksmith

Education

University of Georgia, Franklin College of Arts and Sciences December 2026
B.S. in Computer Science, Focus in Artificial Intelligence; GPA: 3.5/4.0 Athens, GA

Relevant Coursework:
Software Engineering, Data Structures, Algorithms, Systems-Level Programming, Web Programming, Robotics

Technical Skills

Languages: Python, Java, C/C++, JavaScript, TypeScript, SQL, HTML/CSS
Frameworks/Libraries: React, Next.js, Node.js, FastAPI, Flask, Spring Boot, Streamlit, scikit-learn
Databases: PostgreSQL, SQLite
Tools/Cloud: Git, Docker, Railway, Vercel, AWS, GCP, Github Actions, Kafka, Postman

Experience

Lead Software Engineer | Kappa Theta Pi | Athens, GA September 2026 – Present
- Leading a 3-person engineering team building a distributed systems incident monitoring platform that ingests GitHub deployment events and application telemetry through REST APIs and AWS SQS, persists incident data in PostgreSQL, and visualizes service health through a React dashboard
- Developing KTPI, an agentic AI incident investigator that uses MCP tools to correlate recent deployments, error logs, and service health, producing evidence-backed root-cause hypotheses and recommended remediation steps
- Designing for Docker services, automated GitHub Actions deployments to AWS, reliable cloud operation with event processing, retry and duplicate-handling logic, unit and integration testing, and peer code reviews

Software Engineer Intern | UGA Warnell School of Forestry | Athens, GA April 2026 – Present
- Engineered an AI assistant using Python and local LLMs to convert natural-language questions and prompts into validated SQL database queries and interactive map visualizations across 100+ test prompts
- Processed 500K+ rows across large forestry datasets using Pandas, combining data validation, filtering, and the AI model Qwen2.5's output to create visualizations using Streamlit in under 30 seconds for workflows that could take up to 1 hour to complete manually using traditional data tools
- Evaluated 4 Llama and Qwen models and deployed Docker containers on Railway to run the AI assistant, integrating persistent storage, deterministic fallback logic, and geospatial visualization with Leaflet/Folium

Founder & Software Engineer | Spotr | Duluth, GA May – August 2026
- Secured $7500+ in funding to build an AI-powered vehicle sourcing platform for dealerships, using 3+ automotive data APIs to search tens of thousands of listings across hundreds of dealerships and online marketplaces
- Developed API-driven workflows to retrieve, process, filter, and rank vehicle inventory across 10+ dealership-defined criteria, consolidating multiple external data sources into a unified search experience
- Designed the platform to replace per-vehicle auction sourcing fees with a $500/month subscription model, targeting $8.5K–$11.5K in monthly savings for dealerships sourcing 30–40 vehicles per month

Projects

Software Engineer (Forage Experience) | JPMorgan Chase & Co. | Remote February – March 2026
- Engineered a Spring Boot service to receive, validate, and record high-volume financial transactions
- Integrated Kafka-based transaction processing with a SQL database for persistent, validated data storage and reliable downstream access
- Developed a REST API to expose processed financial data, enabling reliable downstream access and communication with external service consumers

Machine-Learning Health Model | **GitHub Repo** August – September 2025
Python, scikit-learn, Pandas, Streamlit, Random Forest
- Engineered an end-to-end health trend prediction pipeline across 1,000 synthetic health records and 8 biometric and lifestyle features, including sleep, activity, caloric intake, blood pressure, and resting heart rate
- Trained and evaluated a 100-tree Random Forest classifier with scikit-learn and Pandas using an 80/20 train-test split, achieving 98% classification accuracy across improving, stable, and declining health-trend categories
- Serialized the trained model with Joblib and integrated real-time inference into an interactive Streamlit web application, enabling users to enter daily health metrics and receive immediate personalized trend predictions