

Harish Sakthivel

AI Systems | Full-Stack Development | LLM Orchestration

US Citizen | hxrish04@gmail.com | (818) 601-0465 | [linkedin.com/in/harishsakthivel04](https://www.linkedin.com/in/harishsakthivel04) | github.com/hxrish04

Information Sciences student who has shipped three production AI systems to cloud, including an evidence-grounded decision engine on AWS ECS with explicit uncertainty handling and bounded agent evaluation.

EDUCATION

University of Tennessee, Knoxville

Knoxville, Tennessee

B.S. in Information Sciences (AI & Data Analytics Concentration)

Expected May 2026

- Major GPA: 3.65/4.00 | Dean's List
- Relevant Coursework: NLP, Artificial Intelligence, Machine Learning, Database Systems

SKILLS

Languages: Python, TypeScript, JavaScript, SQL, Java

AI Systems: LLM orchestration, retrieval-grounded evaluation, structured output validation, uncertainty routing, deterministic fallbacks, bounded agent loops

Backend: FastAPI, Node.js, Express, React, Next.js, REST APIs

Infra: Docker, AWS ECS/Fargate, Azure, Playwright, Prisma, Supabase, Git

PROJECTS

Evident | Evidence-Grounded AI Decision System

Feb 2026 – Apr 2026

- Built and deployed a decision system that ranks research outreach targets with cited public evidence and explicit “insufficient evidence” refusals when support is weak.
- Designed a hybrid pipeline (deterministic extraction/enrichment/prefilter + selective Claude evaluation) to reduce unnecessary model calls and keep early stages fully explainable.
- Implemented a bounded agent loop from evaluation to adaptive retrieval to second-pass review to final decision, capped at one adaptive retrieval per run.
- Productionized on AWS ECS/Fargate with Docker and Secrets Manager, using demo/local modes with per-run caps to keep cost and blast radius controlled.

StackGate | AI-Assisted Internal Developer Platform

Sep 2025 – Dec 2025

- Built a governed PostgreSQL provisioning platform (Next.js 15, Prisma, Claude API) that turns plain-English infra requests into validated, policy-routed tickets.
- Designed a risk and policy engine classifying requests into low/medium/high tiers, auto-approving safe requests and routing others through a role-aware multi-stage approval chain.
- Architected an adapter-based provisioning system enabling clean live-provider swap-in without core logic changes, with a guarded Azure PostgreSQL path for low-risk dev requests.
- Built a complete audit trail capturing every workflow event from submission through provisioning with full requester-facing transparency on routing decisions.

VeriFlow | AI-Powered QA Automation Copilot

Jun 2025 – Aug 2025

- Built an AI-powered QA tool that converts user stories into executable Playwright tests, closing the gap between written requirements and automated coverage.
- Delivered a UI for test generation, browser execution, scheduled reruns, and persisted run history.
- Integrated Claude with Playwright to validate user workflows across live web apps, improving automated coverage.
- Integrated Azure DevOps APIs to pull user stories and work items as direct input to automated test generation.

TEACHING

UTK College of Emerging and Collaborative Studies

Knoxville, Tennessee

Peer Mentor | AI 102: Natural Language Processing

Jan 2026 – May 2026

- Mentored 20+ students through NLP concepts including preprocessing, model training, and evaluation across weekly lab sessions.
- Supported debugging of transformer-based projects covering tokenization, embeddings, and sequence modeling.