

Adithya Narayan

Pittsburgh, PA | 412-737-4602

anaraya2@cs.cmu.edu | [linkedin.com/in/adithya-n-b637b8146](https://www.linkedin.com/in/adithya-n-b637b8146) | [adithyaknarayan.github.io](https://github.com/adithyaknarayan) | [Google Scholar](https://scholar.google.com/citations?user=adithyaknarayan)

Education

Carnegie Mellon University - Robotics Institute

Master of Science in Computer Vision | GPA: 4.11/4.0 | Start: January 2026

August 2024 - December 2025

Pittsburgh, Pennsylvania

Manipal Institute of Technology

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

July 2017 - July 2021

Karnataka, India

Research

Human Sensing Laboratory | Meta Reality Labs

September 2024 – Present

Pittsburgh, PA

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

- Exploring how 2D VLMs gain 3D scene understanding through multi-view reasoning and view selection, integrating Gaussian Splatting + depth representations, and analyzing the emergence of this 3D understanding via mechanistic interpretability.
- Developed an adversarial scene exploration framework combining SE(3) manifold optimization and ordinal-loss functions to expose failure modes in depth estimation and geometry reconstruction (CVPR 2026).

Experience

HeyGen

May 2025 – August 2025

Research Engineering Intern

Los Angeles, California

- Fine-tuned a camera-motion ControlNet using point-cloud renders and pose inputs to guide multi-modal Video Diffusion Models (Wan 2.1), producing smooth, spatially consistent camera trajectories for AR content. [[link](#)]
- Leveraged sparse background optical flow and early fusion to develop a multi-modal classifier to classify camera motions as part of the internal data filtration pipeline - increasing the pass rate by 10% and data purity by 20%.
- Built and parallelized a large-scale SfM pipeline (using VGGSfM and dynamic scene segmentation) to extract and align camera poses from 100K+ in-the-wild videos.

Arintra

March 2023 – July 2024

Machine Learning Engineer

Bangalore, India

- Improved ICD code prediction accuracy by 6% via Retrieval-Augmented Generation (RAG) with medical LLMs leveraging disease comorbidity knowledge.
- Designed a semantic search engine using SapBERT + Qdrant vector DB for medication retrieval, boosting F1 by 11% and generalizability across 4 hospitals.
- Developed a model versioning and deployment system using MLFlow, FastAPI, and GCP, streamlining engineer workflows, cutting model deployment time by ~20–25%, and ensuring reproducible, traceable releases of finetuned models.

Klothed

February 2022 – March 2023

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

New York, USA

- Enhanced texture fidelity for single-view 3D human mesh reconstruction (ECON) using diffusion-based texture synthesis, image super-resolution and synthetic data augmentations.
- Accelerated a 2D image warping pipeline by 90% (2s → 0.2s) through finite-element optimization, enabling real-time inference for AR applications.
- Proposed a synthetic data pipeline in Blender to render diverse clothing conditions (lighting, draping), resulting in ~2% improvement in MSE for image matting models.

Origin Health

November 2020 – November 2021

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

Raffles Quay, Singapore

- Co-authored [[link](#)] a paper, 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 (IEEE ISBI 2022).

Publications

- [Narayan, A.](#)*, Ojha, U*. Theiss, J., Prakash, A., De la Torre, F. Breaking Depth Estimation Models with Semantic Adversarial Attacks. CVPR 2026 (Under Review)
- Shankar, H., [Narayan, A.](#), Jain, S., Singh, A., Devalla, S. K. Biometric Constraint Attention Masks to Obtain Sonographic Measurements. IEEE ISBI 2022. [Paper](#)
- Lad, A., [Narayan, A.](#), Shankar, H., Jain, S., Devalla, S. K. Device-Independent Deep Learning System for the Automated Segmentation of Sonographic Fetal Brain Structures. SPIE Medical Imaging, 2022. [Paper](#)

Skills

Languages: Python, C++, Bash

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