

Adithya Narayan

Pittsburgh, PA 15206

(412) 737-4602 | anaraya2@andrew.cmu.edu | linkedin.com/in/adithya-n-b637b8146

Education

Carnegie Mellon University

Master of Science in Computer Vision – Internship Start Date: May 12, 2025

December 2025

Pittsburgh, Pennsylvania

Manipal Institute of Technology

B. Tech in Electronics and Communication | GPA: 8.92/10

July 2021

Karnataka, India

Research Experience

Human Sensing Laboratory | Meta Reality Labs

September 2024 – Present

Graduate Research Assistant | Advisor: [Prof. Fernando De la Torre](#)

Pittsburgh, PA

- Collaborating with Meta to enhance the robustness of SoTA depth estimation and semantic segmentation models on long-tail distributions by adversarially sampling 3D scene parameters (illumination/sensor parameters) via differentiable renderers.
- Developing a synthetic data pipeline in Blender to animate trajectories, bake and extract texture and depth maps from 3D assets.

CMU Lab4D++ (Opensource)

September 2024 – Present

Research Collaborator | Advisor: [Dr. Gengshan Yang](#) | [link](#)

Pittsburgh, PA

- Investigating the impact of generative image priors on improving the performance of Gaussian Splatting on 4D reconstruction tasks.
- Exploring future directions by building a sim-2-real pipeline in Unity to model free-form human/hand-object/scene interactions.

Professional Experience

Arintra

March 2023 – July 2024

Machine Learning Engineer

Bangalore, India

- Achieved a 6% improvement in ICD code prediction accuracy by applying Retrieval-Augmented Generation (RAG) to medical LLMs, leveraging a human-curated corpus to query and extract relationships between disease co-morbidities.
- Designed and deployed a system that leveraged SapBERT word embeddings in Qdrant, a vector database, to query for semantically similar medications, improving system's F1-score by 11% and improving model generalizability across 4 hospitals.
- Led the development of a CI/CD system using MLFlow, FastAPI, and GCP to manage the lifecycle of ML model deployment, monitor model drift, and track evaluation metrics, ensuring zero downtime while processing 2.8 million patient charts a month.

Klothed

February 2022 – March 2023

Machine Learning Engineer | Advisor: [Prof. James O'Brien](#)

New York, USA

- Built a PyTorch3D pipeline merging sparse-view 3D reconstruction models (e.g., ECON) with diffusion-based inpainting models to generate 3D clothed avatars, driving consumer click-through rates by 12% and boosting monthly revenue by 7%.
- Developed and optimized a 2D image warping (optical flow) pipeline [link](#) by 90% (from 2s to 0.2s per inference) by leveraging finite-element (FEM) techniques, SMPL-X priors and vectorization in SciPy, OpenCV and NumPy.
- Led synthetic data efforts, building a Blender-Bpy pipeline to render diverse clothing conditions (lighting/draping), improving image matting and super-resolution performance by 3% overall, 13% on long-tail cases, and boosting downstream SMPL-X fits.

Origin Health

November 2020 – November 2021

Research Engineer | Advisor: [Dr. Sripad Devalla](#)

Raffles Quay, Singapore

- Co-authored [link](#) an approach combining domain-specific synthetic data and a novel heatmap-based attention mechanism to achieve a 3.8% reduction in MAE for fetal biometry measurements compared to SoTA approaches.
- Spearheaded a project utilizing GANs and domain adaptation to generate synthetic data samples, improving segmentation mIOU by 20% on legacy ultrasound sensors and enabling safe deployment across 20+ clinics and 5 generations of ultrasound scanners.

Research Intern | Advisor: [Dr. Sripad Devalla](#)

- Published a multi-scale kernel backbone [link](#) that improved segmentation mIOU on 10 fine-grained fetal structures by 4%.

Skills

Languages: Python, C++, Java, Bash

Libraries: PyTorch, PyTorch3D, Torch-TensorRT, OpenCV, Tensorflow, Keras, Pandas, Docker, Numpy, SciPy, GCP, Redis

Publications

- **Biometric Constraint Attention Masks to obtain Sonographic Measurements** - Shankar, H., [Narayan, A.](#), Jain, S., Singh, ... & Devalla, S; IEEE 19th International Symposium on Biomedical Imaging (ISBI), 2022. [Paper](#)
- **Device-independent deep learning system for the automated segmentation of sonographic fetal brain structures** - Lad, A., [Narayan, A.](#), Shankar, H., ... & Devalla, S. K. (2022, April); SPIE 2022. [Paper](#)
- **OC11.02: Validation of a deep learning system for the automated segmentation of fetal brain structures** - [Narayan, A.](#), Kaushik, S., Shankar, H., ... & Devalla, S; ISUOG 2021. [Paper](#) | [Supporting Material](#)