

# Anirudh Chakravarthy

📞 (412) 721 8285 • ✉ anirudh.s.chakravarthy@gmail.com  
🌐 anirudh-chakravarthy.github.io • in anirudh-chakravarthy

## Education

### Carnegie Mellon University, School of Computer Science

Master of Science in Computer Vision (Robotics Institute)

GPA: 4.14/4

Pittsburgh, PA

Dec 2022

### Birla Institute of Technology and Science

Bachelor of Engineering in Computer Science (Minor: Data Science)

GPA: 9.55/10; Department Rank: 9/153; Top 2% Merit Scholarship

Pilani, India

May 2021

## Experience

### Cruise LLC

Machine Learning Engineer

- Implemented a transformer and multi-task methods (GradNorm, PCGrad) to improve temporal stability in classification
- Introduced a camera-lidar fusion method to improve gesture recognition by 10% and pedestrian falling classification by 15%
- Developed new fine-grained attributes (children, strollers wheelchairs, etc) with > 90% recall; led end-to-end behavior integration; introduced miners for continual learning from misclassification to improve caution around pedestrians by 5x
- Recognized as an **outstanding engineer** in the team for leading 5 safety-critical model releases and most code contributions

San Francisco, CA

Feb 2023 – Present

### Cruise LLC

Research Intern

- Implemented a multi-task learning framework to adaptively balance task weights during training; improved precision performance on school bus classification by 7% and delivery vehicle classification by 3% over production model
- Formulated a task-incremental method to dynamically balance task gradients while maintaining performance on old tasks, achieving SOTA on CIFAR and Waymo Motion Dataset on task incremental settings; under review at **ECCV'23**

San Francisco, CA

May 2022 – Aug 2022

### CMU Argo AI Center

Research Collaborator (mentor: Dr. Deva Ramanan)

- Introduced Open-World LiDAR Panoptic Segmentation and an evaluation protocol to benchmark open-world generalization
- Developed a panoptic segmentation method combines SOTA semantic segmentation with non-learned hierarchical grouping
- Improved closed-set SOTA by 3% PQ on Semantic-KITTI, and open-world baselines by 20% UQ; under review at **IJCV**

Pittsburgh, PA

Nov 2021 – Dec 2022

### Harvard University

Research Intern (mentor: Dr. Hanspeter Pfister)

- Implemented an inter-frame attention mechanism in PyTorch to propagate missing object instances in a video track; achieved state-of-the-art performance on YouTube-VIS benchmark with 5% improvement on Mask R-CNN pipeline
- Developed a transformer-based tracking and segmentation method by using transformer encoded embeddings from previous frames as tracking cues for the current frame; obtained 33.1% mAP on YouTube-VIS benchmark

Boston, MA

May 2020 – July 2021

## Publications

**Anirudh S Chakravarthy**, Shuai Zheng, Xin Huang, Sachi Hemachandra, Yuning Chai, and Zhao Chen. PROFIT: A PROximal FIne Tuning Optimizer. In *Review at European Conference on Computer Vision (ECCV)*, 2024.

**Anirudh S Chakravarthy**, Meghana Ganesina, Peiyun Hu, Laura Leal-Taixe, Shu Kong, Deva Ramanan, and Aljosa Osep. Lidar Panoptic Segmentation in an Open World. In *Review at International Journal on Computer Vision (IJCV)*, 2024.

**Anirudh S Chakravarthy**, Won-Dong Jang, Zudi Lin, Donglai Wei, Song Bai, and Hanspeter Pfister. Object Propagation via Inter-Frame Attentions for Temporally Stable Video Instance Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2021.

**Anirudh S Chakravarthy\***, Praveen Ravirathinam\*, and Roshan Roy\*. MRSCAtt: A Spatio-Channel Attention-Guided Network for Mars Rover Image Classification. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2021.

Donglai Wei et. al. YouMVOS: An Actor-centric Multi-shot Video Object Segmentation Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

## Projects

### Adversarial Robustness of GANs

Carnegie Mellon University [[project page](#)] [[code](#)]

- Proposed a method to examine latent space robustness in GANs by identifying poorly sampled regions in latent space
- Performed adversarial attacks with in-distribution constraints to arrive at latent vectors which result in noisy reconstructions

Mar 2022 – May 2022

Pittsburgh, PA

### Improving Multi-Person Reconstruction using Temporal Cues

Carnegie Mellon University [[code](#)]

- Incorporated optical flow using early fusion to augment temporal cues into 3D pose estimation networks; improved 3D keypoint accuracy by 0.28% and reduced inter-person interpenetration by 5% on the MuPoTS-3D benchmark

Mar 2022 – May 2022

Pittsburgh, PA