

# Ruchit Patel

## Applied AI/ML Engineer

Chicago, IL | ruchitp26@protonmail.com | (773) 663-0733 | ruchitp.com | linkedin.com/in/ruchit-p | github.com/ruchit-p

Applied AI/ML Engineer with experience taking AI powered systems from idea to production, building agents, automations, and workflows that remove bottlenecks across engineering and operations. Hands on background with LLM APIs, embeddings, retrieval, MCP, RAG, structured tool calling, and end to end services in Python and TypeScript. Built production systems spanning APIs, web scraping, browser automation, vector search, and security sensitive workflows with measurable gains in speed and productivity. US Citizen - Willing to relocate for the right opportunity.

### TECHNICAL SKILLS

**ML & AI:** LLMs (GPT, Qwen, LLaMA), RAG, Vector Embeddings, MCP, LangGraph/LangChain, Qdrant, PyTorch, vLLM, sglang, Ollama, TensorFlow

**Backend & Tools:** FastAPI, Node.js, Next.js, React, Bun, Git, Linux, Hono, Turborepo

**Languages:** Python, TypeScript/JavaScript, Rust, Swift, SQL, C++, Java, Bash

**Infrastructure:** AWS, Azure, GCP, Cloudflare, Vercel, Docker, Kubernetes

**Databases:** PostgreSQL, SQLite, Qdrant, Redis, MongoDB, ClickHouse, Cloudflare D1

### PROFESSIONAL EXPERIENCE

#### Founder & Lead Engineer | Org in a Box

Sep 2025 – Present

- Shipped a production multi-tenant AI workforce platform on Azure: isolated per-agent noVNC sandboxes, Postgres + pgvector memory, Entra SSO, and a pre-warmed Container App pool with sub-second cold starts.
- Built the agent runtime on OpenCode (Bun/TypeScript) with plugins for gateways, memory, RBAC, and subagent delegation; integrated Teams/Slack/Email via a unified gateway with DB-backed pairing.
- Authored agent-tunnel, an HMAC-signed WebSocket relay letting cloud agents invoke local filesystem/shell/desktop actions with per-capability consent, plus a cron + webhook trigger system for always-on workflows.

#### Founder & Lead Engineer | Dynasty Platforms

Feb 2025 – Sep 2025

- Built and shipped Dynasty, a cross-platform private family social network (iOS, Android, web) on Expo React Native and Next.js, with a feed, event invitations, interactive family tree, encrypted document vault, and end-to-end encrypted messaging.
- Designed and implemented a legislative MCP server in Rust with a custom REST API and scraper to ingest real time U.S. government data, extract structured information, and support AI agent tool calling.
- Created Android and iOS MCP servers that autonomously test and improve apps through agentic workflows, reducing development time by 60%.

#### IT Automations Engineer | AdvancedBits

Jul 2024 – Feb 2025

- Built a Retrieval Augmented Generation (RAG) data retrieval pipeline with Python, Qdrant, vector embeddings, and reranking models across 3,000+ documents, reducing information lookup time by 40%.
- Automated end to end data workflows with Azure and n8n for provisioning and lifecycle management across multi step processes, eliminating manual intervention and improving reliability.
- Managed technology escalations across 50+ client environments, driving root cause analysis, debugging, and system ownership for more stable cross functional operations.

#### EECS Lab Engineer | Florida Atlantic University

Sep 2023 – May 2024

- Designed and built a laboratory inventory management application with React and SQL to centralize equipment tracking and student checkout workflows across 23 plus engineering courses.
- Engineered an embedded IoT monitoring system with C++, ESP32, wireless connectivity, camera recording, and a React dashboard to process real time sensor data, reducing unauthorized access events by 40%.

#### L2 IT Engineer | Skynet Enterprises

Mar 2023 – Jun 2023

- Built a real time monitoring dashboard with React and Firebase to aggregate telemetry from 200+ servers, circuits, and databases across 40 clients into a unified operational view.
- Implemented an event driven alerting and automated ticket routing system for outage response workflows, reducing SLA response times by 80%.

### PROJECTS

#### LegisMCP | Rust, Tokio, MCP, Next.js

2025 – 2026

- Built an open-source single-binary MCP server exposing 161 tools across 29 U.S. government data sources (Congress, FEC, SEC, FDA, CDC, EPA, FEMA, IRS/Treasury, plus all 50 state legislatures) to Claude, ChatGPT, Cursor, and Windsurf via stdio and Streamable HTTP transports.

#### Essence Web Scraping Engine | Rust, Tokio, MCP, chromiumoxide

2025 – 2026

- Built an open-source web retrieval engine with two-tier rendering (HTTP fast path with automatic Chromium fallback for JS-heavy SPAs), 6 REST endpoints plus MCP server, and Python/TypeScript SDKs; published to crates.io and benchmarked at a 91% quality win rate and 1.5x faster than Firecrawl across 35 real-world URLs under LLM-judge evaluations.

#### MacWright — macOS Desktop Control MCP Server | TypeScript, Node.js, MCP, Accessibility API, AppleScript

Feb 2026

- Built a macOS automation MCP server for AI agents, exposing 72 tools for full macOS desktop control, screenshots, mouse/keyboard input, scroll, clipboard, window management, and UI automation via Apple Accessibility API and Safari browser integration.

### EDUCATION

Florida Atlantic University — B.S. Computer Science, Minor in Cybersecurity | GPA: 3.8

May 2024