

Yash Thakkar

ythakkar@cmu.edu • yash-thakkar.com • linkedin.com/in/yashthakkar21/ • San Francisco Bay Area

EDUCATION

Carnegie Mellon University, Pittsburgh, PA

Graduation: December 2025

Master of Information Systems Management, **CGPA - 4.04**

Coursework: Cloud Computing, Distributed Systems, Object Oriented Programming, LLM Applications, Deep Learning

Teaching Assistant: Engineering Data-Intensive Scalable Systems, Data Science and Big Data, Agile Methods

Dwarkadas J. Sanghvi College of Engineering, Mumbai, MH, India

Graduated: May 2022

Bachelor of Engineering in Mechanical Engineering, **CGPA - 3.92**

EXPERIENCE

Newfront Insurance, San Mateo, CA

May 2025 - August 2025

Software Engineer Intern

- **AI-Powered Document Processing** - Decreased benefits plan data extraction and review turnaround by >90% and saved \$175k/year through scalable AI automation pipelines using FastAPI, GraphQL, PostgreSQL, Next.js, and OpenAI APIs
- **GraphQL Read Optimization** - Cut DB load and improved latency by enforcing read-path separation of concerns with tight selection sets and lean resolvers that minimize joins
- **Monitoring & Analytics** - Enhanced operational reliability by implementing 20+ Datadog alerts and an analytics dashboard to track adoption, usage, and cost savings, enabling data-driven product improvements
- **LLM Pipeline Optimization** - Improved recall by 7% and true positive accuracy by 5% by chunking extractions, building specialized OCR and extraction agents orchestrated in parallel with bounded concurrency using asyncio and semaphores

Carnegie Mellon University, Pittsburgh, PA

September 2024 - December 2024

Software Engineer

- **Frontend Development** - Boosted user engagement by 23% by developing core user-facing features end-to-end in collaboration with design and product teams using React, Redux, TypeScript, and Tailwind CSS
- **Accessibility Enhancement** - Expanded potential user base by 2.5x by integrating i18n for multilingual support

Advect Systems Pvt. Ltd, Mumbai, India

January 2022 - August 2024

Head of Software Engineering

- **Software Development Life Cycle** - Led the end-to-end development of Nanofactory, a 3D printer management app with a decentralized architecture using Svelte.js, Flask, and WebRTC
- **CI/CD & Automation** - Reduced time-to-market by 40% and improved team productivity by 34% by implementing a CI/CD pipeline and API integrations using Docker, GitHub Actions, Cloudflare, and Python
- **Testing & QA** - Lowered bugs by 13% by expanding E2E and unit test coverage with Playwright and Pytest

PROJECTS

Elastic Web Service on AWS

September 2025

- **Elastic Architecture** - Maintained high availability during spikes across multi-AZ by load-balancing traffic and scaling instances using AWS Application Load Balancer and Auto Scaling with health checks, graceful draining, and warm-up
- **Autoscaling & QoS** - Tuned scaling to reach target RPS 70% faster with instance-hours trimmed by 20% via CPU/RequestCount CloudWatch Alarms

AI-Powered Design Generation

November 2024

- **Generative AI Pipeline** - Reduced Amazon listing time by 39% by developing a platform that converts sketches into production-ready images, using custom diffusion models and fine-tuned LLMs
- **Multi-Model Inference Service** - Enabled efficient processing by containerizing and deploying a scalable FastAPI service on a GCP Compute Instance, integrating Stable Diffusion, LLaMA 3, and BART into a single system

SKILLS

Programming Languages: Python, Java, C++, TypeScript, JavaScript, SQL, Bash/Shell

Tools: GitHub, Linux, Docker, GraphQL, OpenAPI, FastAPI, Node.js, Next.js, Flask, Datadog, Pytest, Kafka

Infrastructure & Database: AWS (EC2, Lambda, S3, ELB, ASG, CloudWatch), PostgreSQL, Terraform, Hadoop, Spark

AI & ML: TensorFlow, PyTorch, Scikit-learn, Matplotlib, Pandas, NumPy, XGBoost, Hugging Face Transformers