

Yash Thakkar

thakkaryash21@gmail.com • yash-thakkar.com • linkedin.com/in/yashthakkar21/ • 412-699-0335

EDUCATION

Carnegie Mellon University, Pittsburgh, PA Dec 2025
Master of Information Systems Management **GPA - 4.04**

Relevant Coursework: Cloud Computing, Distributed Systems, LLM Applications, Machine Learning, MLOps, Java
Teaching Assistant: Engineering Data-Intensive Scalable Systems, Data Science & Big Data, Agile Methods

University of Mumbai, Dwarkadas J. Sanghvi College of Engineering May 2022
Bachelor of Engineering, Mechanical Engineering **GPA - 3.92**

SKILLS

Languages & Libraries: Python, TypeScript, JavaScript, Java, SQL, NumPy, Pandas, Bash/Shell

AI & ML: PyTorch, Hugging Face, LangChain, Prompt Engineering, RAG, Scikit-learn, MLflow

Infra & DB: AWS, GCP, Azure, Docker, Kubernetes, Terraform, Kafka, PostgreSQL, MongoDB, Redis, Neo4j

Web: React.js, Node.js, Next.js, FastAPI, Flask, GraphQL, REST API, Redux, WebRTC, Tailwind, HTML/CSS, Pytest

EXPERIENCE

Newfront Insurance, San Mateo, CA May 2025 – Aug 2025
Software Engineering Intern

- **LLM Document Processing:** Built an end-to-end financial document processing system using FastAPI, GraphQL, PostgreSQL, & OpenAI APIs, **reducing manual review time by over 90% & saving \$175k annually**
- **System Design & UI:** Architected the full system design, led UX reviews, and implemented Next.js components within Newfront's Navigator, **supporting 10k document reviews** for the Total Rewards team
- **AI Pipeline Enhancement:** Improved text extraction accuracy through enhanced prompt engineering & context-aware document chunking, **achieving 7% higher recall & 5% better precision**
- **System Monitoring:** Implemented comprehensive monitoring with **over 20 Datadog alerts & an analytics dashboard** to track system adoption, usage patterns & cost savings, **boosting operational reliability**

Advect Systems, Mumbai, India Jan 2022 – Aug 2024
Founding Software Engineer

- **Full-Stack Product Development:** Architected and built Nanofactory, a full-stack 3D printer management & automation platform using Svelte, Flask, & WebRTC, enabling **1k+ remote manufacturing workflows**
- **CI/CD & Automation:** Built comprehensive CI/CD pipelines with Docker, GitHub Actions, Cloudflare & Python Integrations, **reducing time-to-market by 40% & improving team productivity by 34%**
- **Test Coverage:** Built end-to-end & unit test suites with Playwright & PyTest, **achieving 80% code coverage** across critical application workflows

PROJECTS

Distributed Ride-Hailing Pipeline Nov 2025

- **Event Ingestion:** Built a real-time data pipeline by streaming driver location and ride events through Kafka producers into partitioned topics across distributed brokers, handling **100K+ messages reliably**
- **Stream-based Matching:** Implemented rider-driver matching logic using Samza stream processing jobs with a stateful KeyValueStore management across partitions to support low-latency decisions
- **Test Coverage:** Ensured reliability by applying test-driven development with unit tests for Kafka message routing and Samza stream processing logic, **achieving 100% code coverage** across critical components

Multi-Cloud Microservices Sept 2025

- **Multi-Cloud Failover & Autoscaling:** Developed Java-based microservices deployed across GCP & Azure with global failover & autoscaling, **maintaining high availability**
- **Automated Multi-Cloud Delivery:** Packaged services with Helm & automated release to Google Artifact Registry & Azure Container Registry using Docker & GitHub Actions, streamlining deployments