VEER SANGHAVI

Tempe, Arizona | 602-820-9860 | veersanghavi8@gmail.com | linkedin.com/in/veer-sanghavi | github.com/veer64

EDUCATION

Arizona State University — Tempe, Arizona
Masters in Data Science, Analytics and Engineering (Computing and Decision Analytics) — Aug 2024 - May 2026

Vellore Institute of Technology — Vellore, India
Bachelor of Technology in Computer Science and Engineering — Jul 2020 - Jun 2024

SKILLS

Skills: Data Science, Data Analytics, A/B Testing, Machine Learning, LLMs, VLMs, Agentic AI, Gen AI, RAG, NLP, Deep Learning, Statistical Modeling, Predictive Modeling, Computer Vision, Communication

Programming Languages: Python, SQL, R, C/C++, Java, HTML, CSS

Libraries and Frameworks: scikit-learn, XGBoost, NLTK, spaCy, Transformers (Hugging Face), TextBlob, TensorFlow, Keras, PyTorch, OpenCV, Pillow (PIL), LangGraph, LangChain, Hadoop, Apache Spark, PySpark

Tools and Technologies: Jupyter Notebook, Excel, AWS (SageMaker, EC2, Bedrock, Redshift, S3), GitHub, VS Code, Copilot, Jira

Certifications: Oracle Machine Learning Certified Associate, Udemy (Statistics for DS and Business Analytics, SQL for DS, ML A-Z)

PROFESSIONAL EXPERIENCE

New York Life — New York, New York
Data Science Intern — May 2025 - Aug 2025
Multi-Modal Retriever Design - VLMs, Titan Embeddings & RAG Evaluation

- Built a hybrid retriever for RAG by combining ColQwen2 (vision) and Amazon Titan V2 (text) embeddings, improving recall by 15% with no drop in precision through score interpolation and grid search over Top-K and weights.
- Developed a similarity scoring function for 2D vision embeddings, introducing normalization + adaptive thresholding that boosted precision, improved stability and also handled edge cases.
- Enhanced retriever robustness, reducing response latency by ~0.8s per query (20% faster), automated precision/recall telemetry, and version-controlled benchmarking, ensuring reproducibility and continuous evaluation at scale.

Tech Mahindra, Makers Lab — Pune, India
Artificial Intelligence Intern — Jan 2024 - Jun 2024
AI Generated Face Image Detection System

- Formulated a unique system integrating StyleGAN2 inversion with tailored image comparison tools, improving accuracy for differentiating real from AI-generated images and addressing critical gaps identified during research efforts.
- Extracted latent vectors from real and fake images, regenerated images for comparison, and selected most impactful similarity metrics based on feature importance and accuracy.
- Leveraged advanced ML algorithms, including XGBoost and ANN, to process outputs from StyleGAN2 inversion and image similarity metrics, attaining near 99% accuracy in differentiating fake images from real ones.

Tredence — Remote
Data Science Intern — Dec 2022 - Apri2023

- Collaborated with interdisciplinary teams to classify complex medical terminology into actionable insights, leading to improvements in three key areas by resolving major project issues.
- Employed systematic approaches such as tokenization, POS tagging, lemmatization, and vectorization on over 5 million words from medical texts, refining overall accuracy of subsequent classification algorithms in the NLP model.
- Executed Sequential Forward Selection (SFS) for feature extraction, enabling identification of 15 crucial features from a dataset of over 5000 records; attained 80% model accuracy.
- Link to Publication

PROJECTS

RAG ChatBot for VIT Vellore

- Developed a Retrieval Augmented Generation (RAG) Chatbot for VIT Vellore with Langchain, FAISS, and GPT-3.5, delivering 95% response accuracy in answering queries related to programs, events, admissions, etc.
- Automated common inquiries, reducing administrative response times by 40% and enhancing student engagement.

Drowsiness Detection System – Link to Publication

- Designed an advanced driver safety system analysing real-time video data, achieving a 90% detection accuracy for facial and eye status using OpenCV and Haar Cascade Classifier

Data Sharing Based on Facial Recognition Clusters – Link to Publication

- Engineered an image clustering framework to automate picture delivery, reducing processing time by 30%
- Exceeded 98% accuracy in face matching under varying conditions (lighting, pose, expressions).