

Prateek Mulye

Senior Software Engineer | Backend & Distributed Systems | Applied AI

Santa Barbara, California | prateekmulye@gmail.com

prateekmulye.dev | github.com/prateekmulye | www.linkedin.com/in/prateekmulye

Summary

Software engineer with 11+ years of experience across financial services, data platforms, and manufacturing. I design backend services and recovery workflows in Java, Elixir, and Python. At Agilent, I build manufacturing workflow applications and lead internal AI enablement. My personal AI work covers retrieval, agent orchestration, and streamed responses.

Experience

Agilent Technologies | Senior Software Engineer, Manufacturing

Carpinteria, California | January 2026–Present

- Built reporting and action-tracking applications for Gemba walks (Lean shop-floor observation rounds), reducing manual coordination through Power Apps and Power Automate.
- Both applications are used by colleagues and have drawn interest from global manufacturing and technology leadership for rollout across manufacturing plants.
- Support manufacturing processes and recipe authoring, working with R&D and engineering on process plans.
- Lead internal AI education and enablement, helping engineers adopt AI in their day-to-day work.

HG Insights | Software Engineer

Santa Barbara, California | September 2021–November 2025

- Owned the architecture and phased rollout of a geo-standardization service with Data Solutions and downstream stakeholders. It normalized 15M+ U.S. and Canadian company records into canonical shapes for ingestion and UI workflows.
- Scaled PostgreSQL beyond 30M rows with partitioning, composite indexes, and read replicas, reducing customer-facing dashboard latency by more than 60%.
- Built an idempotent contract-processing pipeline in Elixir and Phoenix with Oban Pro. It handled 80K+ records and reduced reporting turnaround by 60%.
- Orchestrated scheduled, on-demand, and event-triggered company-spend refreshes across corporate hierarchies, moving delivery from a manual monthly batch to a weekly cadence.

Cognizant (client: Capital One) | Software Engineer

Vienna, Virginia | November 2018–September 2021

- Helped decompose a legacy Java application into contract-first services so teams could release changes independently.
- Owned Kafka workflows using Spring Boot and DynamoDB, including dead-letter handling, bounded retries, backoff, and idempotent replay for recovering failed financial-transaction messages.
- Built an OpenResty service-virtualization layer with NGINX and Lua that sustained 12K+ transactions per second in test environments and removed an integration-testing bottleneck.

Earlier experience

Tuutkia | Software Engineer (part-time) | September–November 2018

San Jose, California.

Illinois Department of Public Health | Software Engineer Intern | August 2017–June 2018

Springfield, Illinois. Built a reporting application that replaced manual data processing and reporting procedures.

AurionPro Solutions | Software Engineer | March 2013–September 2015

Pune, India. Built secure login and anti-phishing middleware in Java for a Kenyan bank's online-banking platform under regulatory and latency constraints.

Selected engineering projects

Assay (FinResearch AI) | Community-project contribution

Contributed to the SuperDataScience CP044 FinResearch AI community project. Implemented parallel analyst and bull, bear, risk, and reporter workflows in Python and LangGraph, streamed reports through FastAPI, and an evaluation harness that records judge preferences, token usage, and estimated cost. The reviewed implementation uses pgvector search with keyword fallback.

github.com/prateekmulye/assay

ChatFormula1 | Developing v2

Built a Python/LangGraph pipeline with schema-validated query routing, parallel retrieval, and fallback handling. The developing v2 includes an Elixir/Phoenix streaming gateway and a typed React client that handles duplicate and replayed stream events. Repository implementation is separate from the earlier public v1.

github.com/prateekmulye/ChatFormula1

Technical skills

Languages and APIs: Java, Elixir/Phoenix, Python, TypeScript, SQL, REST, GraphQL, React.

Data and delivery: Kafka, PostgreSQL, DynamoDB, Oban Pro, AWS, Docker, Kubernetes, GitHub Actions.

AI and testing: LangGraph, LangChain, retrieval workflows, Pydantic schemas, evaluation harnesses, JUnit 5, ExUnit, Testcontainers, OpenTelemetry, Datadog.

Manufacturing workflows: Power Apps, Power Automate, workflow applications and process improvement.

Education

M.S. Computer Science | University of Illinois at Springfield | 2018

Post-Graduate Diploma in Advanced Computing | C-DAC, Pune | 2013

Bachelor, Computer Application | Savitribai Phule Pune University | 2009–2012

Course completion

Generative AI with Large Language Models | DeepLearning.AI, Coursera | February 2026

[Verify course completion](#)

LangChain & LangGraph - Develop AI Agents | Udemy | November 2025 (LinkedIn-listed course completion)

[View LinkedIn-listed course credential](#)

Artificial Intelligence A-Z 2025: Agentic AI, Gen AI, and RL | Udemy | October 2025 (LinkedIn-listed course completion)

[View LinkedIn-listed course credential](#)