

# Sampath Garimella

628-309-5036 | sampath.garimella.dev@gmail.com | linkedin.com/in/sampathgarimella | github.com/SampathGarimella

## SUMMARY

Full stack and AI Developer with 4+ years of experience building scalable cloud apps and AI frameworks. Skilled in Python, Swift, React, and RAG pipelines. Built internal automation tools adopted across engineering teams, delivered under tight schedules, and scaled to thousands of users.

## EDUCATION

<b>University at Buffalo – MS, Engineering Science</b> , Buffalo, NY	Jan 2022 – May 2023
Graduate Teaching Assistant, <i>Machine Learning with Search Engines</i>	
Courses: Machine Learning, Deep Learning, Advanced NLP, Search Engines, ML for Large Datasets, Multimodal ML, Computer Vision	
<b>KL University - BTech, Computer Science</b> , India	May 2017 – Apr 2021

## TECHNICAL SKILLS

<b>Programming Languages:</b>	Python, JavaScript, TypeScript, Swift, Java, Go, C++, C#, .NET Core
<b>Frontend:</b>	React, Next.js, Vue.js, Angular, React Native, HTML5, CSS3, Tailwind CSS, Bootstrap
<b>Backend &amp; Architecture:</b>	Node.js, Django, Flask, Spring Boot, GraphQL, REST, gRPC, WebSocket, Microservices, Event-Driven Architecture
<b>AI/ML:</b>	LLMs, RAG, Transformers, NLP, Computer Vision, Hugging Face, LangChain, PyTorch, TensorFlow, scikit-learn, Pandas, NumPy
<b>Data &amp; Search:</b>	Elasticsearch, Pinecone, FAISS, Chroma, PostgreSQL, MySQL, MongoDB, DynamoDB, Redis, Cosmos DB, Kafka
<b>Cloud &amp; MLOps:</b>	AWS (SageMaker, Bedrock, EC2, RDS, S3, Lambda), Azure (AI Studio, Service Bus, Event Hub), GCP (BigQuery), Docker, Kubernetes, Terraform, MLflow, Jenkins, GitHub Actions
<b>Security:</b>	RBAC, OAuth 2.0, JWT, IAM hygiene, TLS, audit logging, OWASP Top 10 awareness
<b>Observability &amp; QA:</b>	Datadog, Grafana, New Relic, Splunk, SonarQube, Postman, Cypress, Selenium, JUnit, Jest
<b>AI-Assisted Development:</b>	Copilot, ChatGPT, Claude Code, Cursor for testing, refactoring, and documentation support

## EXPERIENCE

<b>Mastery Logistics Systems</b> - Atlanta, GA	Oct 2024 – Present
<b>Full Stack Developer (Generative AI)</b>	
<ul style="list-style-type: none"><li>Design and deploy ML-driven route optimization solutions, increasing operational efficiency by 25% and saving \$100K annually</li><li>Build and maintain LLM-powered logistics reporting tools with LangChain, Pinecone, and FAISS, enabling context-aware search and summarization</li><li>Develop and optimize RAG-based recommendation modules for shipment planning, reducing decision time by 60%</li><li>Architect and secure GraphQL/REST APIs serving 500K+ daily requests with 99.99% uptime, supporting scalable operations</li><li>Deploy and monitor AI workflows on AWS Bedrock and SageMaker, optimizing performance and cost efficiency</li><li>Delivered compliance-critical tools on tight deadlines, adopted by multiple logistics teams, driving measurable impact.</li><li>Developed internal AI tooling with Python, React, and Swift to streamline logistics workflows and improve user adoption</li><li>Used code assistants (Copilot, ChatGPT/Claude Code, Cursor) to speed up implementation and test coverage while keeping code reviews strict</li><li>Containerized services with Docker/Kubernetes, built CI/CD pipelines in GitHub Actions, and monitored with Datadog, Grafana, and Splunk.</li></ul>	
<b>UPMC</b> - Pittsburgh, PA	Dec 2023 – Sep 2024
<b>Full Stack Developer</b>	
<ul style="list-style-type: none"><li>Built HIPAA-compliant data platforms with Django REST Framework for 90K+ patients</li><li>Developed React/Next.js reporting modules, improving provider satisfaction scores by 30%</li><li>Collaborated with clinicians and UX teams to design intuitive healthcare dashboards.</li><li>Integrated NLP pipelines to analyze clinical text, reducing manual review time by 50%</li><li>Worked across engineering and clinical teams to ensure successful adoption of tools by thousands of healthcare providers</li><li>Added role-based access, audit trails, and TLS/IAM-secure data handling patterns to protect patient data</li></ul>	
<b>Cigna</b> - Overland Park, KS	Jan 2022 – Nov 2023
<b>Full Stack Developer</b>	
<ul style="list-style-type: none"><li>Developed Kafka-based analytics pipelines and integrated DynamoDB for sub-50 ms query latency</li><li>Developed healthcare features with React Native + Spring Boot, increasing adoption by 15%</li><li>Authored unit/integration tests and targeted refactors, reducing bugs found after release</li><li>Strengthened security with OAuth 2.0 and JWT, and IAM hygiene practices ensuring HIPAA compliance</li></ul>	
<b>Tech Mahindra</b> - India	Jan 2021 – Dec 2021
<b>Web Developer</b>	
<ul style="list-style-type: none"><li>Rebuilt a legacy front end into a reusable Vue.js component library, modernizing the UI and improving load speed by 20%</li><li>Implemented Selenium/Jenkins pipelines, reducing post-release defects by 25%</li></ul>	

## PROJECTS

<b>AI-Powered Logistics Optimization</b> ( <i>Generative AI, LLM, RAG, Vector Search, AWS Bedrock, SageMaker</i> )	Oct 2024 – Present
<ul style="list-style-type: none"><li>Built a generative AI system for shipment scheduling and routing with vector search, cutting planning time by 60% and boosting throughput by 25%</li><li>Integrate RAG pipelines with logistics microservices for real-time decision-making and adaptive routing</li><li>Ship with Docker/Kubernetes, CI/CD checks, and useful service dashboards to support fast iteration</li></ul>	
<b>Document Intelligence &amp; Summarization Tool</b> ( <i>LLM, RAG, Transformers, NLP</i> )	Oct 2024 – Present
<ul style="list-style-type: none"><li>Develop an NLP platform to summarize shipment contracts and compliance documents using transformer-based NLP models and RAG pipelines</li><li>Reduce manual document review time by 55% and improve compliance turnaround speed</li><li>Used assistants to bootstrap tests and docs, handled sensitive content with care in storage and logs</li><li>Delivered automation framework that accelerated compliance workflows and was adopted org-wide</li></ul>	
<b>Intelligent Visual Search Engine</b> ( <i>Multimodal Search, ResNet, Elasticsearch, NLP</i> )	Jan 2022 – May 2023
<ul style="list-style-type: none"><li>Designed a multimodal search engine enabling image-to-image and text retrieval using ResNet embeddings and cosine similarity in Elasticsearch</li><li>Combined visual embeddings with NLP-based caption search for hybrid retrieval, improving result relevance by 35% over keyword-only search</li></ul>	