

YASHWANTH DEVAVARAPU

+91 7893889642 | yashwanthdevavarapu@gmail.com | Portfolio Link: yaswanth169.github.io

EDUCATION

KL University, India : B.Tech in Computer Science and Engineering (Hons.), CGPA 9.6/10

Aug 2021 – May 2025

RESEARCH

Research Collaborator – Stanford University (MS&E)

Jun 2026 – Present

- With **Prof. Ross Shachter**, designed experiments testing whether LLMs can execute decision analysis; discovered that response-field order alone changes the answer (\$60, \$70, \$82.35, \$67.50 with answer-first versus a stable \$76.00 with reasoning-first), a measurement artifact that affects how LLM reasoning is evaluated.
- Built a code-generation setup in which models solved decision problems exactly (128/128) where in-context reasoning failed, and showed that error feedback lifts GPT-4o from 10.2 to 23.0 of 29 tests but never to perfection.

Research Collaborator – Harvard Medical School / MGH

Apr 2026 – Present

- Lead a team under **Dr. Xiangrui (Taylor) Zeng** generating patient-specific 3D brain MRI volumes without paired data (PyTorch, diffusion models); Set the research direction by hypothesizing that gradient guidance would fail where adapter conditioning succeeds, and designed an oracle experiment to test this hypothesis.
- Traced residual blur to conditional-mean averaging over plausible anatomies, localized it to high-frequency wavelet subbands, and designed a wavelet-domain loss that improved sharpness; built the ADNI data pipeline and nnUNet-based segmentation evaluation pipeline, generating all second-phase results.

Research Collaborator – Carnegie Mellon University, Prof. Min Xu's Lab

Jan 2026 – Present

- MIRA (CVPR 2026 workshop): Evaluated 11 frontier VLMs on a benchmark of 713 real images and 3,529 question-answer pairs across 10 scenarios, finding that they scored 39–48% compared with 94% for humans, and that multi-turn dialogue degraded the performance of most models. (Advisor: **Yunbei Zhang**)
- SR-LoRA / CLAC (ICLR'27 submission, with Xingjian Li): built and ran the 480-compositions-per-method GPU benchmark on Stable Diffusion v1.5 (rank-4 LoRAs, DPM-Solver++, CLIP ViT-L/14) and the $\tau/\alpha/\delta$ and $N=3$ ablations; the ten-seed run showed our text-alignment gain was 4.3%, not the 6.8% seen with three seeds. CLAC raises identity consistency 0.839→0.915 (+7.6 pts) over the strongest baseline. (Advisor: **Xingjian Li**)
- TRACE and APISHIFT: Contributed to the development of verifiable reinforcement-learning environments, including environment design, verification logic, and reward mechanisms that score agent actions against ground truth rather than relying on a judge model.

Undergraduate Research Assistant – KL University

Aug 2022 – Mar 2025

- With Dr. K. Ashesh, first-authored an IEEE CONIT 2024 paper applying outlier analysis to credit card fraud detection

Additional Publications:

- Resolving Spatial and Semantic Conflicts in Multi-LoRA Composition via Softmax Routing and Cross-Adapter Attention Conditioning ICLR'27 (submitted)
- MVEB: Massive Video Embedding Benchmark EMNLP'26
- Through the Looking Glass: MIRA, a Mirror Reasoning Benchmark for Vision-Language Models CVPR'26
- EdgeMedBench: Benchmarking Small Language Models for Medical QA IC4AI'26
- ML-Based SDN DDoS Detection: Machine Learning-Enabled IDS CIMA'26
- Credit Card Fraud Detection Using Outlier Analysis and Detection CONIT'24
- TabuLM and KinYaEmbed: Tabular Pre-training and Sentence Embeddings for Kinyarwanda arXiv'26
- Reproducible, Cross-Platform Study of Order-Book Data Structures in a Low-Latency Matching Engine arXiv'26
- Physics-Informed Neural Network Analysis of Rotational Effects on Oscillatory MHD Flow.(submitted) Journal, Scopus
- TRACE: Verifiable RL Environment for Network Troubleshooting Agents In Review
- APISHIFT: Verifiable Agent Learning for API Contract Migration In Review

Research Contributions: Benchmarking, dataset construction, and evaluation frameworks for vision-language models, LLM agents, multi-adapter generative models, medical imaging, decision analysis, and verifiable reinforcement-learning environments.

WORK EXPERIENCE

Software Engineer – Barclays, Pune, India

July 2025 – Present

- Architected and built an 8-microservice platform (Java 21, Spring Boot, PostgreSQL, AWS) that centralizes vulnerability management, FinOps, database operations and infrastructure automation, adopted by 100+ applications and **50+** engineering teams
- Built LangGraph agents with MCP tool integrations that run security scans, detect CVEs, create Jira tickets and generate reports automatically, replacing manual triage across 50+ teams; recognized with the Director's Recognition Award for improving automation reliability and reducing user-prone errors.
- Built automation covering **15,000+** services and **4,000+** applications for container remediation, vulnerability detection and validation pipelines, cutting remediation time from days to under one hour.
- Designed a hybrid SES-Keycloak authorization model with dynamic role provisioning, giving **50+** teams self-service onboarding
- Developed an AWS FinOps module for multi-account cost analytics, anomaly detection, forecasting and savings recommendations; automated GitLab migration and CI/CD onboarding for 100+ repositories

Software Engineering Intern – Barclays, Pune, India

June 2024 – July 2024

- Built a full-stack cache monitoring platform (Spring Boot, Thymeleaf, Apache Ignite) tracking SDS and ESM cache data across distributed systems, with anomaly detection and automated email alerts on cache fluctuations; improved data retrieval speed by **30%**.

SKILLS

- **Languages & Core:** Java, Python, C++, C, SQL, Bash, Linux, Git, GitLab, Docker, Kubernetes, CI/CD
- **AI & Research:** LLMs, Agent Orchestration, RAG, Multi-Agent Systems, Model Benchmarking, AI Evaluation, Decision Analysis, PyTorch, Hugging Face, LangChain, LangGraph, CrewAI, Vector Databases
- **Systems & Cloud:** Spring Boot, FastAPI, Django, REST APIs, Apache Ignite, Kafka, Spark, PostgreSQL, MongoDB, ChromaDB, FAISS, AWS (Bedrock, SageMaker), OpenTelemetry, Microservices, Distributed Systems
- **Selected Projects:** C++ limit order book and matching engine (~3M orders/sec); Python/C++ statistical backtesting framework; RAG document QA and text-to-SQL agents.

SERVICE & LEADERSHIP

- **Technical Committee Member, IEM-HEALS 2026:** reviewed submissions in health science and technology.
- **Core Member, RPA Tech Club, KL University (2023–2025):** ran 12 major events and 40 internal workshops, training 400+ students in robotic process automation.

CERTIFICATES & RECOGNITIONS

- **Top 10, Class of 2025, KL University:** top 10 of the Computer Science and Engineering graduating class.
- **Director’s Recognition Award, Barclays:** Received the for improving automation reliability and reducing human error.
- **Meta PyTorch OpenEnv Hackathon Finalist:** top 100 teams from 31,000+ registrations and 800+ finalist teams.
- **Google DeepMind APAC Research Symposium 2026:** Selected Student Delegate, Bangalore, October 2026.
- **AWS Certified:** Solutions Architect; Cloud Practitioner.