

RUDRA GARG

+91 7303530051 ✉ rudragarg09006@gmail.com 🌐 rudragarg.dev 🌐 rudra-garg 🌐 Rudra-Garg

EDUCATION

Indian Institute of Information Technology, Guwahati

Bachelor of Technology, Computer Science and Engineering

2022 – 2026

CGPA: 8.42 (Sem 7)

TECHNICAL SKILLS

Languages: Python, C++, SQL, Go, JavaScript

Technologies: FastAPI, Flask, Temporal.io, SQLAlchemy, PyTorch, NumPy, React.js, Vue.js

Cloud/Databases: AWS, Google Cloud, PostgreSQL, Redis, Supabase, Firebase, MongoDB

Tools: Docker, Kubernetes, Helm, GitHub Actions, Prometheus, Grafana, Linux, Git

EXPERIENCE

Back End Developer Intern

June 2024 – Aug 2025

GRIG Technologies

Remote

- Architected **alert.ventory**, a resilient microservice using **FastAPI** and **Temporal.io**, implementing priority queues and human-in-the-loop verification for guaranteed enterprise notification delivery.
- Engineered a production-grade **Kubernetes** infrastructure using **Helm Charts**, automating NGINX Ingress routing and SSL/TLS rotation via Cert-Manager, reducing environment setup time by **90%**.
- Developed **request.ventory**, a multi-tenant inventory API using a **Hybrid Database** strategy (**Firestore** for NoSQL & **Supabase** for relational) to handle complex GST and Geo-point validations.
- Secured internal workflows by implementing **RBAC** and **Zoho OAuth2** integration with **JWT** session management across sensitive enterprise routes.
- Streamlined **CI/CD** pipelines using **GitHub Actions** and **Docker**, orchestrating automated deployments to **Google Cloud Run**.

PROJECTS

LOKI | Python, Faster-Whisper, Ollama, Scikit-Learn, CustomTkinter



- Engineered a privacy-centric, local-first voice assistant with **Faster-Whisper** and **Ollama**, achieving **sub-500ms** offline inference.
- Architected a dual-layer NLP pipeline using **SentenceTransformers** and a **CRF** model, achieving a weighted F1-score of **0.9977**.
- Built a modular **Agent framework** with **Voice Activity Detection (VAD)**, decoupling CPU-intensive inference from the UI thread.
- Integrated an asynchronous **Piper TTS** pipeline utilizing **priority queues** to deliver natural speech without blocking the main event loop.

MUCEDS | PyTorch, SUMO, Reinforcement Learning, LSTM



- Designed a **Hierarchical RL** framework (**Double DQN + MADDPG**) to optimize UAV fleet sizing and positioning across 8 global cities using **Eclipse SUMO**.
- Implemented a Spatial-Temporal **LSTM** predictive caching model trained on Azure workload traces, improving cache hit ratio from **18% to 78%**.
- Reduced average task latency by **53%** by proactively shifting task execution from cloud backhaul to local UAV cache.
- Increased system profit margins by **~300%** (vs baseline) by aligning UAV trajectories with realistic traffic flows derived from **OpenStreetMap**.

Portal Gambit | FastAPI, React, Firebase, WebRTC, Docker



- Architected a real-time chess variant with a custom **PortalChess** engine enabling non-Euclidean piece movement via **Firebase Realtime Database**.
- Built a **FastAPI** backend for ELO analytics with a hybrid auth system merging Firebase tokens and custom **JWTs**.
- Integrated **WebRTC (PeerJS)** for P2P voice chat and deployed via automated **Cloud Build** CI/CD pipelines to Google Cloud Run.
- Engineered a synchronized game timer using server-time offsets and **Firestore Transactions**, preventing client-side clock manipulation.

TerraQuest | Vue 3, Go (Gin), PostgreSQL, Docker, WebSockets



- Developed a concurrent multiplayer geography game using **Go (Gin)** and **Gorilla WebSockets** with synchronized real-time lobbies.
- Engineered a Go routine worker pool to validate and seed **81,000+** locations via Google Street View API into **PostgreSQL**.
- Implemented server-side Haversine scoring with composite **PostgreSQL indexes**, ensuring **sub-100ms** interaction latency.
- Managed complex frontend game state using **Pinia** and **Vue 3**, seamlessly integrating Google Maps Street View with the WebSocket layer.