

Jiya Jaiswal

✉ jaiswaljiya001@gmail.com | Delhi, India | +91 87001 87122



SUMMARY

Backend engineer in training, focused on distributed systems and infrastructure that has to hold up under failure. Learn by building end-to-end and stress-testing until things break, then fixing what's actually broken rather than what looks good on paper. Comfortable in Go and Python, working close to the systems layer – consensus, service communication, data partitioning – and looking to bring that instinct for correctness to a production engineering team.

EDUCATION

Kalinga Institute of Industrial Technology (KIIT)

B.Tech in Computer Science Engineering, CGPA: 9.26/10

- Relevant Coursework: DSA, DBMS, COA, OS, OOPj

Bhubaneswar, India

Aug 2024 – Oct 2028

Delhi Public School, Ruby Park

AISSCE (Class XII, Science), Percentage: 95%

Kolkata, India

Aug 2022 – Jun 2024

SKILLS

Languages: Python, Go, C#, Java, JavaScript, SQL

Backend & Distributed Systems: FastAPI, gRPC, Protocol Buffers, HashiCorp Raft, Consistent Hashing, Multi-Agent Pipelines, RAG, Knowledge Graphs

Tools & Infrastructure: Git, Unity, Docker, GitHub Actions, Linux, Unit Testing

Databases: MySQL, Oracle SQL

PROJECTS

• MazeMind – Adaptive AI Maze Game 🎮

May 2026 – Jun 2026

C#, Unity, ScriptableObjects, URP, Git

- Designed an **AI Director** tracking deaths, time-in-section, and hazard hits, using priority-sorted **AdaptationRuleSO** ScriptableObjects to retune hazard density, gem placement, and key spawns – with **AdaptationState** as a shared event bus decoupling the Director from per-room logic.
- Integrated a **procedural CheckboxFloorGenerator** producing connected solvable mazes and an in-game **DecisionLogger** capturing every adaptation event for replay and rule tuning; owned **Room 1 end-to-end** (3 sections, spawn points, section triggers, win/lose flow) integrating teammate-authored enemy mechanics.
- Validated AI Director rule firing across all difficulty transitions through 20+ playthroughs; the full game shipped with **3 playable rooms** featuring escalating mechanics – bullet traps, sliding platforms, forced-fall sequences, and NPC enemies.

• FactLens – Agentic Misinformation Detector 🕵️

MumbaiHacks 2025

Python, FastAPI, Groq LLaMA 3.1, Knowledge Graphs, Docker

- Built a **two-stage claim verification pipeline** labeling social posts True/False/Partly True at up to **96% confidence**, by chaining two Groq LLaMA 3.1 agents – one for claim extraction, one for evidence-grounded verdict reasoning.
- Prevented out-of-evidence hallucinations by constraining LLM verdict reasoning strictly to a **structured knowledge graph** (**kg.json**) via prompt-level context blocking; tested across True/False/Partly-True scenarios on a synthetic crisis dataset.
- Structured the repo with a clean **backend/ extension/ assets/** split for parallel team development and containerized the backend via **Docker** for single-command local setup.

• OmniDict – Distributed Sharded KV Store 🗄️

Jun 2025 – Aug 2025

Go, gRPC, Protocol Buffers, HashiCorp Raft, Consistent Hashing, Docker

- Verified **zero data loss under single-node failure** on a 3-node quorum by building a **kill-node test CLI** that performs mid-operation node termination and confirms all committed writes survive via subsequent **Get** calls.
- Implemented a **consistent-hash ring** guaranteeing **≤25% key remapping** on node join/leave by design; verified distribution behavior by running rebalance operations against a live cluster.
- Implemented the **gRPC server and CLI client** (Put, Get, Join, Delete, TTL) using Protocol Buffers, and authored the **Docker Compose** setup enabling one-command spin-up of a 3–5 node cluster for reproducible failure testing.

ACHIEVEMENTS

- Secured **AIR 646** in KIITEE Entrance Exam and awarded a **\$5000** scholarship
- **Top 30 Finalist** at Tech Zephyr 2025, IIT Bhubaneswar
- **Top 50**, Silicon Labs IoT Challenge 2026 – Smart Healthcare with Edge AI track
- **Top 500 / 3,000+ teams** at MumbaiHacks 2025, Multi-Agent AI track