

DHRUVIL RANGANI

Software Engineer | Python · TypeScript · Go | Full-Stack, APIs, and AI Products in Production
Toronto, ON | dhruvilrangani007@gmail.com | [LinkedIn](#) | [GitHub](#) | Portfolio: [dhruvilrangani.com](#)

SUMMARY

Full-stack engineer who ships products end to end and stays to operate them. Works across Python, TypeScript and Go: REST and GraphQL APIs, relational schema design, React interfaces on top, and the deployment underneath. Has built and run AI products with real users rather than demos, including a tool-calling agent gated in CI on measured correctness and latency, and a retrieval-augmented document platform in production since launch. Picks up unfamiliar stacks on their own terms rather than by analogy, and has a working answer for how these systems fail once people depend on them.

TECHNICAL SKILLS

- **Languages:** Python, TypeScript, JavaScript (ES2022+), Go, Java, SQL; data structures, algorithms, time and space complexity, object-oriented and functional design
- **Backend & APIs:** Node.js and Express, FastAPI, Django, REST and GraphQL API design, microservices, third-party and protocol-level integration (IMAP, SMTP), JWT authentication, concurrent pipelines with worker scheduling, caching and queues; PostgreSQL, MySQL, schema design and migrations (Prisma), relational data modeling, MongoDB, DynamoDB, Redis
- **Frontend:** React (hooks, composition, memoization, re-render behavior), Next.js 14 App Router, Tailwind CSS, React Query, responsive layouts, perceived-latency and loading-state design
- **AI & LLM:** OpenAI API integration, RAG pipelines, tool-calling agents, deterministic guardrails around probabilistic components, eval harness design, CI release gating, prompt and context tuning
- **Cloud & Practice:** AWS (EC2, ECS/Fargate, Lambda, S3, RDS, SQS), GCP, Docker, CI/CD (GitHub Actions), testing (Jest, pytest), TDD to 90%+ coverage, structured logging, production debugging, code review, technical documentation, Git, Agile

EXPERIENCE

DiceMail | Software Engineer

Toronto, ON | April 2025 - Present

- Build and operate an AI-powered, self-hosted email platform end to end on a three-person team, from first conversation through deployment and maintenance, including the LLM layer on the OpenAI API for composition, smart replies and summarization.
- Own the backend: a Node.js and Express REST API with Prisma and PostgreSQL schema design and migrations, JWT authentication, and protocol-level send and receive over IMAP and SMTP against a Postfix and Dovecot mail stack we run ourselves, which means owning the failure modes in production.
- Model threaded conversations in PostgreSQL by reconstructing them from message headers, which is harder than it looks because clients set those headers inconsistently, plus tagging and full-text search.
- Ship weekly, review peers' code and write the documentation the rest of the team works from, contributing to a 30% improvement in feature delivery speed.

Vital Plus | Full-Stack Developer (Intern)

Toronto, ON | January 2024 - April 2024

- Built REST and GraphQL APIs in Node.js, Express and TypeScript on a four-person Agile team, reaching 90% unit-test coverage through test-driven development.
- Designed PostgreSQL schema migrations and automated CI/CD with GitHub Actions, Docker and AWS EC2, enabling zero-downtime deployments with CloudWatch monitoring.

PROJECTS

Voice Order Agent | Python agent with deterministic guardrails and a two-mode eval harness | [View Project](#)

Aug 2026

- Built a tool-calling agent that runs a transaction end to end against a live catalog, on one rule: the model decides what the user meant, never what the system does. Prices come from the system of record, invalid records are rejected by the tool layer, and anything with money or safety exposure is a deterministic gate, not a prompt instruction.
- Built the eval harness as a first-class part of the repo: 12 golden conversations run in two modes against the same expectations, a scripted stand-in model gating every PR at zero cost and a real model nightly, with five release thresholds including escalation recall at 1.00 and p95 latency under 2500ms.
- It caught a defect on its first run that manual testing missed, and the defect was not in the model: a gate evaluated inside a code path that returned early, so one class of action bypassed it.
- Building on the agent framework surfaced a silent failure mode in it, where two parts of the same integration read different environment variable names for the same credential. Fixed upstream: [livekit/agents PR #7090](#)

GoFeedEngine | Concurrent ingestion service in Go, REST API, PostgreSQL | [View Project](#)

May 2024 - Jul 2024

- Distributed Go backend with a concurrent ingestion pipeline, worker scheduling across goroutines, API-key authentication and REST endpoints over PostgreSQL, sustaining 99.9% uptime and 40% lower response latency under load. Written to learn Go on its own terms rather than by analogy.

AI Document Analysis Platform | Python, React, TypeScript, PostgreSQL, RAG | [View Project](#)

Jan 2025 - Mar 2025

- Shipped solo and operated in production since launch: Python backend, React and TypeScript front end, PostgreSQL, Docker, cutting document turnaround by 60%.

EDUCATION

Centennial College | Ontario College Advanced Diploma, Software Engineering Technology (Honours) | GPA 4.01/4.5

2025