PATRICK SUIJK

Software Engineer

(804) 317-8544 | patrick.suijk@gmail.com | LinkedIn | Portfolio

Software engineer and tech lead who builds and operates production systems, currently working at the AI application layer and moving deliberately toward infrastructure and systems engineering. Delivers RAG systems, multi-step agents, and data extraction pipelines for enterprise clients, and owns the cloud infrastructure beneath them: Terraform across AWS and Azure, CI/CD, and production observability. Writes Go across a production backend, open source, and personal systems projects to build depth below the application layer. Long-term focus is the intersection of AI and systems engineering.

TECHNICAL SKILLS

Languages: Go, Python, TypeScript, C, SQL

Infrastructure & Tools: AWS (Lambda, S3, EC2, RDS), Azure, Terraform, Docker, CI/CD, observability/monitoring, PostgreSQL, MS SQL Server, Git

Backend & Frameworks: FastAPI, Node.js, Express, NestJS, React, React Query, TypeORM, Drizzle, NumPy

AI/ML: LLM application development, RAG, agent architecture, prompt engineering, evals; OpenAI, Claude, Gemini

PROFESSIONAL EXPERIENCE

RTS Labs July 2025 – Present

AI Software Engineer / Tech Lead

- **Contribute to a large-scale production platform for field-service scheduling and routing,** with a backend built entirely in Go and integrations across multiple external data sources (Go)
- Architect and deliver AI/LLM applications for enterprise clients across logistics, healthcare, manufacturing, and financial services: RAG systems, multi-step agent workflows, document extraction pipelines, and conversational AI (Python, FastAPI, TypeScript, Node.js, AWS, Azure)
  - Built an automated invoice processing pipeline for a major logistics company that extracted and normalized fields from varied contractor formats, eliminating 87% of manual headcount (Python, AWS Lambda, S3, Claude)
  - Designed a natural-language-to-SQL agent with an iterative refinement workflow over a database populated from 4,000+ PDFs and legacy Excel files for a manufacturing client (Python, FastAPI, MS SQL Server, Azure, Gemini)
- **Own the cloud infrastructure beneath delivered systems:** provision and manage resources with Terraform across AWS and Azure, build CI/CD pipelines, and run production observability on latency distributions, tool-call success rates, and retry patterns, which has surfaced more issues than offline testing
- Lead technical discovery and scoping, write proposals and SOWs, mentor junior developers, and establish coding standards across projects

NimbleBrain Feb 2024 – July 2025

Full Stack Developer

- Owned the backend end-to-end: designed RESTful API routes, modeled the PostgreSQL data layer with TypeORM, and implemented S3 presigned-URL upload flows for direct client-to-storage uploads (TypeScript, Node.js, Express)
- Led a refactor moving business logic out of bloated route handlers into a dedicated service layer, so routes handled HTTP concerns only and domain logic became independently testable and reusable

SYSTEMS & OPEN SOURCE

Golem (Go) Agent harness built from scratch with no framework dependencies; handles tool-use loops, session management, memory, process lifecycle, and file system operations directly.

Ollama (Go, open source) Investigated request-queueing behavior in the model scheduler; reproduced and root-caused an issue where the queue-depth limit is bypassed for already-loaded models, and documented a two-queue diagnosis for the maintainers.

ADDITIONAL PROJECTS

Neuromod (Python, TypeScript) Composable, type-safe LLM inference library with provider-agnostic model integration, step-function pipelines, and Pydantic-based tools.

Cardlang (TypeScript, Node.js, Angular, Postgres, Lambda) Natural language search engine for Magic: The Gathering cards; runs on Neuromod.

EDUCATION

B.S. Computer Science | The Ohio State University GPA: 3.8 | Dec 2023

Microsoft Certified: Azure Fundamentals (AZ-900), 2024