

Soumedhik Bharati

📍 West Bengal, India ✉ soumedhikbharati@gmail.com ☎ +91 8240947878 🌐 soumedhik.github.io
in soumedhik-bharati 📄 soumedhik

Education

Sister Nivedita University, Kolkata, India

Sept. 2022 – Sept. 2026

B.Tech in Computer Science Engineering

(Expected)

- **CGPA:** 8.68/10.0
- **Core CS:** Data Structures & Algorithms, Operating Systems, Computer Networks, DBMS, OOP, Cryptography & Network Security
- **AI/ML:** Machine Learning, Artificial Intelligence, Deep Learning, Optimization Techniques, Digital Image Processing, Computer Vision, Reinforcement Learning
- **Mathematics:** Linear Algebra, Probability & Statistics, Signals & Systems, Numerical Methods
- **Leadership:** Machine Learning Lead, SKEPSIS; Core Technical Team ML Lead, Google Developer Group (GDG) SNU

Publications

Accepted / Published

- **FROST: Factual Reasoning via Optimized Stochastic Trajectories in Large Language Models during Inference.** Soumedhik Bharati et al. **ACL Industry Track 2026** (Accepted, First Author). Achieves 28% cost reduction over Self-Consistency ensembles in production retrieval-reasoning systems.
- **AjamiMorph: Zero-Annotation Morphological Discovery for Hausa Ajami via Multi-Method Consensus** Soumedhik Bharati et al. **EACL 2026 – AbjadNLP Workshop** (Oral, First Author). 99.9% morphological coverage via fully unsupervised pipeline; enables zero-annotation discovery for a low-resource script.
- **Morphological Feature Extraction for Fine-Grained Sorani Kurdish Dialect Identification: A Hybrid Transformer-Linguistic Approach** Soumedhik Bharati et al. **EACL 2026 – AbjadNLP Workshop** (Oral, First Author). 91.9% accuracy; first dedicated dialect classification system for Sorani Kurdish.
- **LLM Orthographic Fragility Under Script Variation.** Soumedhik Bharati et al. **EACL 2026 – AbjadNLP Workshop** (Poster, First Author). Demonstrated 41% performance degradation in LLMs under orthographic perturbations.
- **HCAT-Net: Hierarchical Cross-Attention Transformer for EEG Emotion Classification.** Soumedhik Bharati et al. **IEEE CIACON 2025, First Author** (Best Paper Award, First Author). 99.8% accuracy, 100% ROC-AUC, 99.7% MCC on EEG Brainwave Dataset using ResNet-Transformer with Rotary Positional Embeddings.
- **The Glass Policy: A Quantum-Inspired RL Framework for Brittle Mastery in Upskilling** Soumedhik Bharati et al. **Springer QuantumKol 2025** (Best Paper Award, First Author). A quantum inspired RL framework for addressing the cold start problem and framing personalized employee training as a constrained RL problem.

Preprints

- **Dynamic Ranked List Truncation for Reranking Pipelines via LLM-generated Reference-Documents** Nilanjan et al. [arXiv:2604.09492](https://arxiv.org/abs/2604.09492). 33–66% inference time reduction over top-down and sliding window approaches; 3x throughput improvement on complex query loads.
- **DHV-MNet: An Efficient Transformer for Single-Image Dehazing in Surveillance Applications** Soumedhik Bharati et al. [SSRN Preprint](#). 55% parameter reduction; SOTA performance on RESIDE-6K benchmark. *Under Review.*

Under Review / In Preparation

- **CADET: BiLSTM with Multi-Head Attention for Automated Essay Scoring.** Soumedhik Bharati et al. SOTA QWK of 0.98 on ASAP and AES 2.0 datasets. *Journal under review.*
- **Novel DNA Encoding for Hybrid CNN-BiLSTM Sequence Classification.** Soumedhik Bharati et al. 97.2% accuracy; 35% parameter reduction over one-hot encoding baselines. *Journal under review.*
- **Deep Learning Surrogate for Quantum State Evolution Prediction.** Soumedhik Bharati et al. (with University of Lille). R^2 of 0.94 against ground-truth numerical solver. *Manuscript under review.*
- **Few-Shot Transfer Learning for Multimodal EEG/fMRI Signals.** Soumedhik Bharati et al. (with Xu Lab, Carnegie Mellon University). 12% zero-shot seizure prediction improvement; 95% reduction in labeling requirements. *Manuscript in preparation.*

- **Culturally Sensitive Evaluation Framework for LLM Socio-Cultural Alignment.** Soumedhik Bharati et al. (with IIT Hyderabad). *In preparation.*
- **RL for cultural adaptation and open-ended generation** Maharaj et al. (with IIT Hyderabad). *In preparation.*
- **Prestige Bias through low resource cultures and languages in LLMs** Soumedhik Bharati et al. (with Dr. Arij Riabi). *In preparation.*

Research Experience

Indian Institute of Technology Hyderabad, Sangareddy, India

Jan 2026 – Present

Research Intern

- **Task:** Develop a culturally sensitive evaluation framework to benchmark socio-cultural and regional linguistic alignment in large language models.
- **Action:** Designing annotation schemas and evaluation metrics spanning regional dialects, cultural context, and low-resource language phenomena across Indian languages.
- **Status:** Ongoing; manuscript in preparation.

Xu Lab, Carnegie Mellon University, Pittsburgh, PA (Remote)

Sept 2025 – Present

Undergraduate Researcher

- **Situation:** Neurological signal analysis requires massive labeled datasets, limiting clinical deployment.
- **Task:** Develop few-shot transfer learning for multimodal EEG/fMRI signals to reduce labeling requirements.
- **Action:** Architected transfer learning pipeline for a 1.2B parameter spatio-temporal Transformer pre-trained on 100,000+ hour clinical EEG corpus. Designed self-supervised contrastive learning objective to enhance latent space robustness.
- **Result:** Achieved **12%** improvement in zero-shot seizure prediction with **95%** reduction in labeled data requirements; produced **3x** more robust latent space with **80%** faster fine-tuning convergence across 5 neurological datasets. Manuscript in preparation. Targeting improvement over SOTA baselines (e.g., BrainOmni).

Indian Institute of Technology Kharagpur, Kharagpur, India 

May 2025 – July 2025

Research Intern

- **Situation:** Large-scale IR experiments with LLMs face computational bottlenecks limiting reproducibility.
- **Task:** Design a scalable reranking pipeline for accelerated large-scale experimentation.
- **Action:** Designed PSI-Rank, a multi-stage caching IR reranking pipeline with novel parallelism strategy. Executed 2,400+ large-scale LLM experiments on the Param Vidya HPC cluster using SLURM-based distributed computing.
- **Result:** Achieved **33%** and **66%** inference time reduction over top-down and sliding window approaches respectively; boosted throughput by up to **3x** for complex query loads while maintaining/improving NDCG@k and MAP metrics. Preprint: [arXiv:2604.09492](https://arxiv.org/abs/2604.09492) .

University of Lille, France (Remote)

Feb 2025 – Present

Collaborative Researcher

- **Situation:** Direct numerical integration of the Schrödinger equation is computationally prohibitive for quantum control optimization.
- **Task:** Create a deep learning surrogate for rapid quantum state evolution prediction.
- **Action:** Designed feedforward architecture with batch normalization and optimized dropout to map Hamiltonian control parameters and initial wavefunction to final evolved quantum state. Demonstrated superiority of Huber loss for heavy-tailed error distributions.
- **Result:** Achieved **R² of 0.94** against ground-truth numerical solver, enabling rapid quantum control parameter exploration. Collaborative work with postdoctoral researchers; manuscript under review.

Sister Nivedita University, Kolkata, India

Mar 2024 – Present

Undergraduate Researcher

- **HCAT-Net – EEG Emotion Classification:** Pioneered ResNet-Transformer architecture integrating 1D ResNet blocks with hierarchical Transformer encoder using Rotary Positional Embeddings, cross-attention fusion, and custom Balanced Ordinal Loss. Achieved **99.8% accuracy**, **100% ROC-AUC**, **99.7% MCC**. **Best Paper, IEEE CIACON 2025** .
- **LLM Robustness & Low-Resource NLP – EACL 2026 (2 Orals, 1 Poster):** (i) Demonstrated LLM orthographic fragility: 41% performance degradation under script variation. (ii) **Enabled zero-annotation morphology discovery for Hausa Ajami**  (99.9% coverage). (iii) **Built first 6-way Sorani Kurdish dialect classifier**  (91.9% accuracy).
- **CADET – Automated Essay Scoring:** Engineered BiLSTM with multi-head attention capturing hierarchical essay structure. Achieved SOTA QWK of **0.98** on ASAP and AES 2.0 datasets. Journal under review.
- **RL for Employee Training Optimization:** Implemented Actor-Critic agent with hybrid reward function for

multi-objective curriculum optimization under hard budget constraints. Achieved **82%** success rate over four alternative strategies. Presented at QuantumKol 2025.

- **DNA Sequence Classification:** Engineered novel encoding technique for hybrid CNN-BiLSTM model. Achieved **97.2%** accuracy with **35%** parameter reduction over one-hot encoding baselines. Journal under review.
- **Image Dehazing – Multi-Scale U-Net with Mamba SSM:** Optimized multi-scale U-Net integrated with Mamba state-space model for image-to-image dehazing. Achieved **55%** parameter reduction while attaining SOTA on RESIDE-6K benchmark. [SSRN Preprint](#) ; under review.

Work Experience

Exalt Data (*UChicago Polsky Center Build & Launch 2025*), Remote
Chief Technology Officer (CTO)

Jun. 2025 – Present

- **Situation:** Hyper-focused news platform required scalable ML infrastructure for voter behavior analysis and intelligent content retrieval.
- **Task:** Architect end-to-end ML systems and lead technical strategy across the organization.
- **Action:** Engineered 800+ feature voter behavior pipeline and deployed production-grade hybrid retrieval (BM25 + Faiss dense vectors) with ColBERT-style re-ranking. Directed PEFT-based LLM optimization across multiple domain-specific tasks.
- **Result:** Delivered **22%** improvement in summarization accuracy; established core ML infrastructure serving multiple specialized clients.
- **FROST (ACL Industry Track 2026, First Author – Accepted):** Introduced FROST, an inference-time optimization framework for production retrieval-reasoning, achieving **28% cost reduction** over Self-Consistency ensembles.

Ramakrishna Mission Shilpamandira, Belur Math, India
Visiting Faculty (AI/ML)

Oct 2025 – Present

- **Situation:** Faculty required upskilling in modern LLM and Generative AI techniques.
- **Task:** Lead a 3-month Faculty Development Program (FDP) for professors.
- **Action:** Architected and delivered specialized curriculum bridging theoretical foundations with modern deployment strategies including Quantization and Knowledge Distillation.
- **Result:** Upskilled professors on LLMs, Generative AI, and production deployment strategies, enhancing institutional research capabilities.

Raapid.ai, Remote 
R&D Intern

Apr. 2025 – Jun. 2025

- **HCC Code Extraction:** Designed and implemented novel deep learning architectures for HCC code extraction from unstructured clinical notes, achieving **12%** improvement in extraction accuracy over baseline.
- **Data Pipeline Optimization:** Developed automated entity-linking module and optimized large-scale medical data pipelines via parallelism and query optimization. Increased knowledge graph entity consistency by **18%**; reduced ingestion and preprocessing time by **25%**.

Projects

Automatic Essay Grading System (*SIT ICOE Hackathon Winner, 400+ Teams*) 

- Engineered novel BiLSTM architecture with multi-head attention and custom hierarchical attention layers to capture essay coherence and structure across multiple linguistic levels.
- Developed domain-specific preprocessing pipeline handling 12,978 essays with advanced feature engineering for syntactic and semantic analysis.
- Achieved MAE of 1.166 and QWK of **0.674** on ASAP/AES datasets, outperforming SOTA by **8.3%** in coherence assessment metrics.

Image Enhancement using Autoencoders 

- Architected multi-scale convolutional autoencoder with sub-pixel convolution and skip connections for single-image super-resolution.
- Engineered hybrid perceptual loss combining SSIM-based structural similarity metrics with deep VGG feature representations for pixel-level and perceptual quality.
- Achieved **42.8%** PSNR improvement over bicubic interpolation; SOTA 31.2 dB PSNR and 0.897 SSIM on DIV2K benchmark.

Multi-Modal Face Tracking and Analysis System 

- Created end-to-end pipeline integrating VGG16 transfer learning with custom detection and classification heads for simultaneous face localization and multi-attribute analysis.

- Applied 15+ Albumentations augmentation techniques (rotation, scaling, color jitter, random crops) for generalization across lighting conditions and pose variations.
- Achieved **30+ FPS** at **95.2%** localization accuracy (IoU p 0.5) on embedded devices; **47%** model size reduction via pruning while maintaining accuracy.

Assistive System for Blind People (*Intel OneAPI Hackathon Winner, 100+ Teams*)



- Built multi-task computer vision system integrating YOLOv9 (obstacle detection), MIDAS (monocular depth, MAE < 10cm), and ResNet50 (Indian currency recognition) with Intel OneAPI optimizations.
- Implemented spatial audio feedback and voice synthesis for real-time navigation guidance; applied quantization and model compression for edge deployment.
- Achieved **98.3%** obstacle and **99.4%** currency recognition at **25 FPS** (14.2MB model, <40ms latency); **92%** task success in real-world testing with 50 visually impaired users.

Image-to-Music Synthesis System (Aurora)



- Engineered modular pipeline fusing 512-D Vision Transformer and CLIP embeddings with color histograms (K-means), edge maps (Canny/Sobel), and texture analysis (GLCM) via cross-attention for visual-to-musical parameter mapping.
- Implemented diffusion-based generators for chord progression, rhythm, and melody synthesis with music-theory harmonic constraints; integrated neural vocoder-style synthesizer using residual upsampling and Butterworth filtering.
- Sub-second encoder latency on GPU; achieved **0.88** Pearson correlation for emotional alignment between visual sentiment and musical valence/arousal.

NeuroAdaptive Affective Decoding from EEG



- Developed adaptive deep learning pipeline for real-time affective state decoding from EEG signals, combining temporal and spectral features for multi-class emotion recognition.

Position of Responsibility

Core Technical Team ML Lead

Feb 2024 – Present

Google Developer Group (GDG), Sister Nivedita University

- Led advanced technical workshops on transformer architectures, attention mechanisms, and modern deep learning frameworks with practical PyTorch and TensorFlow implementation.
- Successfully trained **100+** students in deep learning; multiple students progressed to advanced research projects and industry positions.

Machine Learning Lead

Oct 2023 – Present

SKEPSIS, Sister Nivedita University

- Led 5 major research initiatives spanning NLP and Computer Vision, supervising interdisciplinary teams of 4–6 undergraduate researchers across concurrent projects.
- Designed and deployed production Book Recommendation System using collaborative filtering and k-means clustering, processing 10,000+ entries; improved user engagement by **20%**. Mentored **60+** students with **4** successful project completions.

Skills

Languages: Python, R, C, C++, SQL

Frameworks & Libraries: PyTorch, TensorFlow, Hugging Face, Faiss, Elasticsearch, Scikit-learn, ONNX, OpenCV, Weights & Biases (W&B)

Tools: L^AT_EX, Git/GitHub, SLURM, Tableau, Zotero

Specializations: NLP, Computer Vision, Time Series Forecasting, Reinforcement Learning, Information Retrieval, RAG, Vector Databases

Conferences & Talks

- **ACL Industry Track 2026** – Accepted paper: *FROST: Inference-Time Optimization for Production Retrieval-Reasoning* (First Author)
- **EACL 2026 – AbjadNLP Workshop** – 2 Oral presentations + 1 Poster: [Hausa Ajami Morphology Discovery](#) (Oral), [Sorani Kurdish Dialect Classification](#) (Oral), LLM Orthographic Fragility (Poster)
- **IEEE CIACON 2025**, Durgapur – Presented *HCAT-Net for EEG Emotion Classification* (**Best Paper Award & Best Presenter Award**)
- **QuantumKol 2025 (International Conference on Quantum Technology and Applications)**, Kolkata – Presented *The Glass Policy: A Quantum-Inspired RL Framework for Brittle Mastery in Upskilling*

- **ICDMAI 2025** [🔗](#) – Participated and felicitated; Top 3 Teams in associated offline hackathon (1,000+ teams)
- **BRICS International Big Data & AI Conference** [🔗](#) & **AI Working Committee** [🔗](#), Chongqing, China – International attendee

Awards & Achievements

- **2nd Prize, BRICS International Vocational Skills Competition 2024** [🔗](#), Shandong, China – Represented India among 178 international competitors
- **Top 3 Teams, ICDMAI Offline Hackathon 2025** [🔗](#) – from 1,000+ teams
- **Best Presenter Award, IEEE CIACON 2025** [🔗](#) – Recognized for HCAT-Net EEG Emotion Classification presentation
- **Outstanding AI/ML Research Contribution Award, SNU Connect 2025** [🔗](#)
- **1st Place, SAP ICOE Hackathon 2024** [🔗](#) – from 400+ teams
- **1st Place, Intel OneAPI Hackathon 2024** [🔗](#) and **Pragati 2k24 ML Mania Competition** [🔗](#) – from 100+ teams each
- **Selected for India Regional Bootcamp, Google Solution Challenge 2024** [🔗](#)