

Mehul Mangave

San Jose, CA | +1 (408)210-8953 | mehul.mangave@sjsu.edu | [linkedin.com/in/mehul-mangave](https://www.linkedin.com/in/mehul-mangave)

EDUCATION

San Jose State University

Master of Science in Computer Engineering

San Jose, CA, USA

Degree Expected: May 2026

Vishwakarma Institute of Technology

Bachelor of Technology in Electronics and Telecommunications Engineering

Pune, India

Degree Received: May 2023

Relevant coursework: Designing Data Structures, Algorithm Design, Operating Systems, Advanced Computer Design, Systems Software, Artificial Intelligence and Data Engineering.

SKILLS

Languages: Terraform, Python, C, C++, Java, SQL, NoSQL, PostgreSQL, OOP

Frameworks/Tools: Django, React, Node.js, Git, CI/CD, Jenkins, MongoDB, TensorFlow, PyTorch

Platforms: AWS, Azure, Kubernetes, Docker, Linux, Windows, GitHub Actions, Gitlab

AI/ML Tools: LLM / RAG Pipelines, Knowledge Graphs, Vector Search / Embeddings, Hugging Face Transformers

EXPERIENCE

State Street Bank and Trust, Boston, US

June 2025 – Aug 2025

Cloud Engineering Intern

- Owned the full software development lifecycle for an AI-powered internal chatbot using NLP to retrieve answers from Confluence docs, improving knowledge access and reducing support overhead.
- Engineered a **Terraform automation** proof of concept (PoC) that dynamically generates infrastructure templates based on user-defined requirements, utilizing in-house modules to streamline deployment workflows and reduce manual effort by over **60%**.
- Collaborated with the Service Enablement team to architect, develop, and document reusable Terraform modules for provisioning various AWS services, contributing to scalable infrastructure practices and adherence to organizational standards.

Syngenta

June 2023 – July 2024

Cloud Engineer

- Architected multi-tier **cloud infrastructure** that supported increased user demand; and optimized resource allocation, resulting in a 25% decrease in operational costs while maintaining high availability for over 10,000 users.
- Collaborated with senior engineers to migrate over **100 AWS** accounts to **Terraform Cloud**, streamlining infrastructure provisioning and management. This migration significantly improved cloud operations efficiency and saved considerable time.
- Collaborated with the applications team to deploy scalable, low-latency services using **Terraform**, cutting deployment time by 40% and achieving 99.9% uptime.
- Managed AWS **DynamoDB** migration from legacy to new AWS accounts, ensuring a smooth transition. Managed cloud-based **database migration tasks**, optimizing the process for efficiency and **minimal downtime**.
- Led the development of cost-saving strategies in AWS, optimizing underutilized resources and achieving **\$500k in annual savings** by effectively scaling down unnecessary resources.
- Implemented **Sentinel policies** for AWS and Azure services, integrating them with **Terraform Cloud**, which improved resource verification speed by 50%. This integration enhanced overall efficiency, reducing manual checks by 70%.

Syngenta

Aug 2022 – Jun 2023

Software Engineering Intern

- Provisioned cloud infrastructure for 20+ applications using **AWS** services such as **EC2, S3, RDS, DynamoDB, APIGateway, Lambda, ECS, ECR, and ELB** ensuring efficient deployment and scalability.
- Planned and implemented Terraform modules for AWS services that standardized infrastructure management across departments, now utilized by over 50 teams, providing a reliable framework for secure and efficient cloud resource management.
- Developed and documented a solution for AWS database services (DynamoDB, Redshift, DocumentDB, Timestream), which has been successfully adopted by teams, improving performance by 30% and enhancing data processing efficiency by 25%.

PROJECTS

Dynamic Weight Learning for Self-RAG, Team of 2

Aug 2025 – Dec 2025

- Developed a dynamic weighting framework for Self-RAG pipelines to improve relevance of retrieved context, addressing limitations of static weighting across retrieved documents during LLM generation.
- Implemented adaptive scoring algorithms using Python and PyTorch to re-rank and dynamically weight retrieved passages across **1000+ query evaluations**, improving contextual grounding and reducing irrelevant retrieval influence.
- Improved response accuracy and relevance in retrieval-augmented generation workflows by **~18% in evaluation experiments** while maintaining low inference overhead through optimized retrieval scoring logic.

Omnireasoner, Knowledge Graph Reasoning System, Team of 4

Aug 2025 – Mar 2026

- Built a knowledge-graph-augmented reasoning pipeline integrating graph traversal with LLM prompting to improve factual consistency in complex question answering.
- Designed entity linking and traversal workflows over a graph with **10K+ entities and relations**, generating interpretable reasoning paths and reducing hallucinated responses by ~15% during evaluation experiments.