NAYEEM HOSSEN JIM

Dhaka, Bangladesh | nayeemhossenjim@gmail.com | +880-1631-123854
Github | LinkedIn | Portfolio

PROFESSIONAL SUMMARY

AI Engineer and Game Developer with 1+ years of professional experience shipping production-grade Retrieval-Augmented Generation (RAG) systems, adapting large language models (LLMs), and engineering real-time gameplay mechanics in Unreal Engine 5. Proficient in Python, C++, LangChain, LangGraph, and Hugging Face Transformers, with a proven record of applying MLOps and LLMOps practices across the full model lifecycle. Combines deep machine learning expertise with hands-on game programming to produce scalable, measurable outcomes.

TECHNICAL SKILLS

Programming Languages: Python, C++, C, C#

AI and Machine Learning: LangChain, LangGraph, Hugging Face Transformers, TensorFlow, PyTorch, RAG Systems, LLM Adaptation (LoRA, QLoRA), Unsloth, YOLO (Ultralytics)

MLOps and LLMOps: MLflow, DVC, Docker, CI/CD Pipelines, Model Versioning, Experiment Tracking, Production Monitoring, n8n Automation

Game Development: Unreal Engine 5, C++ Gameplay Programming, Blueprints, Unity, Blender

Databases and APIs: PostgreSQL, pgvector, MySQL, FastAPI, REST APIs

Developer Tools: Git, GitHub, VS Code, Google Colab, Kaggle

WORK EXPERIENCE

Junior AI Engineer | Softvence Delta, Dhaka, Bangladesh *October 2025 – Present*

- Architected and shipped production-grade RAG pipelines using LangChain and LangGraph, processing 100,000+ monthly queries and boosting automated issue resolution by 30%.
- Optimized Hugging Face Transformer models through LoRA and QLoRA parameter-efficient adaptation for domain-specific NLP tasks, reducing inference latency by 40% and improving benchmark accuracy.
- Established MLOps standards covering experiment tracking (MLflow), model versioning (DVC), CI/CD automation, and live production monitoring, cutting deployment errors by 25%.
- Automated cross-functional data workflows using n8n by integrating 10+ third-party APIs, eliminating all manual data-handling steps for the engineering team.
- Trained a real-time Bangla Sign Language recognition model on a custom YOLO (Ultralytics) dataset, achieving 87% accuracy across 36 alphabet classes in live-camera evaluation.

Unreal Engine Gameplay Programmer | Studio Aloukik, Dhaka, Bangladesh *November 2025 – Present*

- Engineered core gameplay systems — including melee mechanics, animation state machines, hit detection, and cinematic camera rigs — for 3 simultaneous client projects using C++ and Blueprints in Unreal Engine 5.
- Delivered 4 production-ready client prototypes on schedule, each featuring responsive action-RPG combat, which cut client revision cycles by 35%.

- Authored 5+ game features for a Bangladesh-themed horror title and the Masud Rana narrative project, both in collaboration with Sheba Prokashoni, shipping all work to internal review on time.

PROJECTS

LLM From Scratch | *PyTorch, Python*

- Constructed a GPT-3-architecture language model in PyTorch from the ground up, then fine-tuned it for personal-assistant use cases, demonstrating full-stack mastery of transformer internals and backpropagation.
- Implemented a custom SimpleTokenizerV2 with regex-based text cleaning, vocabulary mapping, and unknown-token handling to support efficient preprocessing for downstream tasks.

Parameter-Efficient LLM Adaptation | *Unsloth, Hugging Face, LoRA, QLoRA*

- Applied LoRA and QLoRA within the Unsloth framework to adapt large language models with 60% lower GPU memory consumption and up to 4x faster convergence versus full-weight fine-tuning.
- Streamlined the model adaptation workflow — encompassing dataset curation, tokenization, and checkpoint management — reducing new job setup time from 4 hours to under 20 minutes.

Real-Time Skin Cancer Detection System | *YOLOv12, Python, JavaScript, FastAPI*

- Developed a full-stack computer vision system using YOLOv1 for skin lesion classification, attaining 91% accuracy across 5 cancer categories on a held-out test set.
- Deployed the solution on a VPS with a FastAPI backend and a custom JavaScript frontend, enabling sub-200 ms clinical inference for real-time medical evaluation.

The Last Gods – Third-Person Action Game | *Unreal Engine 5, C++, Blueprints*

- Built a third-person action game featuring responsive character controls, dynamic environmental interactions, and immersive combat feedback systems in Unreal Engine 5.
- Programmed state-machine-driven enemy boss AI with 8 distinct attack patterns, reducing observed player loop repetition by 50% across playtesting sessions.

EDUCATION

B.Sc. in Computer Science and Engineering | Uttara University, Dhaka, Bangladesh

CGPA: Running

CERTIFICATIONS AND AWARDS

- AWS AI Practitioner Challenge - Udacity
- Advanced Learning Algorithms - Stanford University (Coursera)
- Supervised Machine Learning: Regression and Classification - Stanford University (Coursera)
- 12th Place - Game Jam Competition, BUET CSE Fest 2024
- Vice President - Machine Learning Club, Uttara University (March 2025 - January 2026)