SADIQ MOHAMMAD

@ sadiq.build@gmail.com +91 7416414823 md-sadiq-md MdSadiqMd

WORK EXPERIENCE

Dognosis - R&D Engineer

October 2025 - Present Bangalore, India

- Solving **cancer prediction** by leveraging dogs olfactory capabilities, delivering accurate diagnostics from video-based sniff analysis
- Built **multi-stage video processing pipeline** YOLO-based dog tracking and ROI detection, centroid interpolation at 60 FPS, and stabilized 256x256 follow-cam crop generation via FFmpeg for FERAL (V-JEPA 2) model inference
- Trained and deployed **V3 Fusion LSTM** architecture combining kinematic time-series (head position, velocity, acceleration) with scalar context features (sniff duration, session progress, tortuosity)
- Managed **Raspberry Pi camera fleet** for automated daily session capture by debugging streaming protocols, service orchestration, and multi-network connectivity for reliable field deployment
- Managed cloud and on-prem infrastructure by automated data ingestion, processing, and backup workflows across distributed storage and compute environments
- Built **multi-modal sensor synchronization** matching high-frequency biosignals with video timelines, enabling time-aligned feature extraction for model training pipelines
- Overall worked with Python, OpenCV, YOLO, Airflow, TimescaleDB, PostgreSQL, Linux, FFmpeg, Docker, AWS, Rust

Scale.jobs - Founding Team Engineer

September 2024 - October 2025 Hyderabad, India

- **Only developer** in the company **besides the CTO**
- Wrote AI content generation pipelines using LangChain and OpenAI SDK, while integrating Langfuse for prompt versioning, tracing, evaluation, and production monitoring
- Built **AI Interviewer** that conducted conversational mock interviews different for each and every job role, generated personalized feedback
- Optimized frontend performance using WebAssembly and TanStack, significantly improving rendering speed and overall user experience for large interactive applications
- Overall worked with TypeScript, tRPC, Hono, Prisma, PostgreSQL, QStash, LangChain, Langfuse, OpenAI SDK, TanStack, WebAssembly (WASM), and PostHog

HCL - Tech Intern

June - September 2024 Remote

- Worked in **Resume Application Tracking System (ATS)** team
- Applied **TF-IDF** for keyword density analysis and relevance scoring
- Built OCR-based document processing and text extraction pipelines
- Overall worked with Python, TensorFlow, Scikit-Learn, NLP, TF-IDF, OCR, and document processing

SKILLS

ML & Research: Python, PyTorch, Hugging Face Transformers, Mechanistic Interpretability, Reinforcement Learning, LoRA/PEFT, Computer Vision
ML Infra & DevOps: AWS, Docker, Kubernetes, MLflow, Airflow, Linux, CI/CD

PROJECTS

Circuit-Steering

- Fixes a model's behaviour by changing only its circuit, then audits whether anything else broke
- Pre-registered 393-run study on Tiny Aya Global (5 languages) comparing circuit-restricted steering, LoRA and MEMIT against random, lowest-scoring and unrestricted controls
- Found circuits over 612 heads and MLPs with EAP-IG, validated against random baselines: fact recall is localized, refusal is distributed
- Circuit LoRA cut harmful answers from 12.4% to 7.1% outside English at 1.02$\times$ perplexity, but a same-size random-component LoRA matched it ($p = 0.48$): the circuit was not the cause
- Built a 5-part side-effect audit: 31 of 92 edits passing standard perplexity/MMLU checks still shifted another metric by $\geq$10 points (vs. 2 of 68 noise controls)
- Technologies Used: Python, PyTorch, CircuitKIT, SafeTune, AuditKIT, PEFT

Non-linear Steering

- Steers a frozen language model toward a target behavior without drifting far from its original outputs
- Trained a nonlinear, norm-bounded residual-stream intervention under a strict forward-KL divergence budget
- Designed a two-pass policy-gradient estimator (on-policy rollout + teacher-forced replay) combining a frozen causal probe path with a REINFORCE term for unbiased gradient estimates
- Enforced the KL constraint as a Lagrangian with a dual multiplier, verified against exact gradients on a tiny model (cosine > 0.99)
- Technologies Used: Python, PyTorch, TransformerLens, HuggingFace Transformers

Embodied-Vision

- Sorts egocentric video frames by how hard they are for hand-tracking models
- Built a difficulty-aware classifier that inverts model confidence to triage frames into 5 classes
- Combined a two-pass HybridSampler, a dual-threshold MediaPipe ensemble with CLAHE, and hand-coded occlusion/dexterity/lighting scores
- Bidirectional temporal bridging separates occlusion from no_hands using temporal continuity, fixing BlazePalm's single-frame blindness and raising occlusion recall from 3.6% to 25.6%
- Technologies Used: Python, MediaPipe, OpenCV, NumPy

EDUCATION

Bachelor's of Technology (B.Tech) - Electronics and Communication Engineering

Velagapudi Ramakrishna Siddhartha Engineering College

2021-2025 **CGPA:** 8.00