

# Shreyas Kaldate

New York | +1 (551) 424 8886 | sk12898@nyu.edu | LinkedIn | GitHub | Portfolio

## Education

**New York University** GPA : 3.6/4

Masters of Science in Computer Science

New York City

2025 - 2027

**Savitribai Phule Pune University** GPA : 3.85/4

Bachelors of Engineering in Computer Engineering

Pune, Maharashtra, India

2019 - 2023

## Technical Skills

**Languages:** Python, Java, JavaScript, TypeScript, C++, SQL

**Backend & APIs:** Spring Boot, FastAPI, REST API Design, Microservices, OAuth2/JWT, Kafka

**Cloud & DevOps:** AWS (EC2, S3, Lambda), Docker, Kubernetes, GitHub Actions, CI/CD

**Databases:** PostgreSQL, MySQL, MongoDB, Supabase

**AI & Data:** LangGraph, LangChain, RAG, Scikit-learn, Pandas, NumPy, Anthropic Claude API

**Practices:** Agile/Scrum, OOP, Data Structures & Algorithms, Code Review, Git, Swagger/OpenAPI

## Work Experience

### Jio Platforms Limited

Mumbai, India

*Software Development Engineer-I*

Jan 2024 - July 2025

- Built and shipped **production-ready Java/Spring Boot backend services** for Jio CloudXp at **million-user scale**, writing REST APIs, debugging logic, and contributing to a large backend codebase in Agile sprints.
- Designed a **Kafka-based event pipeline** handling **millions of events/day** – profiled bottlenecks, reasoned through data flow end-to-end, and improved throughput by **40%** at **99.9% uptime**.
- Optimized **REST APIs** with OpenFeign on AWS, cutting latency by **30%** and deploy cycles by **25%** through CI/CD automation and systematic debugging.
- Worked with **MySQL and MongoDB** in production, writing and optimizing queries, managing schema changes, and ensuring data integrity across distributed services.

### Pie Infocom Pvt Ltd

Mumbai, India

*Software Developer Intern*

Nov 2021 - Dec 2021

- Built **RESTful APIs** in Java/Spring Boot with MySQL, wrote unit tests, containerized with Docker, and documented endpoints with **Swagger/OpenAPI** reducing onboarding time by **30%**.
- Debugged backend logic and data flow issues in a live microservices codebase, contributing fixes through structured **Git** workflows and code reviews.

## Projects

### Caregiver Co-Pilot – Agentic Healthcare Backend

Apr 2026 - Present

*LangGraph, FastAPI, Python, Qwen3, GLM-4.5-Air, Qdrant, Supabase PostgreSQL, BGE-M3, Langfuse*

- Building a **FastAPI backend** powering a 3-model LangGraph pipeline – Router, Generator (tool-calling loops for drug checks, vitals logging, symptom triage via RxNav/OpenFDA), and Verifier validating every output before delivery.
- Designed **REST API endpoints** for session management, care recipient data, and agent execution with **Supabase PostgreSQL** for structured data persistence and **JWT auth**.
- Engineered **hybrid RAG** over Qdrant with metadata-filtered retrieval and **human-in-the-loop escalation**; full observability via **Langfuse** tracing across all backend nodes.

### NutriScan AI – Automated Health Data Pipeline

Jan 2026 - Mar 2026

*Next.js, FastAPI, Python, Anthropic Claude API, Tesseract OCR, Supabase PostgreSQL, Docker, MCP*

- Built a **FastAPI backend** that chains OCR, a Python classification engine, and Claude API tool-use into a pipeline delivering structured recommendations in under **60 seconds**.
- Designed **REST APIs** for report ingestion, biomarker querying, and result delivery, backed by **Supabase PostgreSQL** with structured schema design and JWT auth.
- Deployed with **Docker**, full **GitHub Actions CI/CD**, and **MCP-based agents** for autonomous multi-step external API workflows.

### OrbiNasaSense – Anomaly Detection on NASA Telemetry

Sept 2025 - Nov 2025

*Python, PyTorch, Scikit-learn, Pandas, NumPy, NASA Telemetry APIs, Streamlit*

- Built an **ML anomaly detection system** on real NASA telemetry time-series data using Isolation Forest and PyTorch sequence classifiers to flag abnormal spacecraft behavior.
- Designed a **Python data pipeline** querying NASA APIs, cleaning and normalizing raw sensor streams via **Pandas and NumPy**, and benchmarked models on precision and recall across failure patterns.
- Shipped a **Streamlit dashboard** for real-time telemetry visualization with clear documentation of findings, failure modes, and model performance.