

Nishat Ahmad

Machine Learning Developer | Lahore, Punjab, Pakistan
nishatahmad.dev | nishat.ahmed1200@gmail.com | +92 311 6300916
github.com/Nishat-Ahmad | linkedin.com/in/nishat-ahmad

SUMMARY

Software Engineer specializing in scalable machine learning infrastructure, GPU optimization, and resilient backend systems. Proven ability to architect end-to-end MLOps pipelines and accelerate inference engines, bridging the gap between heavy-metal ML and production-ready deployments.

EDUCATION

Ghulam Ishaq Khan Institute of Science & Technology (GIKI) Topi, Pakistan
Bachelor of Science in Artificial Intelligence Sep 2023 – May 2027 (Expected)

CGPA: 3.44 | **Honors:** Dean's Honor List (x2)

Relevant Coursework: Deep Neural Networks, GPU Programming, Design & Analysis of Algorithms, Machine Learning, Operating Systems, Data Structures & Algorithms, Object oriented programming, Database management system

EXPERIENCE

GIKI RoboCup Team | *Software Developer* Swabi, Pakistan | Dec 2025 – Present

- Programmed a video game utilizing OpenGL (GLFW) to model different robot prototypes for simulated testing.
- Configured build automation pipelines using CMake, streamlining the compilation process and reducing build times.
- Deployed and tested robotic software environments natively on Linux (Pop!_OS).

FFC | *Summer Intern* Mirpur Mathelo, Pakistan | July 2022 – Aug 2022

- Analyzed the enterprise IT infrastructure, workflows, and service delivery models to identify efficiency bottlenecks.
- Shadowed senior IT personnel to understand inter-departmental technical support and network operations.

KEY PROJECTS

GPU-Accelerated Sentiment Analysis Engine *CUDA, C++, Python, PyBind11, PyTorch*

- Engineered a custom NLP inference pipeline via C++ and PyBind11 to process 7M reviews, achieving a 1.59x speedup and 13.5x memory reduction over PyTorch baselines.
- Hand-coded 12 CUDA kernels, including a 4x4 Register-Tiled GEMM that outperformed NVIDIA cuBLAS by up to 1.57x on production matrix projections (2048x128).
- Maximized SM occupancy to 98% by utilizing Kernel Fusion, float4 memory vectorization, and warp-shuffle reduction primitives.

Fleet-Vision: Autonomous Vehicle Inspection *PyTorch, FastAPI, LangChain, Docker*

- Designed a dual-model vision pipeline utilizing ResNet-50 (85.3% test accuracy) for vehicle classification and a Swin Transformer for recall-optimized structural damage detection.
- Architected an asynchronous FastAPI and LangChain backend, accelerating multi-angle image inference via parallel `asyncio` execution and integrating OpenRouteService for dynamic ride routing.
- Implemented a custom Grad-CAM solution for real-time damage localization heatmaps and deployed the Dockerized microservice via an automated GitHub Actions CI/CD pipeline.

Thalassa: End-to-End MLOps Platform *Python, Prefect, FastAPI, Docker, XGBoost*

- Constructed a production-grade MLOps platform automating data ingestion and feature engineering, orchestrating a multi-task pipeline of XGBoost, ARIMA, and KMeans models.
- Built an iterative FastAPI simulation engine capable of forecasting returns up to 90 days out by dynamically recalculating lagged features and technical volatilities.
- Architected a gated GitHub Actions CI/CD pipeline and Dockerized deployment system, implementing custom Kolmogorov-Smirnov data drift monitoring to audit feature integrity and prevent model degradation.

HONORS, AWARDS & CERTIFICATIONS

- Academic Honors & Societies:** Dean's Honor List (x2); Member of GIKI Competitive Programming Club and GIKI RoboCup Team.
- Competitions & Hackathons:** ICPC OnSite Participant (2024) & Preliminary (2023); Competed in MIT Hackathon and 2 university-level hackathons.
- Certifications:** Introduction to AI with Python (Harvard), Deep Learning Specialization (DeepLearning.ai), CNNs in TensorFlow.

TECHNICAL SKILLS

- Languages:** Python, C++, SQL, C, CUDA.
- Frameworks & Libraries:** PyTorch, Scikit-learn, FastAPI.
- Tools & Platforms:** Docker, Git, GitHub, MLOps, Linux.
- Domains:** Machine Learning, Deep Learning, Computer Vision (CNNs), Natural Language Processing (NLP), Backend Infrastructure.