Itikela Bhaskar

+91-9014036703 itikelabhaskar@gmail.com itikela-bhaskar itikelabhaskar

Education

International Institute of Information Technology (IIIT), Hyderabad Hyderabad, India
B.Tech + M.S. by Research in Computer Science and Engineering — UGEE Rank: 77 *Jul 2023 – May 2028*

SAHITTII Junior College
Higher Secondary Education: 95.4% *2022*

Viswabharati High School
Secondary Education: 99.16% *2020*

Experience

Undergraduate Researcher Apr 2025 – Present
Centre for Visual Information Technology (CVIT), IIIT Hyderabad

- Member of the Video-Language Group under Prof. Makarand Tapaswi
- Working on multimodal learning, video understanding, and vision-language models

Backend Developer 2025
Hustlr (Startup)

- Contributed to backend development for an early-stage startup
- Worked on API design, server-side logic, and integration with cloud infrastructure

Technical Skills

Languages: C, C++, Python, JavaScript, SQL
AI/ML Frameworks: PyTorch, TensorFlow, scikit-learn
Computer Vision: OpenCV, MediaPipe, Vision Algorithms
Systems & OS: Linux, POSIX, Multi-threading, Virtual Memory, Process Management
Networking: TCP/UDP, Socket Programming, Protocol Design, Packet Analysis (libpcap, Wireshark)
Web & Backend: React, Next.js, Node.js, React Native, Tailwind CSS
Databases & Cloud: PostgreSQL, Supabase, Google Cloud Platform (GCP), BigQuery
Tools: Git, Docker, Make, GCC, CMake, Redis, Streamlit

Hackathons (7× Winner)

1st Place — ISEC 2026 Student Data Science Challenge *Jan 2026*
Team MEGALODON — Software Defect Prediction

- Predicted faulty software modules from static code metrics without execution
- Built a hybrid pipeline: KNN-based duplicate detection + LightGBM/XGBoost/CatBoost ensemble; used GroupKFold for leak-safe validation and SHAP for interpretability

2nd Place — ANRF AISE Hack 2026 (IBM · ANRF · IIIT Hyderabad) *2026*
Team MEGALODON — Flood Inundation Detection from SAR & Multispectral Imagery

- Tackled flood vs. permanent-water disambiguation on only ∼79 labeled patches with severe class imbalance; optical data was cloud-occluded, requiring SAR-primary fusion
- Built a heterogeneous 5-model ensemble with dual-stream architectures separating SAR signals from topographic priors (DEM, GSWE); majority-voting final prediction
- Physics-aware fusion and explicit flood/water class separation reduced implausible predictions; ensemble robustness beat any single ConvNeXt model under small-data constraints
- Ranked 2nd among 2000+ participants across India; awarded **INR 1.25 lakh** prize

3rd Place — Lloyds Banking Agentic Hackathon *Dec 2025*
Team MEGALODON — Agentic Data Quality System on GCP

- Designed an autonomous multi-agent pipeline enabling non-technical users to validate financial datasets at scale
- Agents performed schema validation, anomaly detection, and auditable decision-making on GCP/BigQuery; mentored by Google and Lloyds engineers on ADK and cloud architecture

1st Place — Qualcomm Megathon *Oct 2025*

On-Device AI Image Editor

- Built a fully on-device generative AI image editor with zero cloud dependency, accelerated on the Snapdragon NPU using Qualcomm QIDK
- Optimized inference pipeline for power efficiency and latency on mobile hardware

2nd Place — StarHack (SUD Life) *Nov 2025*

YouMatter — Gamified Insurance Application

- Developed a gamified insurance app to improve user engagement and retention
- Designed an adaptive cross-platform UI using React Native and Expo

2nd Place — Qualcomm Megathon *Oct 2024*

Real-Time Pose Detection System

- Developed a real-time human pose detection system optimized for Snapdragon mobile hardware using Qualcomm QIDK
- Balanced inference speed, accuracy, and power efficiency for on-device deployment

Projects

Med Veda | *Android, Kotlin, Jetpack Compose, llama.cpp, MedGemma* 2025

- Built an on-device Android medical assistant powered by MedGemma 1.5 4B Multimodal (Q4_K_M GGUF) — 100% local inference, zero cloud egress, full PHI privacy compliance
- Integrated a custom llama.cpp backend with multimodal support for simultaneous text and X-ray image analysis; used SigLIP vision encoder for medical image understanding
- Features longitudinal patient record management (Room DB), background model download (∼2.8GB GGUF), and hardware-accelerated inference on Snapdragon NPU
- Submitted to the Kaggle Google MedGemma Challenge; fine-tuned with QLoRA (r=32) for SOAP note formatting and vernacular translation

Semantic XOR-QA: Cross-Lingual Question Answering | *Python, NLP, LLMs* 2025

- Built an end-to-end cross-lingual open-domain QA pipeline on XOR-TyDi QA, answering low-resource language queries using English retrieval sources
- Hybrid retrieval with BM25 + cross-encoder reranking; query-aware semantic subgraph pruning to compress context and reduce LLM noise vs. standard flattening
- Graph-guided multi-view answer selection combining graph facts, triplet evidence, and raw passage extraction for multilingual LLM (Qwen-2.5) generation
- Productized research notebook into a modular Python package with CLI for automated benchmarking and dataset evaluation

Demand Paging & Virtual Memory (xv6 OS) | *C, OS Internals* Mar 2025

- Extended xv6 with demand paging, FIFO page replacement, and per-process swapping
- Implemented `memstat()` syscall and managed swap space with atomic guarantees

Languages

Telugu (Native) · Hindi (Proficient) · English (Proficient)

Certifications

Mastering OpenCV with Python — OpenCV University
Google Cybersecurity Specialization
Google AI Essentials