

Goutham Raju Kosuru Srinivasa

gouthamraju224@gmail.com | +1 (602) 696-9216 | linkedin.com/in/gouthamraju11 | github.com/Gouthamraju11

Education

Arizona State University, Tempe, AZ Master of Science in Information Technology	Dec 2025 GPA: 4.0/4.0
--	--------------------------

Technical Skills

Languages: Python, Java, JavaScript/TypeScript, C#, HTML, CSS

Framework/Backend: React.js, Angular, Next.js, .NET Core, Node.js, FastAPI, Django, REST APIs, Microservices, Git

Cloud/DevOps: AWS (Lambda, EC2, S3, DynamoDB), Azure (App Services, SQL, DevOps), Docker, CI/CD

Data/ML: TensorFlow, PyTorch, LangChain, SQL/NoSQL (MySQL, MongoDB, DynamoDB), Pandas, NumPy

Experience

Neudesic (An IBM Company) — Associate Software Engineer

Apr 2022 – Nov 2023

Tech: .NET Core, Angular, React.js, Next.js, TypeScript, Azure, Node.js, REST APIs, xUnit, Jest, Agile/Scrum

- **Developed** a full-stack field-service management system for a Fortune 500 client using Angular and .NET Core, serving **2,000+ daily active users** with **99.5% uptime** across **3 client-facing web applications**.
- **Built** 5+ microservice APIs using .NET Core and **Next.js BFF (Backend-for-Frontend) pattern**, processing **15,000+ daily requests** and integrating with **2 legacy systems** for seamless data flow.
- **Engineered** responsive UI components in **Angular and React.js** with **TypeScript**, implementing **15+ reusable component modules** and reducing frontend development time by **30%** through shared component library.
- **Optimized** backend caching layers and database query patterns by refactoring **12+ API endpoints** and implementing **SQL Server query optimization**, improving API throughput by **25%** and reducing response latency by **30%** across all client platforms.
- **Implemented** microservice modules for high availability using **Azure App Services** with sharded **Azure SQL databases across 2 instances**, enabling distributed transaction processing and reliable concurrent operations.
- **Established** CI/CD pipelines with **Azure DevOps** using YAML-based automation and parallelized job execution across **8+ services**, reducing deployment time by **60%** (from 30 to 12 minutes) and increasing release frequency by **40%**.
- **Enhanced** code reliability by implementing **dependency injection patterns** in **15+ service classes** and expanding unit test coverage from **25% to 75%** (**300+ test cases** added using xUnit and Jest), reducing post-deployment regression issues by **40%**.
- **Integrated** distributed tracing with **Azure Application Insights** and Log Analytics for real-time performance monitoring across **10+ key endpoints**, cutting production issue resolution time by **30%** (from 3 hours to 2 hours average).

The Sparks Foundation — Software Engineer Intern

Aug 2021 – Oct 2021

Tech: AWS CDK, Lambda, API Gateway, DynamoDB, Python, Docker, CloudWatch

- **Built** a serverless transaction monitoring API using AWS Lambda, API Gateway, and DynamoDB with event-driven architecture, handling **1,000+ daily transactions** with consistent **sub-300ms response times**.
- **Automated** infrastructure provisioning with AWS CDK (Infrastructure-as-Code), reducing environment setup time from **2 hours to 20 minutes** and enabling **repeatable deployments across 2 environments** (dev and staging).
- **Containerized** Lambda services with Docker for reproducible builds, improving **deployment consistency** and reducing configuration errors during CI/CD integration.
- **Implemented** CloudWatch monitoring with **5+ key metrics** (latency, errors, invocation count) and structured logging for real-time alerting, enabling faster debugging of issues during concurrent Lambda executions.

Projects

Multi-Agent AI Workflow Automation Platform

Aug 2025

Tech: AWS Bedrock, SageMaker, LangChain, CrewAI, FastAPI, React.js, FAISS

- **Developed** an intelligent workflow system using **LangChain and CrewAI** with **3 specialized agents** (extraction, summarization, decision-making) for autonomous document processing, reducing manual processing time by **60%** across **500+ test documents**.
- **Integrated** AWS Bedrock (Claude 3) and **FAISS vector database** for semantic search, building a **FastAPI backend** with **5+ API endpoints** achieving **sub-2 second response times** for document queries.

Book Recommendation System

Dec 2024

Tech: Python, Pandas, NumPy, SciPy, Collaborative Filtering

- **Built** a collaborative filtering recommendation engine using **cosine similarity** on **60,000+ user-book interactions**, achieving **85% accuracy** in predicting user preferences on test data.
- **Optimized** similarity computation using **NumPy vectorization**, improving performance by **10x** (from 30s to 3s), and implemented **caching** for top 100 popular books, reducing recommendation latency to **under 500ms**.

Reader and Object Detector for the Visually Impaired

Jul 2022

Tech: Python, TensorFlow, YOLO, OpenCV, Tesseract OCR, Raspberry Pi, Text-to-Speech

- **Developed** a dual-mode assistive system using **YOLOv3** for real-time object detection and **Tesseract OCR** for text extraction on **Raspberry Pi 4**, identifying **20+ common objects** and reading printed text with **text-to-speech output**.
- **Optimized** the inference pipeline to achieve **15+ FPS** on embedded hardware through **TensorFlow Lite quantization** and efficient preprocessing, enabling real-time assistance for visually impaired users in everyday environments.

Certifications & Publications

"Reader and Object Detector for Blind" - International Research Journal of Modernization in Engineering Technology and Science (IRJMETS), Volume 04, Issue 07, July 2022 — Impact Factor: 6.752