

SAMUEL AKINWUMI ADENIJI

AI for Healthcare · Medical Vision–Language Models · Medical Image Analysis

sadeniji@andrew.cmu.edu | Carnegie Mellon University | Website · Google Scholar · LinkedIn · GitHub

RESEARCH INTERESTS

Medical vision–language models for 3D imaging; token-efficient and multi-scale architectures for volumetric data; representation learning for tumour classification and radiology report generation; deployment of imaging AI in low-resource and African clinical settings.

EDUCATION

Carnegie Mellon University — M.S., Engineering Artificial Intelligence *Aug 2025 – May 2027*

Concentration: AI in Healthcare (Research) · QPA: 3.94/4.00

Federal University of Technology, Minna, Nigeria — B.Eng., Mechatronics Engineering *Sep 2018 – Jan 2025*

First Class Honors (4.60/5.00), Concentration: Embedded Systems and Machine Learning.

RESEARCH EXPERIENCE

Research Intern — Language and Spatial Laboratories *June 2026 – Present*

Carnegie Mellon University

- Conducting research on vision–language models for 3D medical imaging, developing ICTC, an instruction-conditioned token-compression architecture for chest-CT report generation.
- Designed and implemented the full research pipeline: data ingestion, dual-stage compression projector, LoRA-adapted language backbone, and YAML-configured multi-task training.
- Built the evaluation harness benchmarking multi-label classification and report-generation quality across six visual token budgets.

Research Intern (Volunteer) — AI for Healthcare Research Lab *Jan 2026 – May 2026*

Carnegie Mellon University Africa

- Developed machine learning methods for cardiovascular disease detection from ECG signals.
- Supported studies on aortic stenosis, aortic regurgitation and related cardiac conditions.

Research Assistant (Intern) *Oct 2023 – Jan 2025*

Artificial Intelligence for Clean Energy Ltd., Minna, Nigeria

- Built a Flask and Random Forest energy-monitoring platform (Google Cloud, 99% uptime), improving monitoring accuracy by 30%.
- Piloted drones for field data acquisition, reducing collection time by 30% and improving data quality by 12%.

TEACHING

Teaching Assistant — Cloud Computing (15-319 / 15-619)

Spring 2026

Carnegie Mellon University

- Supported a 200+ student course spanning undergraduate, MS, and PhD levels; held office hours, delivered course-overview sessions, and graded assignments at scale.

Teaching Assistant — AI System Design (04-652)

Fall 2026

Carnegie Mellon University Africa · Selected / Incoming

- Selected as Teaching Assistant for the Fall 2026 offering (edge-to-cloud AI systems: embedded ML on Raspberry Pi/Arduino and cloud-based model analysis).

Peer Instructor (Volunteer) — Introduction to Machine Learning

Spring 2026

Carnegie Mellon University

- Tutored classmates through core ML topics (regression, classification, model evaluation, and the underlying linear algebra and probability) in small-group sessions.

RESEARCH PROJECTS

ICTC: Instruction-Conditioned Token Compression for 3D CT Vision— Language Models

June 2026 – Present

Language and Spatial Laboratories, CMU

Stack: Python, PyTorch, Vision Transformers, Perceiver, LoRA/PEFT, HuggingFace, CT-RATE v2

- Designed a dual-stage token-compression architecture for 3D chest-CT report generation: a frozen ViT-Base slice encoder, a task-agnostic Perceiver-style intra-slice stage compressing each slice to $K=32$ latents, and a FiLM-conditioned inter-slice stage that compresses to M instruction-aware visual tokens consumed by a LoRA-adapted OPT-1.3B backbone.
- Trained only the projector, visual bridge and LoRA adapters (13.4M parameters; 0.24% of the LLM), keeping the ViT encoder and language backbone frozen for stability and efficiency.
- Built a multi-task instruction pipeline (report generation, entity-conditioned response, binary abnormality QA) with YAML-configured training and a 368-test suite for reproducibility.
- Developed a six-budget evaluation harness ($M=16-512$) reporting multi-label AUROC/F1 over 18 abnormality labels alongside RadGraph-F1 and RaTEScore for generation quality.
- Completed first end-to-end training and evaluation run on A100; diagnosed a greedy-decoding degeneracy and resolved it at inference time without retraining.

FPN-Mamba: Multi-Scale Feature Pyramid and Bidirectional Mamba for Pediatric Medulloblastoma MRI Classification

Feb 2026 – Present

Stack: Python, PyTorch, Feature Pyramid Networks, Mamba (state-space models), scikit-learn

- Developed a multi-scale FPN combined with a bidirectional Mamba block for tumour subtype classification on a 761-image pediatric medulloblastoma MRI dataset, achieving AUC 99.35% and F1 94.54%.
- Designed augmentation, normalisation and class-reweighting pipelines to address subtype imbalance and improve minority-class recall.
- Prepared as a methods and feasibility paper with ablations and statistical reporting for a MICCAI 2026 workshop, written in LaTeX under double-blind conventions.

White Blood Cells as a Source of False-Positive Malaria Detection in Thick Blood Smears

Feb 2026 – June 2026

Stack: Python, PyTorch, Ultralytics YOLOv8, OpenCV, SciPy, DBSCAN, Weights & Biases

- Designed controlled detection experiments comparing parasite-only and parasite-plus-WBC supervision on Giemsa-stained thick blood-smear images.
- Conducted spatial error analysis using Euclidean distance metrics and DBSCAN clustering to test whether false positives are systematically associated with white blood cell structures.
- Built an evaluation pipeline combining precision metrics, IoU filtering and Mann–Whitney U testing to quantify the effect of WBC supervision on diagnostic reliability.

Condition-Aware ICU Early-Warning System

Feb 2026 – May 2026

Stack: Python, PyTorch, TensorFlow, MIMIC-IV, FastAPI, Docker, WebSockets, scikit-learn

- Built and benchmarked six disease-conditioned deterioration-risk architectures (LSTM, GRU, TCN, BiLSTM with attention, Transformer, TFT-style) over MIMIC-IV ICU time-series, integrating vital signs, laboratory abnormalities, diagnosis conditioning and treatment exposures; the best model (BiLSTM+attention, ROC-AUC 0.76) is the one served.
- Packaged the model as a FastAPI service with a WebSocket endpoint for streaming inference, containerised with Docker and deployable via Docker Swarm; verified end to end and structured for reproducibility under PhysioNet data governance.

PUBLICATIONS

1. K. E. Jack, **S. A. Adeniji**. “Convergence of Blockchain, Smart Microgrid and Energy Market.” *Electrica*, 24(2), 367–374, 2024. DOI: 10.5152/electrica.2024.23132.
2. K. E. Jack, **S. A. Adeniji**. “Development of an Artificial Intelligence-Based Infant’s Emotion Recognition System for a Caregiving Robot.” DOI: 10.5455/NJEAS.209325.

ACADEMIC SERVICE AND RESEARCH TRAINING

- **SPRINT / SPARK AI for African Medical Imaging Knowledge Translation**, 2026 Cohort — state-of-the-art methods for glioma brain-tumour segmentation and translation of imaging AI into African clinical practice; led a capstone on malaria parasite detection.

TECHNICAL SKILLS

- **Languages and Platforms:** Python, Linux/Bash, Git, L^AT_EX; Microsoft Azure, Google Cloud Platform.
- **Deep Learning:** PyTorch, TensorFlow, Vision Transformers, Mamba / state-space models, CNNs and Feature Pyramid Networks, vision-language models, LoRA / parameter-efficient fine-tuning, HuggingFace.
- **Medical Imaging and Data:** 3D CT and MRI pipelines, CT-RATE, MIMIC-IV, MONAI-style preprocessing, OpenCV, YOLO, scikit-learn, Weights & Biases, Docker.
- **Foundations:** Linear algebra, probability and statistics, mathematical modelling, experimental design.

AWARDS AND SCHOLARSHIPS

- **Mastercard Foundation Scholarship**, Carnegie Mellon University (2025) — toward the M.S. in Engineering AI.
- **MINDS Scholarship** (Mandela Institute of Developmental Studies), Carnegie Mellon University (2026).