

Resume

Nafis Aslam

nafisaslam1819@gmail.com nafisaslam.com github.com/NafisAslam70 India · Open to Remote

→ [Live Demo: Real-Time Attention Monitoring \(YOLO + MediaPipe\)](#) → [Preprint: DOI 10.5281/zenodo.19266394](#)

PyTorch · YOLO · OpenCV · MediaPipe · CNNs · Transformers · GANs · Flask · Docker · Scikit-learn · Hugging Face

*A highly motivated and disciplined **Machine Learning/Computer Vision engineer** with a degree from **Universiti Sains Malaysia (USM)** majoring in AI, **16-month MIT graduate-level ML training**, and multiple shipped systems — including a **real-time attention monitoring system (YOLO + MediaPipe, <50ms CPU inference)** with a **first-author preprint on Zenodo**. I build systems that work outside the lab: fast, deployable, and grounded in real problems.*

Education

Massachusetts Institute of Technology

MITx / edX (online)

MicroMasters in Statistics and Data Science

Aug 2024 – Dec 2025

- **16-month, graduate-level program** (4 courses, 10–15 hrs/week) · self-funded alongside full degree
- Data Analysis (93%), Machine Learning (83%), Probability (78%), Statistics (60%), Capstone (73%)

Universiti Sains Malaysia

Penang, Malaysia

BSc (Hons.) Computer Science — Major: Intelligent Computing (AI)

2021 – 2025

- CGPA: 8.3/10 (3.32/4.00) Major GPA: 9.55/10 (3.82/4.00)
- Final 4 Semesters: 3.49 → 3.65 → 3.73 → 3.82
- Dean's List — 3 consecutive final semesters Gold Award — PIXEL 2025 (top 5%)

Technical Skills

Computer Vision

YOLO (v9/v11), OpenCV, MediaPipe, real-time detection, gaze/pose estimation, image classification, transfer learning (ResNet, EfficientNet), ByteTrack (familiar), temporal modelling

Deep Learning

PyTorch, CNNs, Transformers, GANs, Diffusion Models, MLP design, backpropagation, pruning & quantisation, on-device optimisation

Machine Learning

Scikit-learn, Pandas, NumPy, Ensemble Methods, Time-Series (ARIMA, GBM), Clustering, feature engineering, cross-validation

Algorithms & Optimisation

Model pruning, latency optimisation, heuristic design, hyperparameter tuning

MLOps & Deployment

Docker, Flask, Streamlit, ONNX (basic), on-device inference, cloud VMs (Azure/GCP), Git

Other Tools

Jupyter, Linux, Hugging Face, PostgreSQL, CUDA (basic), Weights & Biases (basic), Next.js

continued on next page →

Projects & Experience

- 2024–2025 **DeepWork AI / DeepLens Engine for Focus (DLEF)** *Gold Award, PIXEL 2025*
Stack: YOLOv11n-cls, MediaPipe (gaze, EAR), OpenCV, PyTorch, Next.js, Flask
- Real-time CV system classifying 6 attention/behaviour states (Focused, Phone, Absent, Drowsy, Looking Away, Bad Posture) from webcam feed.
 - Engineered a novel **20/30s temporal heuristic** — models attention as a time-series signal rather than per-frame output, suppressing false positives from momentary distraction and producing stable, actionable classifications.
 - **CPU-only on-device inference at <50ms latency**, zero video storage; shipped as full-stack privacy-preserving web app named DeepWorkAI, Pomodoro timer, real-time focus overlay, and session analytics.
 - Published first-author preprint (DOI: 10.5281/zenodo.19266394); won Gold Award (top 5%) at USM PIXEL 2025.
- 2026–present **Computer Vision & Deep Learning Fellow** — WorldQuant University (online)
- Wildlife camera-trap classification (Côte d’Ivoire) — multi-class CNN from scratch, PyTorch, end-to-end on cloud VM.
 - Crop disease classification (Uganda, 5 classes) — fine-tuned ResNet backbone with transfer learning; callbacks for training optimisation.
 - MLP regression & binary classification (Concrete Strength, Heart Disease) — manual forward/backprop then PyTorch; explored depth, non-linearity, optimiser behaviour.
- Mar 2026– **Machine Learning Intern** — Unified Mentor (Remote)
present
- Built time-series forecasting pipeline (ARIMA, Random Forest, Gradient Boosting) on 3+ years of weekly government care load data; 3 models benchmarked end-to-end.
 - Developed unsupervised learner segmentation (K-Means, Hierarchical clustering) with feature engineering from clickstream logs; co-authored internal research reports.
- Oct 2025–Feb **Software Engineer** — Meed Public School, India (Freelance)
2026
- Built full-stack platform (Next.js, Drizzle, NeonDB, Twilio) for 30+ users with real-time dashboards, role-based access, and LLM assistant with offline inference.
- Mar 2023–Aug **Software Engineering Intern** — USM Speech Processing Group
2023
- End-to-end stock analysis application (React, Node.js); real-time data processing, interactive visualisation, and data pipeline construction.

Achievements

- **Published:** First-author preprint — DeepLens Engine for Focus, Zenodo 2026. DOI: 10.5281/zenodo.19266394
- **Gold Honor (Top 2%)** — International Youth Math Challenge 2020
- **Competitive Programming:** Active on LeetCode (~50 problems solved); active on Kaggle
- **Leadership:** Superintendent, Meed Public School (2018–2022) — directed academics and operations for 300+ residential students; structural attention gaps observed here directly motivated the DLEF research
- **Languages:** English (fluent), Hindi (native), Urdu (native), Bahasa Malaysia (basic)