Farhan Sheth

farhansheth.acad@gmail.com https://phantom-fs.github.io
https://www.linkedin.com/in/farhan-sheth/ https://github.com/Phantom-fs
Google Scholar ORCID

Research Profile

AI researcher specializing in multimodal, trustworthy, and applied AI, with work spanning privacy-preserving clinical LLMs, vision-language and speech foundation models, synthetic-media forensics, agricultural intelligence, and scientific machine learning. Research interests include multimodal representation learning, geometric learning, explainable and trustworthy AI, healthcare AI, audio deepfake detection, computer vision, and natural language processing.

Education

Manipal University Jaipur, India **2021 - 2025**
Bachelor of Technology (B.Tech.) in Computer Science and Engineering (CSE) *CGPA: 8.53/10.0*

Courses: Artificial Intelligence, Machine Learning, Data Science, Data Mining & Data Warehousing, Soft Computing, Discrete Math, Linear Algebra, Calculus, Probability & Statistics, Relational Database Management System (RDBMS), Computer Architecture, Automata Theory, Object-Oriented Programming (OOP), Data Structures & Algorithms (DSA)

Experience

Nanyang Technological University (NTU), Singapore Jan. 2025 – Present
Research Assistant, Advisor: Dr. Si Yong Yeo

- Developed `HERALD`, a model-agnostic, client-side cryptographic redaction framework for clinical LLM deployment, with medical NER and POS-driven policies to encrypt sensitive tokens while preserving surrounding context.
- Evaluated selective token protection over clinical classification and medical QA workloads, showing that partial securing recovers substantially more downstream utility than fully secured baselines.
- Mapped practical privacy-utility implications for cloud/API-assisted note drafting, triage support, coding workflows, and cross-institution clinical NLP without requiring downstream model architecture changes.

AI4A Lab, India Nov. 2023 – Present
Member; Project Lead (June 2024 – Present) *ai4a-lab.github.io*

- Lead undergraduate research collaborations on applied and trustworthy AI, with projects spanning accessibility, education, speech forensics, agriculture, healthcare privacy, and thermophysical-property prediction.
- Coordinated open research artifacts across datasets, codebases, and web applications, including Agro Companion, Herbify, nanofluid prediction datasets, and deployment-oriented application servers.
- Translated research outputs into reproducible repositories and public demonstrations using Python, PyTorch, Hugging Face, Flask, Google APIs, MongoDB, and GitHub-centered project workflows.

Ulster University, United Kingdom June 2024 – Nov. 2025
Research Intern, Advisor: Dr. Muskaan Singh

- Contributed to speech-forensics research on synthetic speech attribution, open-set detection, and cross-paradigm audio deepfake detection using speech foundation models and geometry-aware representation learning.
- Developed non-Euclidean fusion strategies over hyperbolic, spherical, and Euclidean spaces for synthesis-invariant speech embeddings; work led to IJCNLP-AACL 2025 Main and Findings papers.
- Co-developed **SIGNAL**, combining graph neural networks with k-nearest-neighbor open-set inference to trace synthetic speech to source generators and detect unseen synthesizers; published in EACL 2026.

Manipal University Jaipur, India Jan. 2024 – Oct. 2025
Project Lead, Supervisors: Prof. Sandeep Chaurasia and collaborators

- Led a grant-supported multimodal interview-assessment project for personality and performance prediction, building data-collection and annotation infrastructure for naturalistic video interviews.
- Compiled **RecruitView**, with **2,011** interview clips from **300+** participants and **27,000** pairwise comparative judgments across **12** dimensions covering Big Five traits and interview-performance metrics.
- Designed **Cross-Modal Regression with Manifold Fusion (CRMF)** framework over video, audio, and text embeddings; achieved up to **11.4%** Spearman and **6.0%** concordance-index gains while training with **40–50%** fewer parameters than large multimodal baselines.

National Institute of Technology (NIT) Jalandhar, India June 2024 – July 2024

Research Intern, Advisor: Dr. Rajesh Singla

- Studied machine-learning methods for brain-computer interface (BCI) systems, with emphasis on neural-signal acquisition, preprocessing, and real-time adaptive-feedback workflows for medical applications.
- Reviewed and implemented data-collection procedures for improving neural-signal consistency, then analyzed downstream modeling requirements for robust BCI inference.

Universidad de Las Palmas de Gran Canaria (ULPGC), Spain Dec. 2023 – May 2024

Research Assistant, Advisor: Prof. Carlos M. Travieso-González

- Developed computer-vision models for geological image classification over **950** images covering **19** rock types, using standardized preprocessing and transfer learning with CNN backbones.
- Fine-tuned ResNet, Inception, DenseNet, MobileNetV3, and EfficientNetV2 architectures with k-fold validation; the resulting framework achieved **98%+** top-1 accuracy for rock-type identification.
- Co-authored the resulting Applied Computing and Geosciences article, connecting geological image datasets with reproducible deep-learning classification pipelines.

Selected Publications

- **Sheth, F.***, Yang, Z.*, Lan, Y., & Yeo, S. Y. (2026). Selective Token-Level Cryptographic Redaction for Privacy-Preserving Clinical Deployment of Large Language Models. Under ***review*** at *Artificial Intelligence Review*. *arXiv preprint*. arXiv:2606.03399

- Gupta, A. K.*, **Sheth, F.***, Shaikh, H., Kumar, D., Puniya, A., Panwar, D., Chaurasia, S., & Mathur, P. (2025). RecruitView: A Multimodal Dataset for Predicting Personality and Interview Performance for Human Resources Applications. Under ***review*** at *EMNLP 2026*. *arXiv preprint*. arXiv.2512.00450

- Girish*, Akhtar, M. M.*, **Sheth, F.***, Singh, M., Gerard, J., McClean, P., & Wong-Lin, K. (2026). Synergizing Zero-Shot Cross-Lingual Alzheimer Detection with Language-Invariant Multimodal Bi-Geometric Adversarial Learning. ***Accepted*** to *Interspeech 2026*. *arXiv preprint*. arXiv:2606.17254.

- **Sheth, F.**, Akhtar, M. M., Girish, Singh, M., & Davey, A. (2026). Uc-PrUn: Uncertainty-Calibrated Machine Unlearning using Vision–Language Models for Clinical Decision Support. *ACM Transactions on Computing for Healthcare*. 10.1145/3820497

- Akhtar, M. M.*, Girish*, **Sheth, F.***, & Singh, M. (2026). Bridging Attribution and Open-Set Detection using Graph-Augmented Instance Learning in Synthetic Speech. *Proceedings of EACL 2026*, 5404–5413. 10.18653/v1/2026.eacl-long.250

- **Sheth, F.***, Mathur, P.*, Shaikh, H., Kumar, D., Mishra, S., & Gupta, A. K. (2026). Learning heat: High-fidelity experimental and Koo–Kleinstreuer–Li thermal conductivity predictions in nanofluids via advanced data augmentation and metaheuristic search. *International Communications in Heat and Mass Transfer*, 171, 110022. 10.1016/j.icheatmasstransfer.2025.110022

- **Sheth, F.***, Girish*, Akhtar, M. M.*, & Singh, M. (2025). Curved Worlds, Clear Boundaries: Generalizing Speech Deepfake Detection using Hyperbolic and Spherical Geometry Spaces. *Proceedings of IJCNLP-AACL 2025*, 1923–1932. 10.18653/v1/2025.ijcnlp-long.104

- Girish*, Akhtar, M. M.*, **Sheth, F.**, & Singh, M. (2025). Towards Attribution of Generators and Emotional Manipulation in Cross-Lingual Synthetic Speech using Geometric Learning. *Findings of IJCNLP-AACL 2025*, 635–645. 10.18653/v1/2025.findings-ijcnlp.37

- **Sheth, F.**, Mathur, P., Gupta, A. K., & Chaurasia, S. (2025). An Advanced Artificial Intelligence Framework Integrating Ensembled Convolutional Neural Networks and Vision Transformers for Precise Soil Classification with Adaptive Fuzzy Logic-Based Crop Recommendations. *Engineering Applications of Artificial Intelligence*, 158, 111425. 10.1016/j.engappai.2025.111425

(***** - equal contribution) — Full publication list: Google Scholar

Projects

RecruitView & CRMF | *PyTorch, Transformers, NLP, Web APIs, MongoDB* Dataset | Code

- Built an end-to-end data and modeling pipeline for automated personality and interview-performance assessment, **with over 300+** participants **on** HR-oriented interview prompts, producing **2,011** naturalistic video clips.
- Developed QA-Labeler as a responsive psychologist-informed platform **over** pairwise video comparisons, collecting **27,000** comparative judgments across **12** personality and interview-performance dimensions.
- Designed CRMF with manifold-specific experts and adaptive routing over hyperbolic, spherical, and Euclidean spaces, achieving up to 11.4% higher Spearman correlation and 6.0% higher concordance index with 40–50% fewer trainable parameters than large multimodal models.

HERALD: Privacy-Preserving Clinical LLM | *PyTorch, Transformers, NLP, Cryptography, LLMs* Code

- Designed a token-level cryptographic redaction pipeline **including** medical NER and POS-driven selection **over** clinical text, encrypting sensitive tokens while preserving contextual utility.
- Implemented deterministic ciphertext replacement **on** the client side so protected content remains encrypted during storage, transmission, and model processing without retraining or modifying downstream models.
- Benchmarked `HERALD` **on** public classification and medical question-answering datasets, recovering performance close to plaintext while improving usability relative to fully encrypted baselines.

Agro Companion | *Python, PyTorch, Keras, OpenCV, Flask, MATLAB, Fuzzy Logic* Website | Code

- Developed an AI-powered agricultural tool for identifying **7** soil types and recommending **20+** crops using image-based soil classification with geological and environmental factors.
- Expanded the soil dataset from **1,189** to **8,413** images using CycleGAN augmentation (**7.1**×; **+607.6%**), then fine-tuned CNN and ViT families with grid search, k-fold validation, and ensemble learning.
- Developed a **fuzzy logic-based crop recommendation system** suggesting crops based on **five geological** and **four environmental factors**, with a top-1 accuracy of **92%**, enhancing decision-making for farmers.

Herbify | *PyTorch, OpenCV, Scikit-Learn, Flask, CNNs, Vision Transformers* Website | Code

- Curated and standardized the **Herbify** dataset **with** rigorous cleaning, novel PAHD preprocessing algorithm, and quality control **over 6,104** images representing **91** herb species.
- Fine-tuned CNN and ViT models **on** herb-identification tasks; EfficientNetV2-Large + ViT-L/16 ensemble reached near perfect performance, along with strong precision and high confidence.
- Translated the research framework into an AI-assisted application for real-time herb identification, returning scientific names and resemblance probabilities through an accessible web interface.

Skills

Programming: Python, Java, C/C++, MATLAB, Dart, SQL, HTML, CSS, JavaScript, LaTeX
Machine Learning: PyTorch, Hugging Face Transformers, TensorFlow, Keras, Scikit-Learn, pandas, NumPy, Matplotlib, OpenCV, Albumentations
Methods: Deep learning, computer vision, natural language processing, speech processing, Multimodal learning, LLM/VLM fine-tuning, speech foundation models, geometric learning, CNNs, ViTs, GNNs, manifold learning, fuzzy logic, optimization, data augmentation, evaluation design
Tools: Git/GitHub, Jupyter, Google Colab, Flask, REST APIs, Docker, AWS, Heroku, VS Code, IntelliJ IDEA/PyCharm

Teaching and Service

Undergraduate Teaching Assistant: Supported **50+** students through weekly labs, doubt sessions, quiz preparation, and assignment grading for Data Structures and Algorithms and Data Science and Machine Learning courses.
Peer Review: Reviewer, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2026).
Community Volunteering: Participated in cleanups, fundraising for meals and essential supplies, and social-awareness initiatives on pollution.

Honors and Awards

Student Excellence Award for Research, Undergraduate (**six awards**: Mar. 2024, Apr. 2024, Sept. 2024, Feb. 2025, July 2025, Aug. 2025).
Best Minor Project in Research Category, Undergraduate (Sept. 2024).