

MOHAMMED SAMEER SYED

Tucson, Arizona

+1-520-250-6750 [Portfolio](#) [GitHub](#)

AI Engineer & Researcher Building Models, Startups, and Breakthroughs in Machine Learning

Summary

AI Engineer and Researcher with 3+ years of experience in applied machine learning and deep learning. Focused on building healthcare AI systems, multi-agent architectures, and multimodal LLMs. Experienced with architecting end-to-end pipelines and deploying production-grade AI solutions specializing in scientific AI modelling. Proven record of co-founding startups, publishing research, and filing patents at the intersection of AI, health, and science.

Experience

Roshan AI

Founder

February 2026 – Present

Tucson, Arizona, USA

- Developing **ShifaMind**, a novel interpretable clinical AI system outperforming frontier and baseline models on ICD-10 coding, incorporating a **Multiplicative Concept Bottleneck (MCB)** architecture with generative design extensions.
- Leading end-to-end research and productionization of clinical NLP systems; coordinating with academic collaborators and preparing research outputs for top-tier venues.

Quelea Social

Co-founder / Lead AI Engineer

October 2024 – Present

Singapore (Remote)

- Designed and trained a multimodal multi-agent video understanding system using OpenAI agents, OpenCV, and HuggingFace Transformers; built a **NextTGPT**-inspired model trained on 5k+ influencer videos achieving SOTA performance, leveraging a **mixture of experts (MoE)** for model scaling; prepared for arXiv publication.
- Fine-tuned and deployed **Qwen 2.5 Vision** on 7k data points using **QLoRA** and self-supervised learning on RunPod, significantly improving overlay text extraction accuracy in production pipelines.

University of Arizona

Graduate Research Assistant

January 2025 – February 2026

Tucson, Arizona, USA

- Designed and architected end-to-end **ML pipelines** integrating RAG, multi-agent LLM systems, and knowledge graphs across the full data lifecycle ingestion, preprocessing, training, and cloud deployment on **AWS** establishing scalable system architecture bridging siloed data environments into accessible, digitized information flows.
- Developed digital phenotyping pipelines by integrating mobile-sensing platforms (**MindLamp**, bug fixes and data integrity enhancements) and deploying **RADAR-Base** on AWS EC2 for continuous ingestion of sensor, behavioral, and survey data; built knowledge graphs, KNN clusters, and clinical text embeddings for behavioral modeling.
- Engineered a multi-agent architecture unifying wearable data, passive smartphone signals, clinical text, and survey streams into a real-time analytics framework.

Lumenci

Software Engineer

January 2024 – December 2024

Gurgaon, Haryana, India

- Architected and deployed an automated **ML-driven portfolio analysis system** for financial/technology IP data, designing end-to-end pipelines that reduced consultant effort from 40–50 hours per project to minutes; engineered feature extraction workflows and scalable data pipelines to digitize previously analog processes.
- Accelerated large-scale portfolio analytics by **82%** using GCP GPU infrastructure; designed and built a multi-agent AI workbench integrating a **mixture-of-experts (MoE)** architecture, improving efficiency by **70%** and boosting model accuracy from 6% to **72%** tracking KPIs and ROI throughout deployment.

Indian Institute of Technology, Indore

Research Intern

January 2023 – October 2023

Indore, India

- Developed a **Graph Neural Network (GNN)** architecture from scratch on a decade of imagery from NASA's Solar and Heliospheric Observatory (SOHO) satellite, achieving **91% accuracy** in time series prediction.
- Led a team of 3 researchers and experimented with CNN, RNN, GAN, and GNN models for time series forecasting and image classification, selecting optimal architectures based on performance metrics.

Education

University of Arizona

Master of Science in Information Science: Machine Learning **GPA: 4.0/4.0**

January 2024 – December 2025

Tucson, Arizona, USA

- Distinguished Graduate Scholars list
- Recipient of **Best Graduate Capstone Award** and **REAP Grant** Awardee for clinical AI research.

National Institute of Technology

Bachelor of Technology in Electronics and Communication Engineering

June 2019 – May 2023

Srinagar, India

Projects

ShifaMind Multiplicative Concept Bottleneck for Interpretable ICD-10 Coding | *MCB, BioClinicalBERT, XAI Site*

- Built an interpretable clinical AI system on a **Multiplicative Concept Bottleneck (MCB)** architecture: 160 learnable concept queries attend via cross-attention over **6,144-token** discharge notes using **BioClinical-ModernBERT** (153M params), enforcing concept-mediated predictions with zero residual bypass guaranteeing faithful explanations by construction.
- Achieved **Macro-F1 of 0.7122** on 113,918 MIMIC-IV admissions (50 ICD-10 codes), outperforming LAAT (0.7114), KEPT (0.6870), and frontier LLMs (GPT-5.4: 0.4122), while being the **only model providing full concept-mediated interpretability** in the comparison.
- Validated interpretability via three architecturally grounded metrics: **CSTPR 0.7042** (vs. Vanilla CBM 0.1472), **CIM 36.5× larger gradient norms**, and **CCR 0.8361** (vs. 0.3614).
- Recognized with University's **REAP Grant** and **Best Graduate Project**; submitted to **NeurIPS 2026**.

IonTheFold Electrostatic-Aware AI Protein Folding | *ESM-2, APBS, ProteinMPNN, Deep Learning Website*

- Co-founded a computational biology research initiative reimagining **protein and antibody design** with AI for drug discovery, collaborating with Grinter Lab (University of Melbourne) and Knott Lab (Monash University).
- Designed a fusion architecture combining **ESM-2, APBS, and ProteinMPNN** with a novel **electrostatic-aware loss function**, achieving a **6.53%** improvement on charged proteins and a $4.14\% \pm 0.026$ overall gain in sequence recovery, surpassing SOTA methods across diverse protein families.

COMET-SEE Sungrazing Comet Detection System | *EfficientNet-B0, Difference Imaging, Solar Physics Demo*

- Built an AI system to automatically detect sungrazing comets in NASA SOHO/LASCO coronagraph images using **difference imaging** and an **EfficientNet-B0** classifier; engineered a full data pipeline (FITS ingestion, motion-enhancing difference frames, max-projection, alignment) processing **25k+ images across 665 sequences**.
- Deployed a **97.7% accurate** real-time Gradio app on HuggingFace Spaces with ~ 2 s inference per sequence.

BioSphereAI | *Streamlit, FastAPI, RAG, Multi-Agents, 3D Rendering, Blender, HuggingFace Devpost*

- Developed an AI-powered ecological explorer with real-time environmental data, RAG-based multi-agent insights, and interactive 2D/3D simulations of ecological systems with adjustable parameters (CO₂, humidity, temperature); integrated automated image and 3D model generation pipelines with Blender.

Skills

Languages: Python, C++, C, JavaScript, SQL, Bash, HTML, CSS

Cloud & DevOps: AWS (EC2, S3), GCP, Azure, Docker, Kubernetes, RunPod, Firebase, Supabase, CI/CD

Machine Learning: Linear/Logistic Regression, SVM, Decision Trees, Random Forests, PCA, Anomaly Detection, Time-Series Forecasting

Deep Learning & Models: CNNs, RNNs, GANs, GNNs, Transformers, Cross-Attention, Concept Bottleneck Models, EfficientNet, CLIP, ESM-2, ProteinMPNN, ModernBERT, ARIMA

AI Systems & Agentic Frameworks: Multi-Agent Systems, RAG, Prompt Engineering, Fine-tuning (QLoRA/PEFT), Explainable AI (XAI), Hallucination Mitigation, LLM Inference, API Orchestration (OpenAI, Anthropic, HuggingFace, Llama, Qwen), MCP

Data & Tools: PyTorch, TensorFlow, Keras, scikit-learn, XGBoost, Langchain, NumPy, Pandas, OpenCV, MediaPipe, Pinecone (Vector DB), MongoDB, FastAPI, Flask, Streamlit, Gradio, n8n, Langflow, Ollama, Crew AI

Soft Skills: Leadership, Technical Ownership, Problem Solving, Cross-functional Collaboration

Research Output & Achievements

ShifaMind: Paper submitted to **NeurIPS 2026** Multiplicative Concept Bottleneck for Interpretable ICD-10 Coding (Macro-F1 0.712 with full concept-mediated interpretability)

ASYM: Paper submitted to **EMNLP 2026** Safety Asymmetry Score in Tool-Using LLMs; channel-specific vulnerability metric, mechanistic findings via causal activation patching.

SpectralGAN + FRATransformer: Submitted to **IEEE Transactions on Power Delivery 2026**

SpectralGAN-Augmented Transformer Neural Network for Power Transformer Winding Fault Diagnosis via Frequency Response Analysis; 85.7% accuracy (+23.8 pp over SVM baseline) under strict 21-fold LOOCV

IonTheFold: Forthcoming arXiv preprint (Q3 2026) on AI-driven electrostatic-aware charged protein and antibody design, in collaboration with University of Melbourne and Monash University

Patent: First Inventor on a patent (Publication No. 47/2025; Application No. 202511081185, published Nov 21, 2025) for an autonomous self-landing rocket control system integrating a custom flight computer (ARCHER 1.0), real-time sensor fusion, PCB design, embedded firmware, and MATLAB/Simulink-based simulation pipelines.