ADITHYA N MURTHY

AI / Machine Learning Engineer — Computer Vision, 3D Vision, LLM Systems, Robotics

adithyanmurthy@gmail.com | github.com/Adithyanmurthy | linkedin.com/in/adithya-n-murthy | devpost.com/adithyanmurthy

adithyanmurthy.com · Bangalore, India · Available immediately · Remote or on-site

Final-year Computer Science engineer building machine-learning systems end to end — live 3D reconstruction, video analytics, legal NLP and verified LLM agents, benchmarked on real data. Published at IEEE CONECCT 2026; first author of a paper submitted to NLLP @ EMNLP 2026; secondary track in robotics and PCB design.

EXPERIENCE

Research Assistant — RV University, Core Engineering — Jan – May 2026

ARI Sync — real-time local-first AI coaching engine · Guide: Prof. Monesh N

- Converted instructional video into live form correction, running the entire ML stack locally on Apple Silicon — two 7B LLMs, Whisper and MediaPipe in **under 10 GB with no cloud dependency**
- Engineered the real-time feedback pipeline: 30 FPS pose tracking, ~2.5 ms median form-check latency, ~1.9 s from error onset to audible correction
- Architected async ingestion on FastAPI, Celery, PostgreSQL and Redis; React front-end running MediaPipe in WebAssembly for client-side pose detection

Research Fellow — RV University, Academics & Research — Feb – May 2024

Quantum Sciences · Mentor: Bharath Manchikodi

- Completed a semester-long quantum computing fellowship, earning the **Microsoft Azure Quantum certification** and working hands-on with Azure Quantum services and Q# programming
- Studied quantum mechanics, gates and algorithms via the Linux Foundation course; authored a 50-page report on applications in cryptography and optimization

PUBLICATIONS

Explainable LSTM Framework for Ego-Contingent Trajectory Prediction in Multi-Agent Systems — Jul 2026

IEEE CONECCT 2026, Bengaluru — presented; IEEE Xplore indexing pending

- Conditional sequence-to-sequence multi-input LSTM encoder–decoder treating a hypothetical future ego-plan as an interventional input, moving motion forecasting from correlation toward causal reasoning, explained with SHAP
- Scaled training from 10 to all **850 nuScenes scenes (1.2M rows)** cutting ADE 13% (0.515 → 0.449) at 26.5 ms inference

When the Metric Measures the Bug: Silent Failure Modes in Diachronic Legal Concept Embedding — Aug 2026

First author — long paper submitted to NLLP 2026 @ EMNLP 2026; under revision

- Showed four silent embedding-pipeline defects that inverted a study's conclusion (**22 of 25 legal concepts** appeared to shift meaning before the fix, 1 of 25 after); pooling choice alone swings the verdict from 1/25 to 24/25

PROJECTS

Live3D — live metric 3D reconstruction from any camera — 2026

- Grows a metric 3D model live in the browser from a phone, drone or CCTV feed: Depth Anything 3 depth, verified loop closure (LightGlue), GTSAM Sim(3) pose graph, Open3D TSDF, calibrated confidence — **no cloud GPU**
- **108-run benchmark on 36 datasets** (TUM, EuRoC, KITTI): sub-second live delay on 34/36, camera-path error ≤0.52% on TUM/EuRoC; 12-minute 4K street test; privacy blurring, 29-class labelling, GLB export

GLaw / LawVerse — AI legal intelligence for Indian law · team lead, capstone — 2026

- Built a **231,908-document (666M-word)** legal corpus and a 115,537-case citation graph with good-law labels; a Semantic Movement Index over legal-RoBERTa embeddings tracks doctrine drift across decades
- Shipped a Flutter app, Next.js site and FastAPI gateway with LLM document review in 12 languages: 98.4% type accuracy, 95.8% serious-defect recall on a 60-case synthetic eval set; 590+ tests

One Camera Brain — turns any existing camera into a smart camera — 2026

- 21k-line edge-to-cloud video-analytics platform: rule packs, YOLOX/RF-DETR ONNX detection with tracking, VLM alert verification, signed evidence ledger and a virtual ONVIF camera for existing software
- Evaluated on **313 labelled real videos**: 98% precision/recall on intrusion, 42/42 tamper events, 9.8 ms per 1080p frame; 439 passing tests; cloud tier load-tested to 8,000 nodes

AI E & PAMO — verified AI agents and local guardrails for coding agents 2026

- **AI E:** built five role agents (finance, data, HR, QA, DevOps) on a shared runtime with approval gates, sandboxed execution and MCP tools; deterministic engines compute every figure and a second model family must agree before an answer is certified — the data agent scored **98.4% on 250 questions with zero wrong answers certified**
- **PAMO:** built a macOS daemon and Tauri app that answers coding-agent permission prompts using safety rules and a local LLM; across **3,699 real decisions** it held 258 risky actions (secrets, payments, irreversible operations) for review

Textegrity & VidyaTid — content-integrity NLP and NEET exam prep · vidyatid.in 2025 – present

- Led Textegrity, fine-tuning seven transformers (DeBERTa-v3, RoBERTa, XLM-R, Longformer, SBERT, a cross-encoder and 2.8B Flan-T5-XL) for AI-text and plagiarism detection, with a stacked meta-classifier reaching **97.9% held-out accuracy (F1 0.984)**. Grew VidyaTid from a buildathon prototype into a production-hardened NEET platform — 67k-line Flask backend on PostgreSQL/pgvector, Redis and Docker with a Next.js 16 front-end — serving a verified bank of **39,385 questions** in English and Hindi

CreeperX & RC Aircraft — robotics and aeromodelling 2024 – 2026

- Designed CreeperX's **142-component, 4-layer controller PCB** as code (SKiDL/KiCad) — three 5 A rails, reverse-polarity protection, battery telemetry, disarmed-at-boot servos and a JLCPCB fab package — with TFLite object detection on a Raspberry Pi 4. Designed and built a high-payload BLDC RC aircraft, competing in the Boeing National Aeromodelling Competition at IIT Kharagpur in 2024 and 2025

More projects — full list at adithyanmurthy.com 2024 – 2025

- **AI Raitha Mitra** — 38-class crop-disease CNN on 87,867 leaf images, 83% smaller as TFLite, with Kannada advice via Gemini; presented at AI SoCSE Day · **Auto Image Enhancer** — self-trained blind ×4 super-resolution (DIV2K + Flickr2K)

LEADERSHIP & AWARDS

- **Hackathon Organizer — MLX (Machine Learning Xtreme) 2025** · Devpost, Feb 2025: created and hosted a global ML hackathon with **259 registrants from 28 countries** and 31 submissions, leading 3 co-managers
- **Third Prize**, HPCC Hackathon 2024 (RV University × HPCC Systems) · Runner-up, Building Apps with Python, RV University · Boeing National Aeromodelling Competition, zonal rounds at Kshitij, IIT Kharagpur (2024 and 2025)
- Hackathons: NxtWave × OpenAI GenAI Buildathon, Meta × Scaler OpenEnv, CodeFury 7.0, CodeFest'25 CTF

CERTIFICATIONS & TRAINING

- Microsoft — Get Started with Azure Quantum · Linux Foundation — Fundamentals of Quantum Computing (LFQ101)
- IBM — Big Data Foundations Level 2 (Big Data 101, Hadoop 101, Spark Fundamentals I) · IBM — Containers, Kubernetes & OpenShift, with Containers & Kubernetes Essentials badge
- NPTEL — Affective Computing · NPTEL — Design & Implementation of Human-Computer Interfaces · OpenAI Academy × NxtWave GenAI workshop · RV University workshops: Digital Twinning (Unity3D), Machine Learning (GDSC)

TECHNICAL SKILLS

Languages	Python, C++, Java, JavaScript, TypeScript, SQL, Dart
ML & AI	PyTorch, TensorFlow, Hugging Face Transformers, LLM fine-tuning, RAG, LLM agents & MCP, ONNX Runtime, scikit-learn, SHAP
Vision & 3D	OpenCV, MediaPipe, YOLOX, RF-DETR, Depth Anything, GTSAM, Open3D, LightGlue, three.js
Engineering	FastAPI, Flask, React, Next.js, Flutter, PostgreSQL + pgvector, MongoDB, Redis, ChromaDB, WebRTC
Infrastructure	Docker, Kubernetes, Prometheus, GitHub Actions, Git, Microsoft Azure
Robotics	ROS 2, KiCad / SKiDL PCB design, ESP32, Arduino, Raspberry Pi, sensor integration

EDUCATION

B.Tech (Hons.), Computer Science & Engineering — RV University, Bangalore 2023 – 2027

Specialization in Artificial Intelligence & Machine Learning · Minor in Robotics & Automation · GPA 3.745 / 5.00