

# Anish Sudarshan Gada

(930) 274-9415 | gadaanish@gmail.com | linkedin.com/in/anish-gada | github.com/Anish478

## EDUCATION

### Indiana University Bloomington

Master of Science in Computer Science | GPA: 3.82/4.0

Bloomington, IN

Aug 2024 – May 2026

## TECHNICAL SKILLS

**Languages:** Python, Java, JavaScript, R, C++, C, HTML, CSS

**Frameworks:** Spring Boot, Angular, FastAPI, Next.js, LangChain, LangGraph

**Libraries:** TensorFlow, PyTorch, scikit-learn, OpenCV, SimpleITK, elastix, Pandas, NumPy, Matplotlib

**Machine Learning & AI:** Transformers, XGBoost, GANs, RAG, LLMs, sentence-transformers

**Cloud & DevOps:** AWS (Lambda, S3, EC2), GCP (Cloud Run), Docker, Kubernetes, RabbitMQ, GitHub Actions

**Databases & Tools:** PostgreSQL, MySQL, Qdrant, Redis, SQLite, Git, Gradle, Jupyter Notebooks

## EXPERIENCE

### Indiana University Bloomington

Teaching Assistant – Information Representation

Bloomington, IN

Aug 2025 – Present

- Instruct **90+** undergraduate students in database design, **SQL** query optimization, and data visualization using **ERD modeling** and SQL implementation exercises, grading **100+** assignments with detailed feedback

### Indiana University Indianapolis

Machine Learning Research Intern

Indianapolis, IN

May 2025 – Aug 2025

- Accelerated dataset processing by **99.8%** using **Python multiprocessing** with 6 concurrent workers, enabling real-time preprocessing across 95-patient prostate MRI dataset contributing to **CHIMERA 2025** Challenge at MICCAI
- Engineered multimodal MRI preprocessing pipeline for **T2W/ADC/HBV** sequences using **ITK-Elastix** co-registration with **3000 spatial samples**, **Nyul** intensity normalization, and prostate gland bounding box masking, reducing intensity variation by **52.4%** and validating **100% spatial registration accuracy** through **SimpleITK QA** framework

### SandLogic Technologies

Deep Learning Intern

Bangalore, India

Jan 2024 – Aug 2024

- Architected **LINGOFORGE** LLM Studio using **BPE/WordPiece** tokenization and **TensorFlow/PyTorch**, enabling RAG-based agentic orchestration across enterprise domains with **60%** model size reduction through knowledge distillation
- Deployed **3+** **AI solutions** across SandLogic's TXTR and EdgeMatrix ecosystem using **YOLOv8** and **OpenCV**, enabling **real-time inference** and driving **92%+ mAP** on pipe detection and **94%+** accuracy on CNN image classification

### S.V.G Enterprises

Software Engineer

Kalaburagi, India

Jun 2022 – Dec 2023

- Engineered proprietary inventory management system using **Python** and **MySQL** with normalized database schema, digitizing **2000+** **daily transactions** across procurement, storage, and distribution pipeline enabling **role-based access** for warehouse, procurement, and distribution teams to eliminate manual record-keeping
- Developed automated billing module with **real-time analytics dashboard** using **JavaScript** and **Chart.js**, integrating REST API endpoints to reduce processing time by **85%** and eliminate **95%+** manual errors across **50+** **product SKUs**

## PROJECTS

### Agentic RAG DocQuery | FastAPI, LangGraph, Qdrant, Next.js, Docker, Google Cloud Run

Mar 2026 – Present

- Architected **LangGraph**-based agentic RAG system with **Qdrant** vector store and **stateful reasoning loop**, ingesting **PDF/DOCX/Markdown** and web URLs via **LlamaIndex** loaders, enabling multi-hop query decomposition into 2–4 sub-questions with source-cited answers streamed via **FastAPI SSE** and **hop trace** transparency
- Deployed full-stack document intelligence platform on **Google Cloud Run** with GitHub Actions CI/CD pipeline, integrating **LlamaIndex** ingestion, **sentence-transformers** embeddings, and switchable **Groq/Gemini/Mistral** LLM providers

### Room Reservation | Spring Boot, Angular, PostgreSQL, Kubernetes, RabbitMQ, Redis, Docker

Aug 2025 – Dec 2025

- Architected **5-microservice** room booking platform using Spring Boot and Angular, enabling **RabbitMQ** async event-driven communication across User, Room, and Reservation services with role-based access control for **4 user types**
- Deployed cloud-native system on **Kubernetes** with **Redis** distributed session caching and **auto-scaling to 5M** concurrent users, sustaining **6,063 requests/second** at **1.06ms** average response time across **1.6M+** requests in k6 load testing

## PUBLICATIONS

### PROFUSEme: Prostate Cancer BCR Prediction via Fused Multi-modal Embeddings

Sep 2025

- Published at Cornell University arXiv (arXiv:2509.14051) for CHIMERA 2025 Challenge at MICCAI – C-index of **0.861**