

# Omkar Chhabra

562-353-0254 | [omkarchhabra38@gmail.com](mailto:omkarchhabra38@gmail.com) | [linkedin.com/in/omkar-chhabra](https://www.linkedin.com/in/omkar-chhabra) | [workwithomkar.com](https://workwithomkar.com)

## EDUCATION

---

### California State University, Long Beach

Bachelor of Science in Computer Science – GPA: 3.4

Long Beach, CA

Jan. 2022 – Dec. 2025

## EXPERIENCE

---

### Software Engineer

Adobe Inc.

May 2024 – Present

San Jose, CA

- Built and shipped a document intelligence feature using Python and LangChain with a RAG pipeline over internal knowledge bases, reducing average support escalation time by **35%**.
- Developed 6 REST API endpoints in Node.js integrated with PostgreSQL to serve structured metadata for Adobe's document processing pipeline, handling **10K+ daily requests**.
- Containerized 3 microservices using Docker and deployed to AWS EC2 behind an Application Load Balancer; set up GitHub Actions CI/CD pipeline reducing deployment cycle from 2 days to **under 2 hours**.
- Added Prometheus metrics and 4 Grafana dashboards to monitor service health, latency, and error rates across production environments.
- Participated in bi-weekly code reviews, sprint planning, and architecture discussions; resolved **12+ bug reports** across the codebase within agreed SLA windows.

### Computer Science Tutor

COE Tutoring Center, California State University, Long Beach

Aug. 2025 – Dec. 2025

Long Beach, CA

- Tutored 20+ undergraduate students in Data Structures, Algorithms, and Systems Programming each week.
- Developed review materials covering recursion, graph traversal, and time/space complexity; improved student exam pass rates by an estimated **15%**.

## PROJECTS

---

### ProfScore – Rate My Professor AI | Next.js, TypeScript, Tailwind CSS, LLM API, Pinecone

- Built an AI-powered professor rating app handling **100+ requests/min** with structured summaries and sentiment signals using LLM inference.
- Implemented Pinecone vector search with text embeddings for retrieval-based responses, improving query-to-review relevance over keyword search.
- Standardized a 3-part structured output (summary, pros/cons, recommendations) to reduce hallucination noise and improve response consistency.

### Beach Bank – Full-Stack Banking App | Next.js 14, TypeScript, PostgreSQL, Plaid, Dwolla, JWT

- Built **10+ backend API routes** for accounts, transactions, and transfers across 3 core banking flows with end-to-end input validation.
- Integrated Plaid and Dwolla for ACH transfers secured via JWT sessions and 2-step token verification, meeting financial data security standards.
- Designed a **6-table PostgreSQL schema** and tuned 8 queries with indexes and joins, achieving consistent sub-100ms read times.

### DevSecOps – Netflix Clone on AWS | Jenkins, Docker, Kubernetes, Helm, ArgoCD, Trivy, SonarQube, Prometheus

- Configured a **4-stage Jenkins CI/CD pipeline** with Trivy container scanning and SonarQube quality gates, blocking **15+ vulnerabilities** pre-deployment.
- Deployed to Kubernetes using 3 Helm charts managed via ArgoCD GitOps for declarative, repeatable releases with automated rollback.
- Added **5 Grafana dashboards** backed by Prometheus metrics to monitor uptime, pod health, resource usage, and deploy events in real time.

## TECHNICAL SKILLS

---

**Languages:** Python, Java, TypeScript, JavaScript, Go, Kotlin, SQL, C/C++, Bash

**Backend:** Node.js, Django, Spring Boot, REST APIs, JWT, OAuth, API Design

**Frontend:** React, Next.js 14, Tailwind CSS

**AI / ML:** RAG, LangChain, LLM Evaluation, Embeddings, Pinecone (Vector DB), Prompt Engineering

**Cloud & DevOps:** AWS (EC2, S3, IAM), Docker, Kubernetes, Helm, ArgoCD, GitOps, GitHub Actions, Jenkins, CI/CD

**Data & Observability:** PostgreSQL, MongoDB, Prometheus, Grafana, SonarQube, Trivy, Linux