

A K M Asifuzzaman

AI/ML Researcher & Full-Stack AI Engineer

Chattogram, Bangladesh asif@asifzaman.dev asifzaman.dev github.com/A-K-M-Asifuzzaman
linkedin.com/in/a-k-m-asifuzzaman-6027442a7 huggingface.co/AsifZaman1912

RESEARCH INTERESTS

Vision-language models · Computer vision · Multimodal learning · Visual grounding · Trustworthy AI · Robustness and generalisation · Explainable AI

Seeking a research-focused, preferably funded MSc starting Fall 2027, as the first step towards doctoral research.

EDUCATION

BSc in Computer Science and Engineering

2023 - February 2027 (expected)

East Delta University, Chattogram, Bangladesh

CGPA 3.79 / 4.00

RESEARCH PROJECTS

InsightVLM — Naming the statistical finding a chart shows — and checking it against the plotted data

[Manuscript in preparation](#)

Team project — Group Hocruz, Phitron AI/ML Thesis Track. A K M Asifuzzaman is the first-listed author on the manuscript draft and maintains the repository and the public demo. The repository does not document a per-member split of the work, so none is claimed here.

- 86.7% real test charts ($n = 1,192$), 95% CI [84.6, 88.6]
- 46.8% best zero-shot VLM on the same charts (Qwen3-VL-8B)
- 71.6% charts re-drawn by a plotting library never seen in training
- Current vision-language models cannot do this out of the box: the best zero-shot model reaches 46.8%, and the smallest is barely above a rule that reads only the chart type.

FoodGenome AI — A 97% food classifier that also says, per image, when not to trust it

[Completed research project](#)

Sole author: data pipeline, backbone selection, probes, calibration, conformal prediction, the retrieval-grounded nutrition layer and the deployed application.

- 97.16% top-1, Food-101 test ($n = 25,250$)
- 0.051 $\rightarrow$ 0.006 expected calibration error after temperature scaling
- 99.56% conformal coverage at $\alpha = 0.01$, 1.54 candidates on average
- A parameter-free probability average beat a trained 3.97M-parameter fusion head (97.16% vs 96.97%).

RoboDesign Atlas — Domain-adapting CLIP to retrieve engineering figures, and auditing what the corpus contains

Completed research project

Sole author: dataset streaming and deduplication, fine-tuning, evaluation, significance testing, corpus audit and the in-browser demo. The underlying dataset is Le et al.'s.

- 83 → 26 median rank of the correct figure, text→image, 5,000 candidates
- 12.40 → 17.24 Recall@1, text→image (exact McNemar $p \approx 7 \times 10^{-31}$)
- 5.58% captions that match robotics terms and nothing else
- Unfreezing a quarter of CLIP's parameters for under an hour moves the correct figure's median rank from 83 to 26.

FairShift-PFN — Fairness side-effects of label-free test-time adaptation in tabular foundation models

Manuscript in preparation

Sole author: research question, theory, library, experiments, pre-registration and manuscript.

- 0 utility-harm runs of the gate on every real set (pre-registered)
- p 0.55–0.94 label-free estimate vs true equalized-odds change, held-out real sets
- $\sim 10^3$ labels needed to certify the ~ 1 pp gains on real shifts (budget tested: ≤ 400)
- Label-free prior-correcting adaptations widen equalized-odds gaps on census, hospital and lending shifts and for all three models; group-wise prior correction is the worst offender, and accuracy-only checks miss it.

CauseGround — Visual grounding and reliable multimodal reasoning

Ongoing research direction; no results yet.

MANUSCRIPTS

No published, accepted, submitted or preprint papers.

InsightVLM-20K: chart images paired with statistically verified findings. With Robiul, Fahim Ali, Iftehad, Md Arfat Yunus. [Manuscript in preparation](#) — Data-descriptor draft with open placeholders. Not submitted to any venue.

What Can Be Certified Without Labels? Fairness Side-Effects of Test-Time Adaptation in Tabular Foundation Models. Sole author. [Manuscript in preparation](#) — Complete draft after internal revision. Not submitted to any venue.

EXPERIENCE

Founder, TrinolIT Solutions

Sep 2026 – present

Dhaka, Bangladesh (remote)

Software and AI studio building bilingual (Bangla/English) business software. Delivered Nexora AI, a multi-tenant ERP/POS platform; Trinol OS and a Chromium-based browser are in development.

SELECTED ENGINEERING PROJECTS

- DocFlow AI — 100% citation verification, 30 facts checked

- NWP-Core — 60.7% top-5 accuracy, 19.5M params, 0 servers
- ChatgaiyyaLM — 20.5 → 65.6 chrF, Chittagonian → Bangla, n = 120
- DermAI — 36.4pp provenance shortcut removed
- Recurra — 0.892 test PR-AUC, features known at diagnosis (base rate 0.286)

TECHNICAL SKILLS (DEMONSTRATED IN PROJECTS)

RESEARCH METHODOLOGY

Calibration (temperature scaling, ECE) · Conformal prediction · Bootstrap CIs & paired significance tests · Leakage-safe splits & dataset auditing · Explainability (SHAP, counterfactuals)

AI / MACHINE LEARNING

Python · PyTorch · Hugging Face Transformers · PEFT / LoRA fine-tuning · scikit-learn · XGBoost / LightGBM · LSTM / RNN · ONNX export & in-browser inference

COMPUTER VISION & MULTIMODAL AI

CNNs (ResNet, DenseNet, EfficientNet) · Vision transformers (SigLIP, EVA-02) · CLIP / contrastive image-text models · Vision-language model evaluation · YOLO object detection · Grad-CAM attribution

LLMS, RAG & AI AGENTS

Retrieval-augmented generation · Hybrid dense + BM25 retrieval · Cross-encoder reranking · Multi-agent orchestration · OpenAI API · Anthropic API · Speech recognition (Whisper)

DATABASES & RETRIEVAL

PostgreSQL (incl. row-level security) · Qdrant · FAISS · Redis · Neo4j / graph retrieval · MongoDB · Pandas / NumPy

CERTIFICATIONS

Docker for Machine Learning — CampusX, Aug 2026 · FastAPI for Machine Learning — CampusX, Jun 2026 · Deep Learning: Beginner to Advanced — Codebasics, Jan 2026 · Explainable AI (XAI) — CampusX, 2025 · AI Engineer Core: LLM, RAG, QLoRA, Agents — Udemy, 2025 · NLP with Python — Udemy, 2025 · Deep Learning A-Z 2025 — Udemy, 2025 · A Deep Understanding of Deep Learning — Udemy, 2025 · Machine Learning A-Z (AI, Python & R) — Udemy, Apr 2025 · Supervised ML: Regression & Classification — DeepLearning.AI / Coursera, Feb 2025 · Python for Data Science, AI & Development — IBM / Coursera, Feb 2025 · Master ML for Data Science — Codebasics, Jan 2025 · Complete Web Development — Programming Hero, Jun 2024

REFERENCES

Available on request.