

MUKKTINAADH RAGHAVARAPU

+91-9676349758 | mukktinaadh@gmail.com | linkedin.com/in/mukktinaadh | github.com/mukktinaadh | rmnportfolio.vercel.app

EDUCATION

Parul University

B.Tech Computer Science Engineering — Big Data Analytics

Vadodara, Gujarat

2022 - 2026

PROFESSIONAL EXPERIENCE

AI/ML Intern — Product Engineer

Feb 2026 – May 2026

MyOnsiteHealthcare

Onsite

- Designed and shipped an **AI-powered healthcare scheduling system** — **live in production**; RAG pipeline over patient records and staff availability auto-assigns field workers to visits at **90%+ accuracy**, cutting manual dispatcher workload by **~60%** across 300+ daily visit assignments.
- Built multimodal-compatible **AI agent** (LangChain + LLM) executing scheduling commands via natural language — override assignments, reschedule visits, update patient data, answer schedule queries — integrated end-to-end with FastAPI + PostgreSQL + Redis on AWS.
- Engineered **agentic CI/CD automation system**: Cursor AI ingests user script + pipeline template + preset prompt, autonomously adapts the pipeline to the project via **GitLab MCP**, pushes code, triggers runs, detects failures, auto-rolls back, and re-attempts with fixes — zero manual intervention.

Volunteer AI Engineer

Oct 2024 – Sep 2025

Google Developer Group (GDG)

Remote

- Built Java/Python ML tooling adopted by **200+ students** across 5+ hackathons; improved model training accuracy **10–15%** via optimized preprocessing pipelines. Mentored 15+ contributors on W&B experiment tracking and MLOps practices.

FEATURED PROJECTS

PyroQueue | Python, Redis, FastAPI, Prometheus, Grafana, Docker

Jun 2026

- Built distributed task queue from scratch on Redis — priority queues via sorted sets, exponential backoff retry (3x), dead-letter queue, and live FastAPI monitoring dashboard; **no Celery or RQ**.
- Load tested at **893 req/sec** with **p99 <170ms** under 100 concurrent workers, 0% failure rate; instrumented with Prometheus + Grafana tracking queue depth, worker throughput, and failure rate.

Veriblock-FL | Python, PyTorch, gRPC, Federated Learning, zk-SNARKs, Blockchain

Apr 2025

- Built privacy-preserving federated learning framework for decentralized model training across isolated data silos; implemented **zk-SNARK** cryptographic verification to prevent gradient poisoning attacks from adversarial clients.
- Designed Byzantine fault-tolerant coordination (circuit breakers, dead-letter queues); system maintained liveness under **30% simulated node failure**. gRPC layer reduced inter-node overhead by 10% vs REST baseline.

Agentic Research Paper Evaluator | Python, LangChain, RAG, FAISS, FastAPI

Apr 2025

- Designed multi-agent RAG pipeline with semantic retrieval (FAISS), LLM-based methodology critique, and structured scoring agents; implemented citation graph traversal for automated related-work discovery.

TECHNICAL SKILLS

ML / DL: PyTorch, TensorFlow, HuggingFace Transformers, scikit-learn, CNNs, LSTMs, GANs, Federated Learning

NLP / GenAI: BERT, BioClinicalBERT, Fine-tuning (LoRA/QLoRA), RAG Pipelines, LangChain, Vector Search (FAISS, ChromaDB)

MLOps: MLflow, Weights & Biases, DVC, Docker, Kubernetes, CI/CD, Data Drift Monitoring, A/B Testing

Cloud / Backend: AWS (EC2, S3, Lambda), FastAPI, PostgreSQL, Redis, Kafka

Languages: Python, Java, JavaScript, SQL, C

RESEARCH & PUBLICATIONS

[**NeurIPS 2026 — Under Review**] *SCAR: SLO-Constrained Anytime Reasoning for Language Model Agents* — framework for controlling LLM inference under strict per-query latency SLOs; achieves better accuracy-violation tradeoffs than test-time scaling baselines on MATH-500, AIME 2025, GPQA. openreview.net #18900

[**Published**] *Veriblock-FL: A Blockchain-Based Verifiable Federated Learning Framework Using zk-SNARKs for Privacy-Preserving Model Aggregation* — IJSDR Vol. 11 Issue 4, Apr 2026 | Impact Factor: 9.15 | Paper ID: IJSDR2604263

[**Preprint**] *NeuroChainOps — Privacy-Preserving Blockchain MLOps for Federated NAS* — sciety.org (ResearchSquare)

ACHIEVEMENTS

Deployed **6+ end-to-end ML systems** with MLOps infrastructure (versioning, drift monitoring, CI/CD) in production environments.