

Inayat Arshad

NLP Researcher & Applied AI Engineer

Islamabad, Pakistan | +92-333-5744954 | enayyat156@gmail.com

[LinkedIn](#) | [GitHub](#) | [Portfolio](#) | [ORCID](#) | [ResearchGate](#)

SUMMARY

Computer Science researcher and Applied AI Engineer working across machine learning, artificial intelligence, multimodal systems, and language technologies. Research experience spans low-resource learning, explainable AI, human-model agreement, multimodal evaluation, and reliable AI systems.

EDUCATION

Pakistan Institute of Engineering & Applied Sciences (PIEAS) <i>B.S. in Computer and Information Sciences, CGPA: 3.72/4.00 (92.5%, Top 2%)</i>	Islamabad, Pakistan Sep 2022 – Jul 2026
• Relevant coursework: Fundamentals of Machine Learning, Information Storage and Retrieval, Signal and Image Processing, Applied Machine Learning for Bioinformatics, Probability and Statistics, Linear Algebra, Parallel and Distributed Processing.	
Punjab Group of Colleges F-8/2 <i>Intermediate in Computer Science, 1040/1100 (94.5%, Ranked 7th nationwide)</i>	Islamabad, Pakistan Sep 2020 – Jul 2022
Siddeeq Public School <i>Matriculation (Science), 1073/1100 (97.5%)</i>	Islamabad, Pakistan Mar 2018 – Mar 2020

PUBLICATIONS AND RESEARCH

TRUST-X: Real-Time Explainable Decision-Making for RL-Based UAV Control | *IEEE, Published*

- Designed an explainable decision-making framework coupling surrogate decision trees (91% fidelity), SHAP-based attribution, and a Bayesian trust feedback loop for deep reinforcement learning; achieved 94% mission success versus a 79% TD3 baseline with 12 ms explanation latency.

MUTEX: Multilingual Transformers and CRF for Urdu Toxic Span Detection | *Springer, Under Review*

- Curated URTOX, the first token-level (BIO) toxic span corpus for Urdu — 14,338 manually annotated samples across social media, news, and YouTube comments ($\kappa = 0.82$, $\alpha = 0.81$).
- Established the first supervised baseline for Urdu toxic span detection with an XLM-RoBERTa + CRF sequence labeler at 60.0 token-level F1, benchmarked against mBERT (59.2) and BiLSTM-CRF (56.0) under code-switching and rich morphological variation.

Advancing Urdu Toxicity Detection: Span Models and a Multimodal Audio Extension | *Elsevier, Under Review*

- Built URTOX-MM, the first paired text-audio Urdu toxicity dataset (14,338 TTS-synthesized clips), and URTOX-HumanAudio, a 2,000-clip real-speech evaluation set stratified across four Pakistani regional accents.
- Raised token-level F1 from 60.0 to 67.0 through training-procedure refinements alone — no architecture change — with subword label propagation contributing 4 of the 7 points and lifting recall from 58.0 to 71.5.
- Achieved 83.2% weighted F1 on synthesized and 77.1% on real speech via late fusion of MMS-300M (Urdu adapter) with XLM-RoBERTa; language-matched pretraining degraded only 6.5 pp from synthetic to real speech versus 9.0 pp for English-pretrained wav2vec 2.0.

RESEARCH EXPERIENCE

AI/ML Researcher <i>PIEAS University, Supervisor: Dr. Ijaz Hussain</i>	Sep 2025 – Present Islamabad, Pakistan
• Lead research on low-resource NLP, multilingual representation learning, toxic-span detection, multimodal speech-text modeling, and language-model evaluation. • Curated and annotated the URTOX corpus with token-level toxicity spans, establishing an annotation protocol and achieving $\kappa = 0.82$ inter-annotator agreement. • Investigated transformer and sequence-labeling architectures for Urdu, including XLM-RoBERTa, mBERT, BiLSTM-CRF, and CRF-based decoding. • Extended the Urdu toxicity benchmark into speech by pairing span-level annotations with spoken Urdu and evaluating robustness across regional/accent variation.	

- Investigated zero- and few-shot LLM baselines against supervised token-level models to examine performance gaps between language models and human annotations.

AI Research Intern (Model Optimization)

National Aerospace Science & Technology Park (NASTP)

Jun 2025 – Sep 2025

Islamabad, Pakistan

- Investigated Fourier-based hybrid Transformer architectures for efficient sequence modeling and summarization, benchmarking 1D/2D FFT blocks and analyzing internal representations using CKA and SVD.
- Built and evaluated predictive-maintenance models using LSTM and Transformer architectures across aircraft engine health parameters.

Teaching Assistant – Computer and Information Sciences

Pakistan Institute of Engineering & Applied Sciences

Dec 2025 – May 2026

Islamabad, Pakistan

- Led laboratory sessions for 50+ undergraduate students covering Python, NLP pipelines, machine learning models, search algorithms, and sequence-processing methods.

SELECTED NLP & MULTIMODAL PROJECTS

URTOX: Urdu Toxic Span Detection | Python, PyTorch, XLM-R, CRF

- Built an end-to-end low-resource NLP benchmark with token-level BIO annotations, transformer fine-tuning, CRF decoding, SHAP attribution, and attention visualization for interpretable toxicity spans.

When Fusion Breaks: Multimodal Robustness Benchmark | PyTorch, Multimodal Learning

- Developed a reproducible framework comparing Early/Late Fusion, TFN, LMF, and MulT across text, audio, and visual modalities under graded corruption.
- Measured robustness to ASR errors, audio noise, visual occlusion, modality dropout, and temporal misalignment using retention curves, AUROC, and modality-reliance metrics.

LACUNA: NLP Research Gap Discovery Engine | NLP, Bibliometrics, Information Retrieval

- Built a bibliometric engine indexing 16,605 peer-reviewed publications across 73 languages, 26 NLP tasks, and 1,722 venues to quantify disparities in language-resource coverage.
- Selected as a Top 25 finalist among 250+ international teams in the NeuroLogic '26 Global NLP Datathon.

Multi-Agent RAG System with Self-Correcting LLM Pipelines | LangGraph, LLMs, Retrieval

- Built a nine-agent RAG architecture for query routing, retrieval, verification, and multi-step reasoning with hybrid retrieval and self-correction loops.

PROFESSIONAL EXPERIENCE

Head of Applied AI

Tech GIS

May 2025 – Present

Islamabad, Pakistan

- Developing **VYBERO**, a conversational AI voice agent for enterprise clients; built real-time ASR, intent classification, and entity-extraction pipelines integrated with FastAPI backend services and a Next.js frontend.
- Engineered full-stack AI architectures connecting PostgreSQL, vector databases, REST/WebSocket APIs, and human-in-the-loop review interfaces for reliable data processing.

Machine Learning Intern (Generative AI)

NAECO Blue GmbH

Jun 2023 – Sep 2023

Remote

- Built LangChain-based agentic pipelines processing weather, energy-production, and consumption data, reducing manual extraction time by approximately 60%.

HACKATHONS & COMPETITIONS

NeuroLogic '26 Global NLP Datathon: Top 25 finalist out of 250+ international teams; used LACUNA to identify underexplored language and task gaps.

National AI Hackathon '26: Presented MUTEX, a multimodal text–audio toxic span pipeline for low-resource Urdu content.

FortyGuard Hackathon '26: Represented NAECO Blue GmbH; built LangChain agent workflows for automated energy-load forecasting.

LEADERSHIP, MENTORSHIP & COMMUNITY

Technical Module Head

Google Developer Student Clubs (GDSC PIEAS)

Oct 2023 – Oct 2024

Islamabad, Pakistan

- Led technical curriculum planning and delivered workshops on applied AI/ML, NLP pipelines, and data processing to 150+ undergraduate students.

Executive Board Member

IEEE Computer Society PIEAS Chapter

Sep 2023 – Sep 2025

Islamabad, Pakistan

- Organized technical symposiums, computational workshops, and coding hackathons connecting students with emerging AI and machine-learning research.

Invited Technical Speaker – Adversarial Machine Learning

ICT Club, MNS UET

Jun 2025 – Jul 2025

Multan, Pakistan

- Delivered technical lectures on adversarial machine learning, model interpretability, and robust AI deployment to 100+ engineering students.

Area Fellowship Coordinator

Rotary Club of Islamabad (District 3272, Pakistan)

2024 – 2025

Islamabad, Pakistan

- Coordinated regional civic initiatives; recognized for community health and polio eradication initiatives.

HONOURS AND AWARDS

KAUST Workshop: Selected to attend the KAUST Workshop on Distributed Training in the Era of Large Models, Saudi Arabia (November 2025), among an international cohort of students and researchers.

NeuroLogic '26 Global NLP Datathon: Top 25 finalist out of 250+ international teams for LACUNA, an NLP research-gap discovery system.

Academic Standing: Graduated in the **Top 2%** of the PIEAS Computer and Information Sciences cohort with a CGPA of **3.72/4.00**.

PIEAS Entrance Examination: Ranked 12th nationwide across all departments in the undergraduate entrance examination.

National Board Recognition: Ranked **7th nationwide** in BISE Rawalpindi examinations and received national recognition for academic performance.

TECHNICAL SKILLS

NLP & Language Modeling: Token Classification, Sequence Labeling, CRF, XLM-RoBERTa, mBERT, Multilingual NLP, Low-Resource NLP, Semantic Retrieval, LLM Evaluation

Speech & Multimodal AI: ASR, MMS-300M, Text-Audio Modeling, Multimodal Fusion, Early/Late Fusion, TFN, LMF, MulT, Modality Robustness

Explainable & Trustworthy AI: SHAP, LIME, Surrogate Models, Attention Visualization, Uncertainty, Human-in-the-Loop Evaluation

Machine Learning: PyTorch, TensorFlow, scikit-learn, Transformers, NumPy, Pandas

LLMs & Generative AI: Hugging Face, LangChain, LangGraph, RAG, LoRA Fine-Tuning, Multi-Agent Systems

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

Engineering: FastAPI, Flask, PostgreSQL, ChromaDB, FAISS, Docker, Git, WebSockets, W&B

LANGUAGES

Urdu: Native **English:** Proficient (**IELTS Academic 8.0**)

CERTIFICATIONS

- **Natural Language Processing Specialization**, DeepLearning.AI
- **Machine Learning Specialization**, Stanford University / DeepLearning.AI
- **Generative AI with Large Language Models**, DeepLearning.AI
- **Agentic AI**, DeepLearning.AI