

SIDDHESH MORE

mosiddhesh@gmail.com | linkedin.com/in/sidessh | github.com/sidesshmore | sidesshmore.com | +1 (480) 567-6235

Education

Arizona State University | *Master of Science in Computer Science, GPA: 4.0/4.0* Expected May 2027

University of Mumbai | *Bachelor of Engineering in AI & Data Science, GPA: 3.72/4.0* Aug 2021 – May 2025

Technical Skills

Languages/Technologies: Python, JavaScript/TypeScript, C/C++, SQL, HTML/CSS, Dart
Frameworks/Libraries: React, Next.js, Node.js, Django, Flask, FastAPI, Flutter, TensorFlow, LangChain, Streamlit
Databases/Cloud: PostgreSQL, MongoDB, Redis, Firebase, AWS (Lambda, S3, EC2, RDS), Docker, Kubernetes
Tools: Git, GitHub Actions, Postman, FAISS, Supabase, Render, Vercel, PostHog
Coursework: Data Structures & Algorithms, Machine Learning, Deep Learning, Database Management Systems, Distributed Systems, Cloud Computing, Natural Language Processing

Work Experience

CIS Research Aide | *Arizona State University* Dec 2025 – Present

- Built production **FastAPI** microservices with **RAG pipeline** using LangChain and **FAISS** vector database, deployed on **AWS Lambda** with **CI/CD** via GitHub Actions to handle 500+ concurrent queries daily.
- Engineered chat-history retrieval system with **PostgreSQL** backend implementing **connection pooling** and **Redis caching** layer, reducing average query latency from 800ms to 320ms (60% improvement).

Software Development Intern | *Ayurogyu Saukhyam Foundation* Sept 2024 – Mar 2025

- Created cross-platform mobile application using **Flutter** and **Supabase**, implementing **OAuth 2.0 authentication** with Google Sign-In and **real-time data synchronization** serving 300+ users across rural India.
- Designed accessible UI with **multi-language support** (English, Hindi, Marathi) and **voice navigation** using Flutter **TTS/STT**, reducing onboarding completion time from 8 minutes to 5 minutes (40% improvement).

Python Development Intern | *Digibranders Private Limited* Jun 2023 – Aug 2023

- Developed **RESTful APIs** for ERP system using **Django REST Framework** with **Redis caching** and **Celery** for async task processing, handling 50K+ daily requests with average response time of 60ms.
- Optimized PostgreSQL queries by adding **composite indices** on frequently-joined columns and refactoring **Django** ORM queries to use **select_related/prefetch_related**, reducing report generation time from 45 seconds to 6 seconds.

Projects

Text2SQL System - State-of-the-Art on Spider 1.0 | *Python, PostgreSQL, FAISS, LangChain* [Link](#)

- Architected **dual-similarity retrieval system** using FAISS vector database with sentence-transformers embeddings, combining **semantic search** with **keyword matching** to achieve 91.8% execution accuracy on Spider 1.0 benchmark.
- Implemented **query decomposition** using NatSQL intermediate representation and **self-correction mechanism** with GPT-4, outperforming MiniSeek (91.2%) and DAIL-SQL (86.6%) by 5.2 points while reducing syntax errors by 73%.

Referrly - Full-Stack Referral Platform | *Flutter, Next.js, Express.js, PostgreSQL, Docker* [Link](#)

- Launched full-stack job referral platform with Flutter mobile app (1,800+ users) and Next.js web application (3,600+ visitors), scaling to 200 monthly active users with **99.9% uptime** and **automated deployment pipeline**.
- Established **microservices architecture** with Express.js backend, PostgreSQL database, and Dockerized deployment on Render, implementing **JWT authentication**, **role-based access control**, and **real-time job matching algorithm**.

Aqua Trace - Water Footprint Tracking App | *Flutter, TensorFlow Lite, Firebase* [Link](#)

- Published Android app on Google Play Store with 1,000+ downloads across 50 countries, achieving **120ms image classification** and **local SQLite caching** for offline functionality.
- Trained **MobileNetV2 CNN** on custom dataset of 6,000+ food images across 42 categories, achieving 91% classification accuracy with 120ms inference time using **TensorFlow Lite** on-device model.

Publications

ExamAce: Leveraging LLMs and RAG for Intelligent Academic Assistance | *IEEE ICCNT 2025* Jul 2025

- Deployed a **vector datastore-backed QA system** using **FAISS** and **nomic-embed-text** embeddings with **Binary and Scalar Quantization**, enabling fast, hallucination-resistant answers grounded in curated engineering reference books.
- Constructed a **multi-agent RAG pipeline** (Llama 3.1, Llama 2, Mistral) with a novel **confidence score metric** achieving **0.65 cosine similarity** and **91% accuracy**, reducing hallucinations in LLM-generated academic responses.