

# Sarah Pendhari

(412)-689-2849 | [spendhar@andrew.cmu.edu](mailto:spendhar@andrew.cmu.edu) | [linkedin.com/in/sarah-pendhari](https://www.linkedin.com/in/sarah-pendhari) | [Portfolio](#) | [Google Scholar](#)

## EDUCATION

### Carnegie Mellon University

Master of Science, Computer Vision, School of Computer Science GPA: 3.33/4.0

Coursework: Visual Learning and Recognition, Adv Computer Vision, ML, Adv DL (PhD), Robot Localization & Mapping

Pittsburgh, PA

Dec. 2026

### University of Mumbai

Bachelor of Engineering, Computer Engineering GPA: 9.22/10.0 | Principal's Excellence Award (top 0.03%).

Mumbai, India

Jul. 2025

## SKILLS

**Programming Languages:** Python, C, C++

**Tools and Frameworks:** PyTorch, Tensorflow, NumPy, Pandas, SciPy, OpenCV, Keras, Git, Linux, Docker, AWS (EC2), SQL

**AI Related Tech Stack:** NLP, LLMs, VLMs, RAG, Diffusion, Fine-tuning, Multi-modal, Image/Video Generation

## EXPERIENCE

### Team Chiron – The AIRLab at CMU | Prof. John Galeotti (Capstone Project)

Research Assistant

Pittsburgh, PA

Jan. 2026-Present

- Developing a VLM for UAV RGB video as part of the **DARPA Triage Challenge** (U.S. federal-funded), enabling person tracking, re-identification, and temporally grounded event understanding in complex, real-world scenes.
- Leading large-scale data pipeline and annotation setup (Label Studio) for **UAV video**, supporting robust evaluation of multimodal perception and long-horizon reasoning.

### CMU Tepper School of Business | Prof. Mohsen Foroughifar

Research Assistant

Pittsburgh, PA

Nov. 2025-Jan 2026

- Built a predictive video narrative model combining VLMs, LLM-based future hypothesis generation, and probabilistic uncertainty estimation, scaled on the Pittsburgh Supercomputing Center (PSC) to quantify anticipation (suspense) and prediction error (surprise) in video.

### Carnegie Mellon University – Xu Lab | Prof. Min Xu

Research Assistant

Pittsburgh, PA

Oct. 2025-Present

- ModelFabric (ICML 2026, submitted):** Proposed an explainability-guided selective shielding framework for **VLA models on edge robots**, enabling secure, low-latency deployment via optimized partial TEE execution (up to **4.7x faster** than full secure inference).

### University Of Mumbai | Prof. Nazneen Pendhari

Research Intern

Mumbai, India

Jan. 2022 – Mar. 2025

- Project 1 – ColourViTGAN:** Investigated data-efficient visual representation learning, developing a ViT-CycleGAN hybrid that improves semantic colorization quality (**+2.2 PSNR, 0.982 SSIM**) for **240K+ images**, enabling low-latency deployment. Best Paper IEEE ICUIS Link
- Project 2 – Retinal Diagnosis via FSL:** Designed attention-based Few Shot Learning models for medical imaging, achieving **0.947 AUC** with only **5% labeled data**, scaling efficiently via **mixed-precision multi-GPU training** (V100). Best Paper IEEE IATMSI Link

### IIT Bombay | Prof. Surya Durbha

Summer Intern

Mumbai, India

Jun. 2024 – Aug. 2024

- Programmed embedded C/C++ firmware on ESP32**, optimizing LoRa connectivity for real-time streaming of 10K+ daily sensor signals.
- Built a **low-latency** forecasting model for environmental signals, achieving **+32.3%** over baseline via adv signal processing and ensemble learning, worked in a team of 5.

### Wondrlab India Pvt. Ltd

Software Developer Intern

Mumbai, India

Jun. 2023 – Sep. 2023

- Optimized Airflow-based ML data pipelines, boosting ETL throughput by **9.2%** with parallel orchestration and feature reduction.
- Built a high-concurrency scraping + MongoDB ingestion system enabling **real-time analytics dashboards** for cross-team insights.

## ACADEMIC PROJECTS

### CardioCare: AI for Cardiac Risk, Final Year Thesis | University of Mumbai

Sep. 2024 – Jun. 2025

- Built a multimodal cardiac diagnosis assistant using **RAG on patient records** for information retrieval + **ECG + medical text**, achieving **76.8% EM and 0.83 F1**, outperforming GPT-3.5 by **15%** on clinical benchmarks.
- Fine-tuned an **LSTM arrhythmia classifier (48 MIT-BIH classes)** with chain-of-thought prompting, reaching **0.875 AUC** and leading a **4-member team** from model design to validation.

## PUBLICATIONS

- ModelFabric: Explainability-Guided Selective Shielding for Secure Vision-Language-Action Models on Edge Robots, ICML '26 Submitted
- Attention-Enhanced Prototypical Networks for Few-Shot Microaneurysm Detection in Diabetic Retinopathy Images, IEEE, May 2025
- ColorViTAN: A hybridised approach using Vision Transformers and CycleGAN to add color to greyscale images, IEEE, Jan 2025
- Benchmarking Deep Learning Models for Automated MRI-based Brain Tumor Detection: In-Depth Analysis of CNN, VGG16, VGG19, ResNet-50, MobileNet, and InceptionV3 – IJCA, Nov 2024
- Advanced Neural Network-Based Color Transformation for Enhanced Visual Perception in Tritanopia Using Deep Learning Algorithms Taylor and Francis, Submitted Jan 2025