

Md Malek Sarker

maleksarkerlimon7@gmail.com • limon-7.github.io • github.com/Limon-7
linkedin.com/in/limon7 • ORCID: 0009-0003-2751-4802 • Dhaka, Bangladesh

RESEARCH INTERESTS

Reinforcement learning for language models • Recommender systems • LLM agents and RAG • ML evaluation

EDUCATION

International University of Business Agriculture and Technology (IUBAT) Dhaka, Bangladesh
B.Sc. in Computer Science and Engineering • CGPA 3.20/4.00 2018

- Final-year project: *Development of an Enterprise Resource Planning (ERP) System*. Supervisor: Md. Rashedul Islam.

RESEARCH PROJECTS

GRPO from Scratch: What a Lenient Grader Teaches a Language Model 2026

- Implemented Group Relative Policy Optimization from scratch in PyTorch (group-normalised advantages, clipped surrogate, KL estimate, resumable training) with unit tests, and used it to fine-tune Qwen2.5-0.5B-Instruct with LoRA on GSM8K.
- Designed a controlled study of reward misspecification: four reward graders (strict; last-number; any-match; strict plus a format bonus) with model, data, hyperparameters and seed held fixed, and every model scored by the strict grader.
- Strict reward raised test accuracy from 9.1% to 52.3%. Lenient rewards improved the arithmetic (last-number accuracy 39.3% to 50.8%) but drove strict accuracy to about 0% as the model abandoned the required answer format. A 0.2 format bonus collapsed outputs to a 9-token guessed answer within ten training steps.

Multimodal Semantic IDs for Generative Recommendation 2026

- Built a TIGER-style generative recommender (RQ-VAE semantic IDs, T5 encoder–decoder trained from scratch, trie-constrained beam search) and compared text-only, image-only, early-fusion and late-fusion IDs on Amazon Reviews 2023 (two categories, up to 147K items, ranking against the full catalogue).
- Late fusion of Sentence-T5 and SigLIP embeddings improved NDCG@10 over text-only IDs by 13–44% (relative) in all four dataset and split settings, with 95% paired-bootstrap intervals excluding zero; early fusion gave no consistent gain.
- Reported the negative results: no generative variant retrieved a cold-start item under a temporal split, and SASRec fed the same frozen features outperformed every generative model (NDCG@10 0.0473 against 0.0332). Checked the implementation against published TIGER and SASRec numbers on Amazon Beauty.

Brain Tumor Detection in MRI: One-Stage, Two-Stage and Transformer Detectors 2026

- Benchmarked YOLOv8m, Faster R-CNN (ResNet-101-FPN) and DETR (ResNet-50) for four-class tumor localisation on 5,249 annotated MRI slices under a shared YOLO-to-COCO data pipeline and COCO evaluation.
- YOLOv8m reached 0.781 mAP@[.5:.95] against 0.562 and 0.508. Analysis by IoU threshold showed Faster R-CNN detects tumors but localises loosely (AP50 0.927, AP75 0.592), and both weaker models score about 0.26 mAP on small lesions.

Smart Research Assistant: Local RAG over Scientific Papers 2026

- Built a fully local RAG pipeline (MiniLM, ChromaDB, Llama 3.2) with source attribution and a reference-aware chunking stage that keeps bibliography fragments out of the index.

MANUSCRIPTS

M. K. S. Likhon and **M. M. Sarker**. “Blur Sensitivity and Generator Dependence of Fingerprint Alteration Detectors Trained on SOCOFing.” Manuscript in preparation, 2026.

Ten detectors reach 0.996–0.998 macro-F1 on SOCOFing, yet a one-pixel Gaussian blur makes eight of nine image models flag 93–97% of genuine prints as altered, and alterations from an independent generator are mostly missed (macro-F1 0.38–0.42).

M. K. S. Likhon and **M. M. Sarker**. “Data Audit and Revised Baselines for Structural Damage-Type Classification on PEER Hub ImageNet.” Manuscript in preparation, 2026.

Audit finds 91 near-duplicate training pairs (30 with conflicting labels) and 23 test images with a near-copy in training; a duplicate-aware protocol yields 77.4% test accuracy against 72.4% for the best published result.

PROFESSIONAL EXPERIENCE

Brac IT Services Ltd.

Senior Software Engineer (R&D)

Dhaka, Bangladesh

May 2025 – Present

- Design and build LLM-based code-generation systems with retrieval-augmented generation, grounding outputs in company-specific architectural patterns and coding standards.
- Build agentic LLM workflows and architect distributed services for enterprise and education platforms.

CholPori Technologies Ltd.

Software Engineer

Dhaka, Bangladesh

Sep 2021 – Apr 2025

- Engineered data pipelines for a personalised educational recommender that models learner progress, lesson similarity and prior learning outcomes for context-aware content delivery.
- Developed core platform modules (.NET Core, Angular, PostgreSQL, RabbitMQ); reduced latency with caching and resolved concurrency issues in inter-service messaging.

Infonet Associates Ltd.

Software Engineer

Dhaka, Bangladesh

Apr 2021 – Sep 2021

- Lead contributor to a system for the Bangladesh Power Development Board; optimised SQL queries over large datasets for analytics.

TECHNICAL SKILLS

Machine learning: PyTorch, Hugging Face Transformers, PEFT (LoRA), LangGraph, sentence-transformers, Detectron2, Ultralytics YOLOv8, Weights & Biases

Methods: Policy-gradient fine-tuning (GRPO), generative retrieval (RQ-VAE, semantic IDs), sequential recommendation, retrieval-augmented generation, object detection, bootstrap significance testing

Programming: Python, C#, C/C++, JavaScript/TypeScript, SQL

Systems: .NET Core, PostgreSQL, MongoDB, ChromaDB, RabbitMQ, Docker, Git, Linux

HONORS AND SERVICE

- Best Employee of the Month, Infonet Associates Ltd., July 2021.
- Volunteer, IUBAT Programming Contest, 2016.