

Hamza Aziz

AI Developer @Global InfoVentures Pvt. Ltd.

hamzaaziz822@gmail.com | +91 8447644188

Sector-23, Sanjay Nagar, Ghaziabad-201002 (U. P.)

LinkedIn | GitHub | Portfolio | Medium



PROFESSIONAL SUMMARY

AI Engineer with 2 years of experience in Computer Vision, Generative AI, and MLOps. Built 4 production grade facial recognition systems across education and enterprise sectors, reducing manual effort by 90%. Achieved >99.99% recognition accuracy via real-time pipelines and automated Milvus DB updates. Developed scalable voice to SQL assistants using Rasa and Gemini LLM. Proficient in GPU inference (Triton, TensorRT, DeepStream) and CI/CD deployments.

SKILLS

- **Programming Languages:** Python, Bash, SQL
- **AI/ML Frameworks:** TensorFlow, PyTorch, InsightFace, YOLO, Scikit-learn, OpenCV, NumPy, Pandas
- **Generative & Conversational AI:** Rasa, Gemini 2.0, Large Language Models (LLMs), NLP, Prompt Engineering, LangChain
- **Model Optimization & Inference:** TensorRT, ONNX, DeepStream SDK
- **Deployment & MLOps:** CI/CD Pipelines, Data Pipeline Automation, Docker, NVIDIA Jetson (Nano, Orin Nano), DGX A100, RTX 4090, Linux, Git
- **Databases & Storage:** Milvus (Vector DB), PostgreSQL, Redis
- **Streaming & Messaging Systems:** Confluent Kafka, Flask-SocketIO
- **Other Tools:** SpeechRecognition, Tesseract OCR, Neo4j, VSCode, PyCharm, Postman

WORK EXPERIENCE

Assistant Engineer – AI

Global InfoVentures Pvt. Ltd. | August 2023 – Present

- Led development of 4 commercial-grade Facial Recognition Attendance Systems on DGX A100, RTX 4090, Jetson Orin Nano, achieving >99.99% accuracy
- Developed G6 Voice Assistant, converting natural language voice queries into precise SQL queries via Rasa, Gemini 2.0, and Milvus Vector DB; handled ambiguous queries with advanced intent disambiguation, achieving near-100% intent classification accuracy
- Mitigated domain shift in facial recognition using a CI/CD-style pipeline combining mobile and CCTV datasets — improved cross-domain recognition by 14%
- Automated nightly face vector updates from 8ft-high CCTV, boosting real-world accuracy from 85% to 99%
- Integrated PostgreSQL for recognition logs and built an API sync module with legacy SIM ERP Portal, enabling real-time attendance updates and SMS alerts for over 2,000+ students, enhancing parent engagement and system transparency

AI Intern

Global InfoVentures Pvt. Ltd. | February 2023 – July 2023

- Conducted 10+ hands-on AI/ML training sessions and mentored 50+ students across CSE, AIML, and Data Science domains
- Guided 10+ object detection projects using YOLOv8 and CNNs (VGG, ResNet, InceptionNet), resulting in successful full-cycle model deployments

EDUCATION

Birla Institute of Technology and Science (BITS), Pilani – WILP

2025 – 2027

Master of Technology (Executive)

- Artificial Intelligence and Machine Learning

Dr. A. P. J. Abdul Kalam Technical University, ABES Institute of Technology

Ghaziabad, UP

Bachelor of Technology

2019 – 2023

- Computer Science and Engineering

PROJECTS

G6 Voice Assistant – Natural Language to SQL Converter

Global InfoVentures Pvt. Ltd.

- Integrated SpeechRecognition, Rasa, and Gemini 2.0 to convert user voice input into structured SQL queries, enabling seamless access to ERP data from modules like Admission, Employee, Fee, and Library
- Achieved ~100% intent classification accuracy using Rasa by training with 50+ voice query variations per module. Devised a secondary intent resolution via contextual follow-ups for ambiguous queries (example, "Hamza Aziz" as student or employee)
- Formulated 100+ field-level schema mappings per module for Gemini and integrated Milvus Vector DB-based suggestions, reducing LLM hallucinations and boosting SQL query precision by over 30%
- Crafted a voice-adaptive summary module that transformed 1000+ SQL outputs into natural language, enhancing end-user experience in production deployment
- Solved heterogeneity in field values (example, department / designation name variations) using vector based field and value disambiguation, improving precision in SQL generation by over 40%

CCTV-Based Face Attendance System

Global InfoVentures Pvt. Ltd.

- Launched on NVIDIA DGX A100, orchestrated 20+ live CCTV streams using scrfd_10g_gnkps and w600k_r50 models (TensorRT Optimized)
- Pioneered dual dataset staging (mobile & CCTV) with Milvus Vector DB to mitigate domain shift, achieving over 99.99% face recognition accuracy across varied environments
- Architected a CI/CD-style data pipeline that extracts facial data from elevated-angle surveillance cameras, filters best-quality faces per person, and auto-updates vector database, improving recognition rates by 25% in practical deployment scenarios

PTZ-Classroom Attendance System

Global InfoVentures Pvt. Ltd.

- Automated attendance tracking by deploying rotating PTZ cameras and integrating with face recognition pipelines, eliminating manual roll calls and reducing proxy attendance by 25%

Edge-Based Office Attendance System

Global InfoVentures Pvt. Ltd.

- Commissioned Jetson Orin Nano for real-time edge inference using DeepStream SDK, Raspberry Pi AI camera, and proximity sensors to enable on-device facial recognition
- Eliminated manual attendance logging and replaced legacy biometric systems with 100% contactless AI-driven tracking — enhanced hygiene and operational efficiency in office-based setups

CERTIFICATES

- LangChain – Develop AI Agents with LangChain & LangGraph | Udemy - 2025
- Generative AI with Large Language Models | Coursera - 2024
- Building Video AI Applications at the Edge on Jetson Nano | NVIDIA Deep Learning Institute - 2024
- Getting Started with AI on Jetson Nano | NVIDIA Deep Learning Institute - 2024
- Artificial Intelligence and Machine Learning | Global InfoVentures Pvt. Ltd. - 2023
- Getting Started with Deep Learning | NVIDIA Deep Learning Institute – 2022