

# Konda Balaji Rao

AI & Machine Learning Engineer

231fa18202@gmail.com | +91 83096 36226 | LinkedIn | GitHub | LeetCode | Portfolio

## SUMMARY

---

Motivated AI and Machine Learning Engineer with hands-on experience in building production-grade deep learning models, RAG-based AI systems, and computer vision pipelines. Proficient in Python, TensorFlow, PyTorch, and cloud platforms (AWS). Skilled in deploying ML solutions using Docker and CI/CD pipelines, with a strong foundation in data preprocessing, model optimization, and MLOps. Seeking an entry-level role to apply AI/ML and cloud engineering expertise to solve real-world problems at scale.

## EDUCATION

---

**B.Tech in Computer Science and Engineering (AIML)** | Vignan's Foundation for Science, Technology and Research

- CGPA: 7.50

**Intermediate** | Sri Chaitanya Junior College

Jan 2018 – Jan 2020

- CGPA: 9.54

**SSC** | Z.P. High School

Jan 2018 – Jan 2018

- CGPA: 9.58

## SKILLS

---

**Programming:** Python, SQL, C, HTML, CSS

**AI/ML & Deep Learning:** TensorFlow, Keras, PyTorch, Scikit-learn, OpenCV, LangChain, LangGraph, Hugging Face, NLP, Computer Vision, RAG

**Cloud & DevOps:** AWS (EC2, S3, Lambda, SageMaker), Docker, CI/CD, Git, GitHub Actions

**Databases & Tools:** PostgreSQL, Redis, Neo4j, FastAPI, Celery, Google Colab, VS Code

**Soft Skills:** Problem Solving, Analytical Thinking, Teamwork, Adaptability, Time Management

## PROJECTS

---

**HeartGuard AI – Medical AI Chatbot**

- Engineered a production-grade medical AI chatbot leveraging LangGraph-based multi-agent orchestration for autonomous query routing.
- Accelerated medical query response time by 60% and achieved 95%+ answer accuracy through a self-correcting RAG pipeline with vector search over 10K+ medical documents.
- Architected secure data handling with real-time PII redaction, JWT authentication, and end-to-end AES encryption.
- Integrated persistent session memory using Redis and PostgreSQL, enabling context-aware conversations across 500+ concurrent sessions.

*Technologies Used: Python, LangGraph, Neo4j, Redis, PostgreSQL, RAG, NLP*

*GitHub:* [github.com/9046balaji/Heart](https://github.com/9046balaji/Heart)

### **PDF Tools – Document Management Platform**

- Built a scalable enterprise PDF platform supporting conversion, merge, split, compression, and OCR across 15+ file formats.
- Reduced manual document handling time by 75% by implementing asynchronous task queues with Celery and real-time status updates via WebSockets.
- Secured document workflows with AES-256 encryption, watermarking, and digital signatures, achieving zero data-breach incidents.

*Technologies Used: Python, Celery, Redis, WebSockets, OCR, AES Encryption*

*GitHub:* [github.com/9046balaji/Pdf-Tools](https://github.com/9046balaji/Pdf-Tools)

### **Hospital Management System**

- Designed and deployed an automated appointment booking system, eliminating 100% of double-booking conflicts through custom scheduling algorithms.
- Optimized resource allocation logic, reducing average patient wait time by 35% across 200+ daily appointments.
- Increased administrative efficiency by 50% by building secure REST APIs with JWT-based role-based access controls.

*Technologies Used: Node.js, PostgreSQL, JWT, Redis, REST APIs*

*GitHub:* [github.com/9046balaji/Hospital-Management-System](https://github.com/9046balaji/Hospital-Management-System)

### **Heart Disease Prediction System**

- Trained and fine-tuned ensemble models (XGBoost, LightGBM) with Optuna hyperparameter optimization, achieving 92%+ ROC-AUC on a 10K+ patient dataset.
- Boosted minority-class recall by 25% by applying ADASYN oversampling and class-weight balancing techniques.
- Deployed the inference pipeline via FastAPI and integrated SHAP-based feature explanations for clinical interpretability.

*Technologies Used: Python, XGBoost, LightGBM, Optuna, SHAP, FastAPI*

*GitHub:* [github.com/9046balaji/collage-projects](https://github.com/9046balaji/collage-projects)

### **Computer Vision & Deep Learning Projects**

- Designed image classification and similarity-detection systems using transfer learning (VGG16), achieving 94%+ validation accuracy.
- Constructed automated edge detection and texture-analysis pipelines using Canny, Sobel, and GLCM, processing 1K+ images per batch.
- Delivered production-grade computer vision models with end-to-end preprocessing, augmentation, and inference pipelines.

*Technologies Used: OpenCV, TensorFlow, VGG16, CNN, Deep Learning*

*GitHub:* [github.com/9046balaji/collage-project](https://github.com/9046balaji/collage-project)