# Rohan Deshpande

610-637-0667 | rbdeshpande@ucla.edu | linkedin.com/in/rohan-deshpande/ | github.com/rohand2290/ | **US Citizen**

## Education

**University of California, Los Angeles** — Los Angeles, CA
*Bachelor of Science, Mathematics/Computer Science* — *Expected June 2028*
**Coursework**: Data Structures/Algorithms, Operating Systems, Computer Organization, Linear Algebra, Probability

## Experience

**Roche** — June 2026 – Sep 2026
*Software Engineering Intern* — *Santa Clara, CA*
- Cut daily database cleanup from **30 to <2 minutes** across **8** tables by building a Spring Boot and Angular tool
- Automated deployments to **6** bare-metal clusters by replacing kubectl scripts with Helm and GitHub Actions
- Eliminated long-held DB connections for **15+** daily jobs by designing a Redis and RabbitMQ worker queue
- Removed separate frontend and backend deployments for **12+** engineers by packaging services in a Podman pod

**VetsEZ Inc.** — Sep 2025 – June 2026
*Software Engineering Intern* — *San Diego, CA*
- Enabled programmatic VistA access for **500+** users via **13+** OAuth-secured Spring/OData endpoints
- Cut latency **30%** across **1,600+** records by redesigning PostgreSQL schemas and caching with Spring Data
- Decreased deployment failures for **5+** services by automating Helm releases to EKS via GitHub Actions
- Reduced timeouts/access errors across **300+** API requests/month via KEDA autoscaling and Prometheus

**Drexel University** — July 2025 – Sep 2025
*Software Engineering Intern* — *Philadelphia, PA*
- Allowed **12+** researchers to explore data without manual SQL via a VPS-hosted Next.js/PostgreSQL app
- Cut graph generation from **30 to 15 minutes** for **150+** nodes and **300+** edges via FastAPI and Neo4j
- Unblocked analysis of **9000+** proteins by cleaning data into Cytoscape-ready CSVs and a Jupyter notebook

## Projects

**ProteinLens** | *Spring Boot, Kafka, Neo4j, Docker, AWS, Kubernetes*
- Validated **70%** of detected clusters against known findings on **50+** node/**100+** edge PPI graphs using EJML
- Processed **25+** graph jobs/minute in order via a job-keyed, 3-partition Kafka topic and Spring workers on EKS
- Improved pathway discovery via applying laplacian construction and eigenvector embedding over STRING DB

**SchemaLoom** | *TypeScript, Zod, Docker, Hono, LangChain, Gemini*
- Converted research/Wikipedia articles into DAG schemas with a **70%** pass rate via Gemini/LangChain
- Standardized releases across environments by deploying reproducible Docker images through GitHub Actions
- Caught API failures before release by automating **10+** endpoint tests with Jest and GitHub Actions

**CytoFlow** | *FastAPI, Java, py4cytoscape, Docker*
- Cut Cytoscape SVG processing from **5 minutes to 30 seconds** via FastAPI/py4cytoscape automation
- Caught SVG-labeling errors before commit across **16+** runs via Git hooks and **8+** pytest checks
- Supported **15** concurrent requests on one server by packaging FastAPI in Docker with Pydantic validation

## Honors & Leadership

**William F. Sharpe Fellow** | *UCLA Alumni Association, Technology Vertical* — 2026–2027
**LA Hacks, Finance Coordinator** | *Southern California's largest hackathon* — 2025–Present

## Technical Skills

**Languages**: Java, Python, TypeScript, SQL, C++, Bash
**Backend/Web**: Spring Boot, Spring Data (JPA), FastAPI, Next.js, Angular, Hono, REST, OData, OAuth2, OpenAPI
**Data/Messaging**: PostgreSQL, MySQL, Neo4j, Redis, Kafka, RabbitMQ, gRPC
**Infrastructure**: AWS (EKS), Docker, Kubernetes, Helm, KEDA, Podman, GitHub Actions, Prometheus, Grafana
**Testing/Tools**: JUnit, Jest, Vitest, pytest, Postman, Git, Maven, Cypress