

YASH CHETAN DOSHI

yash.doshi@tamu.edu | Mobile: +1 (979) 574-9820 | [LinkedIn](#) | [GitHub](#)

EDUCATION

Texas A&M University, Mays Business School

College Station, TX

M. S. in Management Information Systems

May 2027

Coursework: Adv Data Management Systems, Blockchain & Artificial Intelligence, Statistical Foundation for Data Science

University of Mumbai, D. J. Sanghvi College of Engineering

Mumbai, India

B. Tech in Information Technology with Honors in DevOps – GPA: 3.96/4

May 2025

Coursework: Artificial Intelligence, Database Management Systems, Data Warehouse and Mining, Business Analytics, Predictive Analytics, Big Data Analytics, Statistical Analysis, Cloud Engineering, Management Information Systems, DevOps, MLOps

EXPERIENCE

Full Stack Developer, Texas A&M University, College Station, TX

Jan 2026-Present

- Working on end-to-end development and maintenance of full-stack web applications used by students and faculty, from system design through deployment and ongoing support.
- Designed and maintained relational and NoSQL (MongoDB) databases storing student and faculty records, ensuring data accuracy, access control, and scalability.
- Developing teacher-facing app that enables faculty to securely view, manage, & analyze student data through dashboards
- Engineer automation workflows for notifications & reminders, significantly reducing manual coordination for admin staff.

Data Analyst Intern, Acma Computers Ltd, Mumbai, India, [Letter of Recommendation](#)

Jul 2024-Oct 2024

- Designed and maintained Python and SQL-based data ingestion and transformation pipelines supporting analytics and operational monitoring across system integration logs.
- Diagnosed pipeline bottlenecks and data quality failure modes, improving processing efficiency by 15% and contributing to a 10% reduction in system downtime.
- Implemented automated data validation, consistency checks, and reconciliation logic, improving operational data accuracy by 20%.

Software Developer Intern, C-DAC India, Mumbai, India, [Letter of Recommendation](#)

Dec 2023-May 2024

- Architected and built a virtual science simulator from the ground up, utilizing ReactJS for the front-end framework and Rive to engineer high-performance, lightweight vector animations, hosted on AWS EC2 for scalability.
 - Launched the interactive educational tool on the national 'Vidyakash' platform, where it successfully scaled to serve over 2 million students while ensuring high availability and a seamless user experience.
-

PROJECTS

[AetherMart E-commerce Data Platform](#) | AWS, MariaDB, Python, NoSQL, Vector Search, Tableau

December 2025

- Designed and built an end-to-end e-commerce platform using Python and SQL on AWS EC2 instance, integrating structured (MariaDB) and semi-structured (NoSQL) data sources to support intelligent commerce workflows.
- Implemented an AI-powered product discovery and recommendation system using vector embeddings and semantic search, enabling LLM-driven retrieval and ranking via the Gemini API.
- Engineered automated data ingestion, feature preparation, and inference pipelines creating ETL pipelines to support downstream analytics and real-time AI recommendations.
- Designed a highly available distributed database architecture using 3-node MariaDB Galera Cluster, reasoning about replication, consistency, and fault tolerance while enforcing security controls (RBAC, encryption).
- Validated model outputs ensuring reliable data driven insights, surfaced via Tableau dashboards for stakeholders.

[AI Agent Lab](#) | LangChain, LangGraph, CrewAI, Groq, Vector Databases

December 2025

- Designed and implemented multi-agent agentic AI systems using LangChain, LangGraph and CrewAI to orchestrate reasoning, retrieval, and task execution across complex, multi-step workflows.
- Built Retrieval-Augmented Generation (RAG) pipelines with vector databases and LLaMA-based models to enable context-aware, grounded responses over heterogeneous datasets.
- Analyzed agent behavior and failure modes, iterating on prompt structure, tool usage, and control flow to improve reliability and task completion accuracy.
- Integrated high-performance inference using Groq to reduce latency in agentic workflows & support real-time interactions.

[Interpretable Machine Learning in Healthcare – An XAI Approach for Diabetes Prediction](#)

16 June 2025

- Authored a paper on Explainable AI (XAI), developing techniques to enhance model transparency and predictive performance for diabetes prediction, presented at the international conference ICMAAI-25 (Publication forthcoming)
-

SKILLS

- **Languages:** Python, SQL, R, ReactJS, JavaScript, C, HTML/CSS, [Node.js](#)
 - **AI & ML:** Agentic AI, LLMs, RAG, LangGraph, LangChain, CrewAI, XAI, YOLO, Computer Vision, DevOps (Honors), MLOps
 - **Data Science & ML Libraries:** Pandas, NumPy, Scikit-learn, TensorFlow, OpenCV, Matplotlib, PyTorch, XGBoost
 - **Data:** PostgreSQL, AWS (EC2, S3, Lambda), GCP, Azure, CI/CD (GitLab), MySQL, MariaDB, MongoDB, Tableau, Power BI
 - **Certifications:** AWS Certified AI Practitioner, Google Cloud Career Practitioners, AI for Project Managers, n8n
-