

Ronak Raisingani

raisingani.r@northeastern.edu — 732-554-7669 — Aldie, VA
linkedin.com/in/ronakraisingani — github.com/ronak4 — ronakr.dev

SUMMARY

Software Engineer focused on implementing and productionizing existing AI/ML models in reliable, workflow-driven systems. Experienced in Python automation, LLM integration, RAG pipelines, structured outputs, prompt design, caching, retries, and backend orchestration with Kafka/Redis, Docker, and AWS. Built production systems that reduced manual work, improved throughput, and emphasized correctness, observability, and debuggability.

EDUCATION

Northeastern University, Boston, MA

Sep 2021 – Aug 2025

B.S. in Computer Science and Game Development

Relevant Coursework: Algorithms, Data Structures, OOP, Systems, Networks, Software Engineering, Databases

TECHNICAL SKILLS

AI/ML Systems: LLM integration, RAG pipelines, prompt engineering, structured outputs, OpenAI-compatible APIs, guardrails, inference workflows

Languages: Python, TypeScript/JavaScript, Java, C#, C++

Backend/Infra: Node.js/Express, Kafka/Redpanda, Redis, Docker, CI/CD (Jenkins, GitLab CI/CD), Kubernetes

Cloud/Data: AWS (EC2, S3, Lambda, RDS), Azure, SQL, PostgreSQL, MySQL, Microsoft SQL Server, MongoDB

Frontend/Other: React, Next.js, REST APIs, GraphQL, Firebase, PowerBI

WORK EXPERIENCE

Northeastern University ITS, Boston, MA — *Lead Automation Engineer*

Apr 2023 – Aug 2025

- Designed and shipped production automation pipelines using Python and AWS RDS, cutting processing time from 10h to 2h (-80%) and improving throughput for operational workflows.
- Built and operated integrations across multiple third-party APIs, handling pagination, validation, and guardrails to improve reliability and reduce manual review.
- Served as deployment on-call for releases; monitored rollouts and triaged a production deployment hiccup to keep workflows running with minimal disruption.

Northeastern University ITS, Boston, MA — *Pro CxT Technician / CxT Technician*

Aug 2022 – Mar 2023

- Supported 2,500+ users with rapid debugging and clear communication; trained 10+ hires; earned a 96% satisfaction rating through consistent issue resolution.

What Remains of Me, Boston, MA — *Quality Assurance Engineer*

May 2023 – Dec 2023

- Executed system and regression testing in Agile sprints; partnered with developers on bug reproduction, triage, and verification of performance fixes, contributing to a 30% reduction in downtime.

PROJECTS

RAG News Generation System (Python, LLMs, Kafka, Redis, Docker) — GitHub

Oct 2025 – Nov 2025

- Built a RAG-based article generation pipeline using existing LLMs to answer 70 structured questions across 10 congressional bills and assemble final articles in under 10 minutes.
- Implemented Kafka/Redpanda queueing and Redis-backed state/caching with idempotent workers, retries with exponential backoff, and progress telemetry to improve reliability and debugging.
- Added structured outputs, prompt constraints, and failure handling to improve answer consistency and maintain 100% completion in repeated test runs.
- Integrated canonical legislative sources and retrieval filters to ground outputs in bill text, votes, hearings, and sponsor activity, improving traceability and reducing hallucination risk.

Discord Voice Assistant (TypeScript, Node.js, discord.js, PostgreSQL, Local LLM/TTS/STT) *2024 – Present*

- Built and operated a 24/7 Discord bot in TypeScript across 2 servers, supporting YouTube/Spotify/SoundCloud playback, playlists, queue controls, and per-guild runtime state.
- Integrated local LLM and TTS features, including configurable personality modes and voice-driven greetings, with runtime controls for behavior, latency, and fallback handling.
- Engineered resilient runtime infrastructure with reconnect and cleanup logic, retry-based stream resolution, Git-push auto-deploy gated by tests, PostgreSQL persistence, and FFmpeg-based audio effects including crossfade and 3D audio.
- Built a custom greeting pipeline around local TTS and conditioning audio, with timeout protection and degraded-mode fallbacks to keep voice features usable in live server environments.

Strategy Town (React, TypeScript, Node.js) — Live Site

May 2025 – Jun 2025

- Built real-time collaboration features including direct messaging and a synchronized whiteboard with incremental updates and render optimizations for multi-client activity.
- Designed a move-queue to serialize player actions and prevent race conditions; validated behavior with 2–7 concurrent users and structured the mechanism to scale beyond the tested load.