Eishaan Khatri

Research Associate Applicant | Computer Vision Robustness | Test-Time Adaptation | VLMs

+91-8171380034 | eishaankhatri@gmail.com | LinkedIn | GitHub | Portfolio

PROFILE

Final-year B.Tech Mathematics & Computing student at RGIPT, ranked **1st in the department**, with a peer-reviewed ESWA paper and focused computer-vision robustness work in **test-time adaptation** and **CLIP/VLM failure analysis**. Available for in-person RA work in Bangalore after May 2026. Target fit: data-efficient visual learning, domain shift, test-time/continual adaptation, and reliable VLM deployment.

EDUCATION

Rajiv Gandhi Institute of Petroleum Technology, Amethi — Expected May 2026

B.Tech, Mathematics & Computing; Minor, Business Consulting | CPI: 8.60; **Rank 1, Mathematics Department**

Academic signal: *JEE Advanced 2022 qualified, top 1.5 percentile among approx. 1.2M candidates.* **Preparation:** *Linear Algebra, Optimization, Numerical Methods, Statistical Inference; self-study implementations in corruption benchmarks, CLIP/VLMs, test-time adaptation, and calibration.*

SELECTED RESEARCH PROJECTS

Source-Preserved Test-Time Adaptation on CIFAR-10-C — GitHub

PyTorch, ResNet-18, CIFAR-10-C, BN Adapt, Tent, confidence/margin gating, source/adapted prediction mixing

- Implemented source-only inference, BN adaptation, Tent entropy minimization, source anchoring, and gated adaptation across 25 corruption/severity settings without target labels.
- Improved mean accuracy from 69.16% source-only and 82.18% Tent to 82.64% with 75/25 adapted-source mixing; reduced negative-transfer cases from 5 to 3 under mild brightness/fog shifts.

CLIP Zero-Shot Robustness Under Corruption and Domain Shift — GitHub

OpenCLIP ViT-B/32, CIFAR-10-C, CIFAR-100-C, EuroSAT, prompt ensembling, ECE, feature drift

- Evaluated zero-shot robustness with 8-template prompt ensembling, corruption severity analysis, mCE, confidence gap, entropy, and clean-corrupted feature similarity.
- Measured 86.30% clean CIFAR-10 accuracy dropping to 34.30% under severe gaussian noise; feature diagnostics tracked degradation on CIFAR-10-C (rho=0.613) and CIFAR-100-C (rho=0.720). EuroSAT zero-shot: 33.50%.

Edge-VCA: CPU Video Content Analysis for Edge Surveillance — GitHub

YOLOv8-Nano, OpenCV, IoU tracking, polygon-zone analytics, edge inference

- Built a CPU-only video analytics prototype with detection, custom tracking, zone entry/exit, dwell-time alerts, speed estimation, and motion heatmaps.
- Compared 3 edge-suitable detectors at 640x480; YOLOv8-Nano reached 24.5 FPS, 40.8 ms/frame, and 185 MB peak RAM in the project benchmark; documented occlusion, perspective, and low-light failure modes.

PUBLICATIONS AND RESEARCH OUTPUT

Target-Based Offensive Language Detection in Hinglish — ESWA, Vol. 308, May 2026

Low-resource learning, annotation design, multilingual transformer evaluation

- Built a 6,800+ instance Hinglish dataset with target-level supervision for noisy, code-mixed social text; proposed token-level target annotation and evaluated mBERT, XLM-R, and MuRIL.
- Handled dataset construction, annotation design, pretrained-model evaluation, and error analysis under limited-label/noisy-data conditions.

Bridging the Gap: Enhancing Zero-Cost LLMs for Scientific QA — CASML 2025 Poster

Python for Physics research program, scientific ML, retrieval grounding, physics/materials-science QA

- Co-authored CASML poster on retrieval grounding, scientific knowledge-graph context, and domain-specific validation for physics/materials-science QA; paper reports 68% hallucination reduction.

TECHNICAL EXPERIENCE

Indian Institute of Management Shillong — May 2025-Present

Research Analyst Intern | Mentor: Dr. Shrawan Kumar Trivedi | Remote

- Built Python/R modeling workflows with feature engineering, cross-validation, error analysis, SHAP, and LIME for explainable decision-modeling research.

Polka Labs / Polkassembly, Singapore — May-Nov 2025

AI Developer Intern | Remote

- Engineered reproducible AI/data pipelines and structured-output workflows, reducing project integration time from roughly one week to under five hours.

TECHNICAL SKILLS

Deep learning/CV: PyTorch, TensorFlow, OpenCLIP, HuggingFace, YOLOv8, OpenCV, ResNet, CLIP/VLM evaluation, object detection, tracking, edge inference.

Robust ML/research: test-time adaptation, BN adaptation, Tent, corruption benchmarks, ECE/confidence gap, feature-drift analysis, uncertainty estimation, ablations, error analysis.

Programming/tools: Python, R, SQL, C++, Git, Pandas, NumPy, Matplotlib/Seaborn, FastAPI, Streamlit, pytest.

Additional signal: NCA-Language-Models, a 2,100+ line interactive paper reimplementation studying Neural Cellular Automata signals for synthetic pretraining.