Siddarth Akalankam

Troy, MI | 248-825-0990 | Sidaka.business@gmail.com | linkedin.com/in/siddarth-akalankam | github.com/sidakalankam

Education

Michigan State University — East Lansing, MI
Bachelor of Science in Computer Science, Minor in Business — *Aug. 2022 – May 2026*
GPA: 3.76

Experience

Amazon — East Lansing, MI
Capstone Software Engineer — *Jan. 2026 – Apr. 2026*

- Architected a unified observability platform automating anomaly detection across **AWS services** deployed in multiple global regions, saving **2+ hours per day** of manual monitoring for engineering teams.
- Designed a serverless **ETL pipeline** using **EventBridge** and **Lambda** to extract, transform, and load scheduled service metrics into **DynamoDB** for downstream processing.
- Enabled multi-modal detection by integrating **SageMaker** for time-series metrics and **Rekognition Custom Labels** for screenshot analysis, automating severity and priority ranking across both data sources.
- Built an anomaly triage system with direct **CloudWatch** deep links to anomaly timestamps, reducing investigation scope from **10,000+** logs to roughly **50** per incident.
- Enabled secure multi-account observability by implementing **IAM role assumption** for cross-account metric ingestion without credential sharing.

Altair — Troy, MI
Software Engineering Intern — *May 2025 – Aug. 2025*

- Shipped a production **REST API** on **Azure** that abstracted all infrastructure management for data scientists, enabling one-click **ML model** deployment across cloud and edge.
- Cut inference latency from **4s to 100ms** by building a **C++ TensorRT** pipeline for **NVIDIA Jetson** devices, migrating cloud inference from CPU to GPU, and resolving networking bottlenecks.
- Integrated the platform into Altair's existing microservice architecture by containerizing the API with **Docker** and deploying it to **Kubernetes** with custom YAML configurations.

RecycleMe — East Lansing, MI
ML Software Engineering Intern — *Jan. 2024 – Apr. 2024*

- Deployed a **PyTorch ResNet-101** model on **AWS SageMaker** achieving **92% accuracy** across 6 material categories, fully replacing manual classification.
- Reduced external API calls by shifting inference on-device within the **Flutter/Firebase** app, improving real-time response for end users.

Projects

FlashML.dev | *Next.js, FastAPI, ONNX Runtime, AWS ECS, AWS S3, PostgreSQL, RunPod Serverless*

- Built **FlashML**, a **self-serve ML deployment platform** for students and small teams to upload, host, and run inference on any model via a single zip file, abstracting GPU, compute, and storage management.
- Engineered automated **model versioning, validation, and API endpoint generation** at upload time, giving users an instant hosted model registry without infrastructure configuration.

FinForgeAI.app | *React, TypeScript, FastAPI, Plaid, Claude API, AWS Lambda, AWS SQS, PostgreSQL, Pinecone*

- Built **FinForge**, a personal finance AI agent that syncs live bank data via **Plaid**, analyzes spending patterns, and delivers scheduled briefings by email using **AWS Lambda** and **Claude API**.
- Engineered a hybrid retrieval system routing queries across **PostgreSQL** for exact transaction lookups and **Pinecone** for semantic memory, enabling users to ask follow-up questions by replying directly to briefing emails.

Technical Skills

Languages: Python, JavaScript/TypeScript, C/C++, SQL, Java
Cloud & Infra: AWS, Azure, GCP, Docker, Kubernetes, Terraform, CDK, CI/CD
Databases & Backend: PostgreSQL, DynamoDB, NoSQL, REST APIs, gRPC, Spring Boot, FastAPI
AI/ML & Developer Tools: PyTorch, ONNX, SageMaker, RAG, Cursor, Codex, Claude Code