

JAYRAJ M. PAMNANI

jimp10051@nyu.edu | [Linkedin](#) | [Github](#) | [About Me](#) | +1(646) 642-1043

SUMMARY

Software engineer focused on backend systems, cloud infrastructure, CI/CD, and ML-enabled production applications. Experienced in building full-stack products, automating business workflows, deploying cloud-native services, and improving reliability, cost, and delivery velocity across production systems. Strong background in Python, TypeScript, Kubernetes, AWS/GCP, and data/ML infrastructure.

PROFESSIONAL EXPERIENCE

GBCS Group - Software Engineer, Calgary, Canada

11/2025 - 02/2026

- Implemented and optimized cloud infrastructure using Infrastructure as Code (IaC) to improve resource utilization, resulting in a **28% reduction in hosting and maintenance costs**.
- Improved CI/CD pipelines and deployment workflows, **accelerating release cycles by 40%** while maintaining 99.9% system availability.
- Helped redesign microservice boundaries and deployment workflows across a 5-engineer team, reducing technical debt, improving maintainability by 30%, and establishing patterns for future service development.
- Defined deployment standards for Docker/Kubernetes-based services, improving release reliability and reducing manual intervention.

Freelance Software Developer, New York, NY

07/2025 - 09/2025

- Engineered and deployed a Django-based web app that automated data transfer between **QuickBooks and KatanaMRP**, cutting the client's manual bookkeeping workload by **~85%**.
- Built a robust **REST API integration** using Django (backend) and JavaScript (frontend), enabling seamless synchronization of **10,000+ financial and inventory records** with zero data loss.
- Designed and implemented **5+ custom verification layers**, achieving **99.8% data accuracy** before syncing to KatanaMRP. Managed full project source code using Git and GitHub, maintaining clean version history and enabling reliable deployments to AWS EC2.

Swaroop.ai - AI Engineer, Ahmedabad, India

01/2024 - 08/2024

- Integrated OpenAI Whisper (LLM-based model) into production for Speech-to-Text/Transcription, increasing accuracy by 30% and reducing inference latency by 20%.
- Developed and streamlined MLOps pipelines (data preprocessing, training, deployment orchestration) to support real-time synthesis, cutting model turnaround time to under 24 hours.
- Optimized real-time synthesis pipelines by fine-tuning Coqui TTS/STT models across multiple languages, boosting TTS performance by 25% and STT accuracy by 18%.

New York University - Teaching Assistant - Machine Learning, New York, NY

09/2025 - 12/2025

- Guided **50+** graduate students through ML fundamentals, including preprocessing pipelining, supervised/unsupervised learning, Deep Learning, and model optimization techniques.
- Conducted weekly office hours to debug Python code, explained algorithms, taught ML topics, and assisted with PyTorch implementations.

SKILLS

Languages & Backend Systems: Python, Java, Go, TypeScript, FastAPI, Django, Flask, RESTful API Architecture | **Cloud Infrastructure & DevOps:** AWS, GCP, Azure, Docker, Kubernetes, Terraform, GitHub Actions, CI/CD Pipelines | **Data & Machine Learning Engineering:** PostgreSQL, MongoDB, Snowflake, Apache Spark, Kafka, Airflow, MLflow | **Frontend Development:** React, Next.js, TypeScript, shadcn/ui Component Libraries

EDUCATION

New York University | Master's in Computer Engineering | GPA: 3.7/4.0 | May 2026

Relevant Coursework: High Performance Machine Learning, Machine Learning Operations, Deep Learning, Database Systems, Big Data

Parul University | Bachelor's in Computer Science & Engineering with Specialization in AI | GPA: 3.4/4.0 | May 2024

Relevant Coursework: Data Structures and Algorithms, Operating Systems, Compiler Design, GPU computing, Pattern Recognition

PROJECTS

[ActualBudget Transaction Categorizer](#) | Python, FastAPI, PostgreSQL, MLflow, Docker, Kubernetes, Terraform

- Designed a reproducible ML platform around data ingestion, model evaluation, serving, experiment tracking, and GitOps-based infrastructure, enabling repeatable model updates and controlled deployment of categorization services.
- Built a Python-based data platform (API Backend) with PostgreSQL that ingests and categorizes financial transactions using a multi-stage ML classification pipeline.
- Engineered end-to-end data lifecycle management, implementing automated deployment pipelines (CI/CD) and Infrastructure as Code using Docker, Kubernetes, and Terraform for model serving and scaling.
- Designed scalable data storage and retrieval systems supporting low-latency access to transaction histories and analytical outputs.

[HexDrop \(Secure File Transferring WebApp\)](#) | Next.js, TypeScript, Prisma, PostgreSQL, AWS, Docker, K8s

- Developed a secure, full-stack file-sharing application using Next.js and TypeScript, implementing client-side encryption and managing the end-to-end user-facing flow.
- Designed a cloud-native deployment architecture on EKS with HPA, ALB Ingress, RDS PostgreSQL, External Secrets, and GitHub Actions to support secure, reproducible deployments and horizontal scaling.

[EyeConnect](#) | WebRTC, Supabase, OpenRouter AI, React, TypeScript

- Full-stack accessibility platform integrating React/TypeScript (Frontend) and WebRTC for peer-to-peer streaming; awarded 2nd place at NYU Hacks 2025.
- Implemented AI Agent Development patterns using OpenRouter AI for advanced features (e.g., natural language command processing, tool/function calling logic).
- Architected the platform integrating React/TypeScript with shadcn/ui to deliver a high-contrast, accessible UI, applying UX best practices for low-latency visual support.