ALI RAZA

mirzaalirazafsd@gmail.com ali-raza-0a1282281 alyrraza hf/alyrraza alyrraza.vercel.app
AI Research Engineer · Medical Imaging AI · Evaluation & Reproducibility · Computer Vision

RESEARCH INTERESTS

Trustworthy Evaluation of Medical Imaging AI · Model Selection Under Scarce Labels · Medical Image Registration · Vision-Language Models for Radiology · Generative Models · Distribution Shift · MRI Reconstruction

EDUCATION

FAST-NUCES, Faisalabad, Pakistan 2022 – Dec. 2026 (expected)
BSc Computer Science CGPA: 2.9 / 4.0 Thesis: RadGuard, Multimodal Error Detection in AI-Generated Radiology Reports

PUBLICATIONS & ABSTRACTS

- **Raza A**. "Inter-Lobe Boundary Channels and Localised Folding Repair for Expiration to Inspiration Lung CT Registration." *Learn2Reg 2026 Challenge, MICCAI 2026*, Strasbourg. **Accepted** (peer-reviewed method description, single author), September 2026. OpenReview: nFNQaSDcO6
- Khan MM, **Raza A**, et al. "Generalizable Detection of CIEDs Using a Hybrid Transformer and CNN Framework." *ACR 2026, American College of Radiology Annual Meeting*, Washington DC. **Accepted** (ePoster)
- **Raza A**. Extended Learn2Breath registration paper. *Journal of Medical Imaging*, Special Section on Deformable Image Registration (invited submission). **In preparation**

RESEARCH EXPERIENCE

Learn2Breath: Deformable Chest CT Registration, MICCAI 2026 Learn2Reg Challenge 2026
Registration · NLST Cohort · ConvexAdam · TotalSegmentator · MIND-SSC · Dense Displacement Fields · Docker

- Built an optimisation-based, label-free registration pipeline (no trained registration weights) combining self-generated TotalSegmentator lobe and inter-lobe boundary channels with MIND-SSC features in ConvexAdam; lobar Dice **≈0.97**, **%NDV 0.0035** on the validation set, 69.3 s per pair on a Tesla T4.
- Designed a localised Jacobian folding repair stage that removes almost all folding at negligible Dice cost; ranked **3rd on the public validation leaderboard** (top-4 team).
- Found that folding measured on a downsampled field under-reports %NDV relative to full resolution; every reported figure now states its evaluation resolution.
- Single-author method description **accepted after peer review**; invited **oral presentation** and poster at Learn2Reg 2026 (MICCAI 2026, Strasbourg); invited full paper to the *Journal of Medical Imaging*.

Overtuning Under Scarce Evaluation: A Pre-Registered Study of Hyperparameter Selection Transfer 2026
Pre-Registration · Nested Resampling · Permutation Null · Bootstrap CIs · NumPy/SciPy · pytest · Kaggle (Tesla T4)

- Asked whether registration settings chosen as best on **10 annotated NLST pairs** stay best on cases never used for selection; fixed the objective, statistic, splits, **2000 bootstrap resamples** and **1000 permutation replicates** in a written pre-registration before the first run.
- Scored **100 configurations** on the 10 pairs with the frozen pipeline (1000 rows, none failed); replayed every select-on-k, evaluate-on-rest split (k = 3, 5, 7; 120, 252, 120 splits) against a permutation null. For the folding-dominated objective, selection kept **0.99** of the available gain (null median 0.80); for Dice alone the estimate was 0.68 with a wide interval (0.13 to 0.82).
- Label-free check on **30 unlabelled pairs**: the selected setting folded **2.9 percentage points less** than the median of 15 random settings (95% CI 2.4 to 3.4); reported as *no overtuning observed within these bounds*, with limitations stated.

RadGuard: Multimodal Error Detection in AI-Generated Radiology Reports 2025 – 2026
Final Year Thesis · FAST-NUCES BioViL-T · CXR-BERT · MIMIC-CXR · FastAPI · Kubernetes · AWS EC2

- Designed a multimodal VLM detecting 4 error classes (Supported, Hallucinated, Missing, Inaccurate) in AI-generated chest X-ray reports: **Val F1 = 0.66** on 74,060 MIMIC-CXR samples across 14 conditions.
- Proposed **ELRRS** (Error-Level Radiology Report Score), a weighted report reliability metric fusing ResNet-50 and CheXBERT signals, aligned with ACR and RadGraph conventions.
- Implemented **bidirectional cross-attention** (BioViL-T + CXR-BERT) with multi-task pseudo-mask supervision; explainability via **GradCAM + CBAM**; FAISS-based RAG assistant; deployed with FastAPI on Kubernetes (AWS EC2).

CIED Detection Pipeline, ACR 2026 Accepted Abstract 2024 – 2025
DINOv3 · DenseNet-121 · TransU-Net · nnU-Net · SAM · Vision Transformers

- Three-stage pipeline for pacemakers and ICDs in chest X-rays: DINOv3-based segmentation isolates the device; hybrid dual-branch classifier identifies the manufacturer (CLS + pooled patch tokens, 1536d); hierarchical model classifier conditioned on a 64d manufacturer embedding.
- Handled class imbalance with focal loss, weighted sampling, Mixup and manufacturer-conditioned SMOTE in the hybrid feature space.

E2E-VarNet MRI Reconstruction: Uncertainty-Based Drift Signal Under Distribution Shift 2025 – 2026
Reproduction + Extension · Sriram et al. MICCAI 2020 PyTorch · fastMRI · Volume-Level Block Bootstrap · Kaggle/Vast.ai

- Reproduced E2E-VarNet for 4x-accelerated single-coil knee MRI (T = 4/6/8 cascades; SSIM 0.7594/0.7606/0.7607 knee, 0.6622/0.6705/0.6773 brain OOD) and tested three extensions with volume-level clustering, bootstrap CIs and trivial-baseline controls.
- **Headline result:** checkpoint-ensemble uncertainty rises under knee→brain shift, **1.54x, bootstrap 95% CI [1.35x, 1.74x]**, computable on 100% of tested volumes (a competing eSNR signal on only 12/15).
- Corrected my own pseudoreplication (slice-level → volume-level bootstrap), documented 19 debugging findings, and reported the cascade-depth and self-supervised TTA extensions as honest near-null and negative results.

ProgressGAN-CXR: Severity-Conditioned GANs for Disease Progression Simulation 2025 – 2026
DCGAN · WGAN-GP · Spectral GAN · Knowledge Distillation · PyTorch · FastAPI · React · HuggingFace

- Introduced continuous severity conditioning into GAN-based chest X-ray synthesis (Normal → COVID-19); showed FID and Spearman r capture different properties and must be evaluated separately.
- Knowledge distillation from a WGAN-GP teacher: KD generator reaches **FID 143** (matched to DCGAN) with **Spearman r 0.984**. Live: progressgan-cxr.vercel.app

TECHNICAL SKILLS

Medical AI: BioViL-T, CXR-BERT, CheXBERT, TotalSegmentator, ConvexAdam, MIND-SSC, TransU-Net, U-Net, nnU-Net, SAM, DINOv3, ViT
Evaluation & Statistics: Pre-registration, bootstrap CIs, permutation tests, nested resampling, Wilcoxon signed-rank
Generative Models: GANs (DCGAN, WGAN-GP), Knowledge Distillation, Diffusion, Flow Matching
CV, NLP & Agents: YOLO (v8/v11), ResNet, DenseNet, GradCAM, CBAM, Transformers, LoRA/QLoRA, RAG, FAISS, LangGraph, MCP
Stack: PyTorch, MONAI, FastAPI, React, Docker, Kubernetes, AWS EC2, MLflow, pytest, Git, Kaggle, Vast.ai
Datasets: NLST (Learn2Breath), MIMIC-CXR (PhysioNet credentialed), NIH ChestX-ray14, MS-CXR, CheXpert, fastMRI

INDUSTRY & APPLIED AI EXPERIENCE

Freelance AI / ML Engineer | HusQuay Solutions 2023 – Present
Production AI Deployment · Voice AI Agents · Multimodal Systems · RAG Pipelines

- Deployed a GPU-backed **YOLO multimodal detection system** (FastAPI + Vast.ai + React/Vercel) for automated asset depreciation detection in production.
- Built a **voice-driven B2B lead generation agent** (LiveKit STT/TTS, n8n, MCP) and a **voice-cloned PDF RAG assistant**, both production-deployed.
- Built a **3-agent LangGraph system** (FastAPI MCP supervisor, LiveKit voice interview agent, MLflow tracking); passed 20/20 integration tests.

Data Analyst Intern | TAKMIL 2024
Louisville, Kentucky, USA · Remote

- EDA on 100K+ row educational datasets; **Power BI** dashboards; pipeline optimisations improving efficiency by **30%**.

SERVICE & ACTIVITIES

- **Co-Head Hackathon**, ACM Society, Daira 2025 (organised and led a university-wide AI hackathon)
- **Member**, ACM Society Management Team, Daira 2024