

Ahish Deshpande

 [AhishD](#) |  ahishd@andrew.cmu.edu |  412-430-1470 |  [Ahish9009](#) |  [Portfolio](#)

EDUCATION

Carnegie Mellon University
Master of Science in Computer Vision

Aug 2025 - Dec 2026
Pittsburgh, PA

- CGPA: **4.11*/4.0**
- Courses: Advanced Computer Vision, Learning for 3D Vision, Robot Learning, Physics based Rendering, Robot Localization & Mapping, Deep Learning Systems, AI for Scientific Computing

International Institute of Information Technology, Hyderabad

Jul 2018 - Jul 2022

Bachelor of Technology (Honors) in ECE

Hyderabad, IN

- CGPA: **9.38/10.0** (Recognized in the **Dean's list** in 7 semesters)
- Courses: Data Structures, OS, Advanced NLP, Mobile Robotics, Statistical Methods in AI, Distributed Systems

EXPERIENCE

Applied AI/ML SWE Intern
Microsoft

May 2026 - Aug 2026
Redmond, WA

- Built a **memory sanitizer/profiler** from scratch for Microsoft's **MAIA Triton** accelerator compiler for **MLIR** with custom dialects. Used in conjunction with a custom harness to report **4+ optimizations**
- **Optimized** kernels for Google's Gemma 2 model to run **2x** faster on the **MAIA chip** versus PyTorch compilation

Graduate Researcher
Pittsburgh Pirates

Jan 2026 - Present
Pittsburgh, PA

- Finetuned & benchmarked vision models (VLM, VideoMAE) for **interpretable, few shot classification**
- Built an end-to-end **production** ready pipeline for data processing & model inference on a dataset of **10k+** videos

Graduate Researcher
Human Sensing Lab

Aug 2025 - Dec 2025
Pittsburgh, PA

- Developed a framework for using foundation image generative models to guide targeted, sample-efficient data collection, achieving up to a **13% AP50** improvement while drastically reducing required training data

Software Engineer II/III
Google

Jul 2022 - Aug 2025
Bangalore, IN

- Coordinated complex changes across a hierarchy of **8+ systems** to help **double** Google's **TPU** delivery capacity
- Built and maintained big data pipelines to process over **60 billion** records/hour to facilitate resource planning
- Reduced pipeline computation time from **14h to 2h** by optimally partitioning shards to reduce stragglers

Research Intern
Adobe Research

May 2021 - Aug 2021
Remote, IN

- Trained language models using **domain adaptation** & auxiliary tasks for **NER** in long form documents

Undergraduate Researcher

Dec 2020 - Jul 2022

Robotics Research Center & Center for Visual Information Technology, IIIT-H

Hyderabad, IN

- Developed obstacle avoidance algorithms using **optical flow** based **visual servoing**. Advisor: Dr. K Madhava K.

SELECTED PROJECTS

RL Safety & Interpretability | *Reinforcement Learning, Control Barrier Functions, Safety*

Implemented & interpreted Lyapunov and Hamilton-Jacobi CBF guarantees for Actor-Critic and PPO Lagrangian

Custom PyTorch (Needle) | *Systems, Autograd, CUDA Programming*

Implemented a subset of PyTorch from scratch, including autograd, NDArrays, common DL architectures, etc.

Point Cloud Completion | *Point Clouds, Computer Vision, Machine Learning*

Implemented multiple unsupervised ML models for representation to improve performance on point cloud completion.

PUBLICATIONS

Panev, S., Jeon, M., Khindkar, V., **Deshpande, A.**, Melo, C. M. de, Hu, S., Chakraborty, S., and De la Torre, F. (Oct. 2026). "Diagnosing Aerial-View Object Detectors with Foundational Image Generative Models". In: *ECCV*.

SKILLS

Programming Python, C, C++, Java, MATLAB, Bash, HTML, CSS, Javascript

Tools/Frameworks PyTorch, Tensorflow, OpenCV, Open3D, Numpy, Jax, Git, Flask, Docker, Hadoop, Linux

ACHIEVEMENTS

- Served as an **external reviewer** for **CVPR 2021**.