

Akash Mehra

+91 7248740615 | akashmehra.aidev@gmail.com | linkedin.com/in/akash-mehra13 | github.com/itsakki10

PROFILE

AI/ML Engineer with hands-on experience building and deploying computer vision, deep learning, and LLM-powered systems using Python and PyTorch. Delivered production-ready REST APIs, Docker-containerised applications, and real-time inference pipelines across three end-to-end projects. Udbhav 2025 Hackathon (UKAIM) winner. With strong foundations in model training, evaluation, and cloud deployment. Graduating June 2026 — actively interviewing for full-time AI/ML Engineer roles.

TECHNICAL SKILLS

Languages: Python, JavaScript, SQL

ML / DL: Machine Learning, Deep Learning, Computer Vision, CNNs, NLP, Neural Style Transfer, Pose Estimation, Face Recognition, Biometric Authentication, LLMs

Frameworks: PyTorch, OpenCV, MediaPipe, dlib, AdaIN, VGG19, Flask, Streamlit

Databases: PostgreSQL, Supabase, SQLite

DevOps: Docker, REST APIs, Git, GitHub, Render, Netlify, Vercel

Model Evaluation: Feature Engineering, Data Preprocessing, Precision/Recall, Inference Optimisation

PROJECTS

NeuralCanvas — Neural Style Transfer System

PyTorch | VGG19 | Docker | Flask | REST API

GitHub | Live Demo

- Achieved 40% faster style-loss convergence across 500+ test image pairs by implementing VGG19 feature extraction with custom data preprocessing and normalisation pipelines.
- Optimized an end-to-end deep learning inference pipeline supporting dual-image inputs, configurable style intensity control, and real-time rendering with optimized forward-pass performance.
- Deployed production REST API on Render with zero-downtime availability by containerising with Docker and Flask/Gunicorn backend serving a responsive JS frontend.

RepSense AI — AI-Powered Gym Coach

OpenCV | MediaPipe | Llama 3.3 70B | NLP | SQLite

GitHub | Live Demo

- 91% joint-angle classification accuracy across 5 exercise categories, evaluated on a 2,000-rep annotated test set using OpenCV and MediaPipe pose estimation.
- Reduced user form-error rate by half (measured by rep-correction events per session) via automated repetition counting, posture correction, and NLP-driven voice coaching with Llama 3.3 70B.
- Delivered 100% session data retention across multiple test sessions using a SQLite-backed workout history system; deployed multi-modal AI pipeline on Netlify.

VisionVox — Multimodal Biometric Authentication System

dlib | Face Recognition | PostgreSQL | REST API

GitHub | Live Demo

- Attained 92% face recognition accuracy (precision/recall on 750+ sample dataset) using dlib facial and speaker embeddings for real-time biometric identity verification.
- Built a cloud-based attendance platform with real-time face recognition and voice-assisted verification, reducing manual processing using Streamlit and PostgreSQL via Supabase.
- Scaled biometric record storage across concurrent sessions with zero data conflicts using a structured PostgreSQL schema via Supabase for secure cloud storage and retrieval.

EDUCATION

Veer Madho Singh Bhandari Uttarakhand Technical University

Dehradun, India

B.Tech — Artificial Intelligence & Machine Learning

2026

Shivalik Holy Mount Academy , Kashipur

U.S.Nagar, India

Class XII (CBSE): 84.4%

2022

Class X (CBSE): 89.3%

2020

ACHIEVEMENTS & LEADERSHIP

- **1st place** at the **Udbhav 2025 Hackathon (UKAIM)** as part of Team UnlockEd, outperforming state-level competitors in a competitive AI and technology challenge.
- Honored at **Raj Bhawan, Dehradun** by the Hon'ble Governor of Uttarakhand in recognition of outstanding academic and technical achievement.
- Organized and led a college-level **Coding & Debugging Competition** for 100+ participants, designing algorithmic problem sets spanning data structures, graph algorithms, and machine learning, evaluated by a 5-member faculty panel.