Tuhin Acharjee

tuhinacharjee91@gmail.com | +91-7903067070 | linkedin.com/in/tuhin-acharjee-b55220222 | github.com/tuhin-07

Professional Summary

Machine Learning and Computer Vision Engineer with 3.5+ years of experience building real-time video analytics, edge AI, robotic vision, face recognition, and Python backend systems. Skilled in NVIDIA DeepStream, TensorRT optimization, PyTorch deployment, Vision-Language Model fine-tuning, FastAPI, PostgreSQL/pgvector, Docker, and GPU profiling, with measurable gains in throughput, latency, GPU efficiency, and automation.

Technical Skills

Languages: Python, C++
ML/DL: PyTorch, TensorFlow, CNNs, SFT, VLMs, Hugging Face Transformers, vLLM, llama.cpp
Computer Vision: NVIDIA DeepStream, OpenCV, TensorRT,, Object Detection, Tracking, Face Recognition, Age/Gender Estimation, Eye-Gaze Monitoring
Optimization/MLOps: Docker, Docker Compose, Triton Inference Server, CUDA, ONNX Runtime, Asynchronous Inference, Multi-Stream Video Analytics, Edge AI, NVIDIA Nsight Systems, NVIDIA Nsight Compute, Linux, Git
Backend/Data: FastAPI, PostgreSQL, pgvector, Vector Databases, REST APIs, NumPy, Pandas

Work Experience

Computer Vision Engineer Nov 2025 – Present
Adagrad AI *Pune, India*

- Designed and developed the complete real-time people analytics ML pipeline using NVIDIA DeepStream for two-camera edge-device deployments.
- Developed people detection, tracking, entry/exit counting, gender classification, and age estimation workflows, achieving 96% accuracy in gender and age detection through a custom MiVOLO plugin for TensorRT-optimized PyTorch inference.
- Built FastAPI backend services to receive ML engine payloads, process analytics data, and store operational events in PostgreSQL.
- Implemented face recognition storage and retrieval using InsightFace embeddings, pgvector-based vector databases, and KNN similarity search, reducing search latency from 19 ms to 10 ms.
- Implemented asynchronous inference support, increasing real-time multi-stream processing performance from 9 FPS to 25 FPS.
- Developed rule-based computer vision algorithms for stalled vehicle detection, wrong-direction driving detection, and poor-visibility detection across blur, adverse weather, and visual degradations.
- Fine-tuned VLM model using Supervised Fine-Tuning to detect road hazards including fallen objects, fire, and smoke from visual inputs.
- Profiled inference workloads with NVIDIA Nsight Systems and Nsight Compute, reducing GPU stall time from 1150 ms to 1 ms through bottleneck analysis and optimization.
- Developed InsightFace-based candidate verification, real-time seatbelt violation detection, and MediaPipe-based eye-gaze monitoring to track candidate attention and movement during driving tests.
- Containerized and deployed ML services, backend components using Docker across edge-device workflows.

Software Development Engineer Mar 2023 – Oct 2025
Rosa Technology *Jamshedpur, India*

- Designed and developed an automated robotic rebar testing system, improving operational efficiency by 60% over manual workflows.
- Programmed robotic arm automation for pick-and-place operations and integrated automation logic for autonomous rebar handling during testing.
- Developed CNN-based computer vision pipelines for rebar localization, elongation measurement, breakage detection, and quality assessment.
- Implemented client-defined formula-based decision logic and integrated machine learning, computer vision, and automation components using Python, OpenCV, and TensorFlow.

Projects

Football Analysis System
Python, YOLOv8, OpenCV, K-Means, Optical Flow

- Developed an end-to-end football analytics pipeline for automated player and ball tracking from broadcast video footage.
- Implemented YOLOv8-based real-time player and ball detection across video streams.

Education

RVS College of Engineering and Technology 2019 – 2023
B.Tech in Computer Science and Engineering *CGPA: 8.5/10*