Rain, Rain, Go Away:
Computer Vision in Adverse Weather

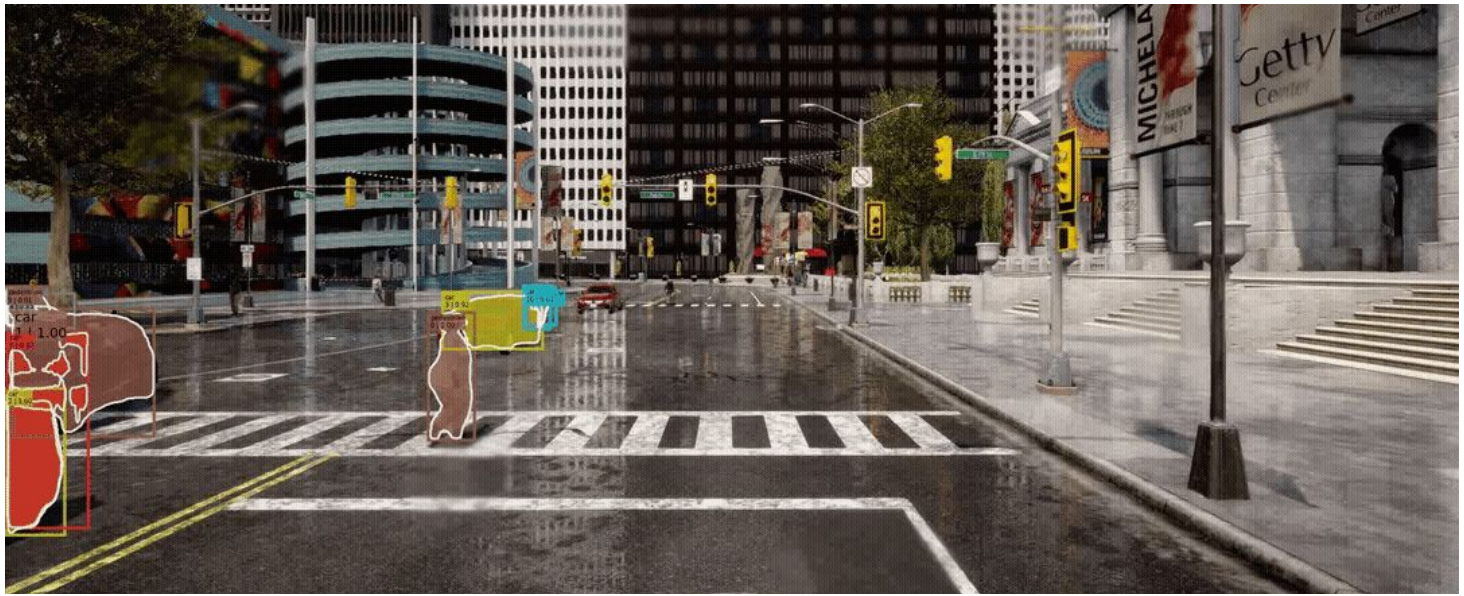

Ben Zhao, David Toledo, Ethan Masadde, Josh Meier, Julian Tanguma, Nathaniel Li

When a computer is in the driver's seat, it's just as confused as you

Presentation Roadmap

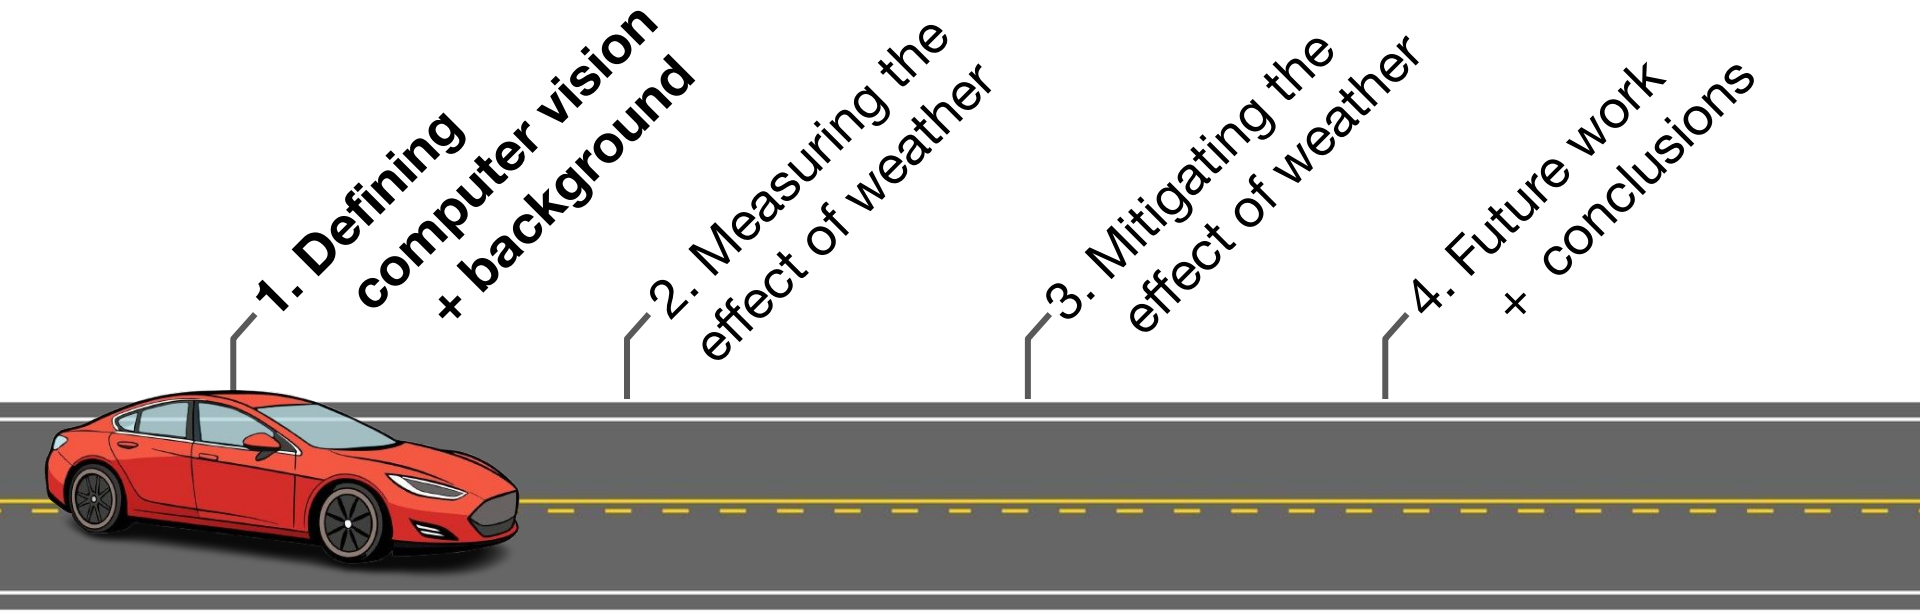

How do autonomous vehicles work

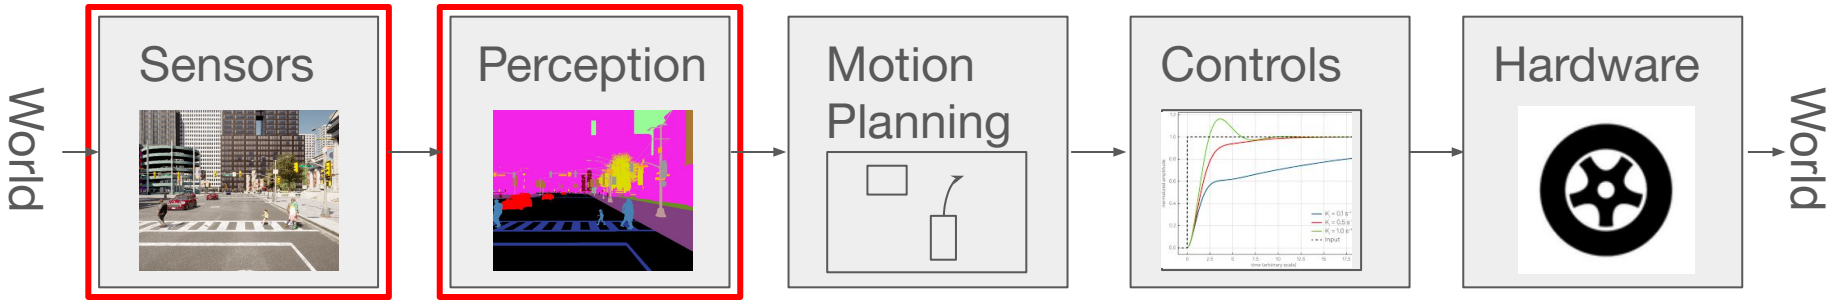

The Eyes of the Vehicle

LiDAR

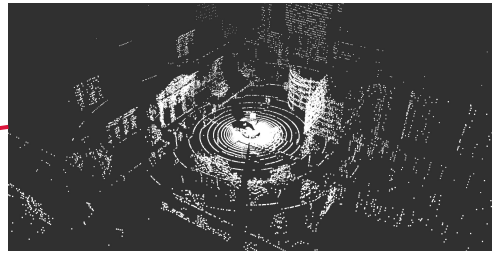

Radar

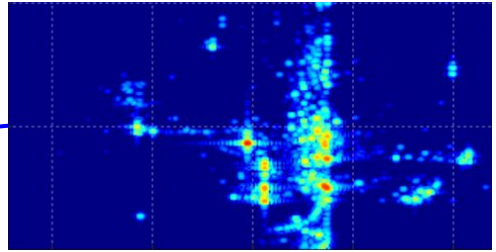

Camera

Making Sense of Pixels and Points

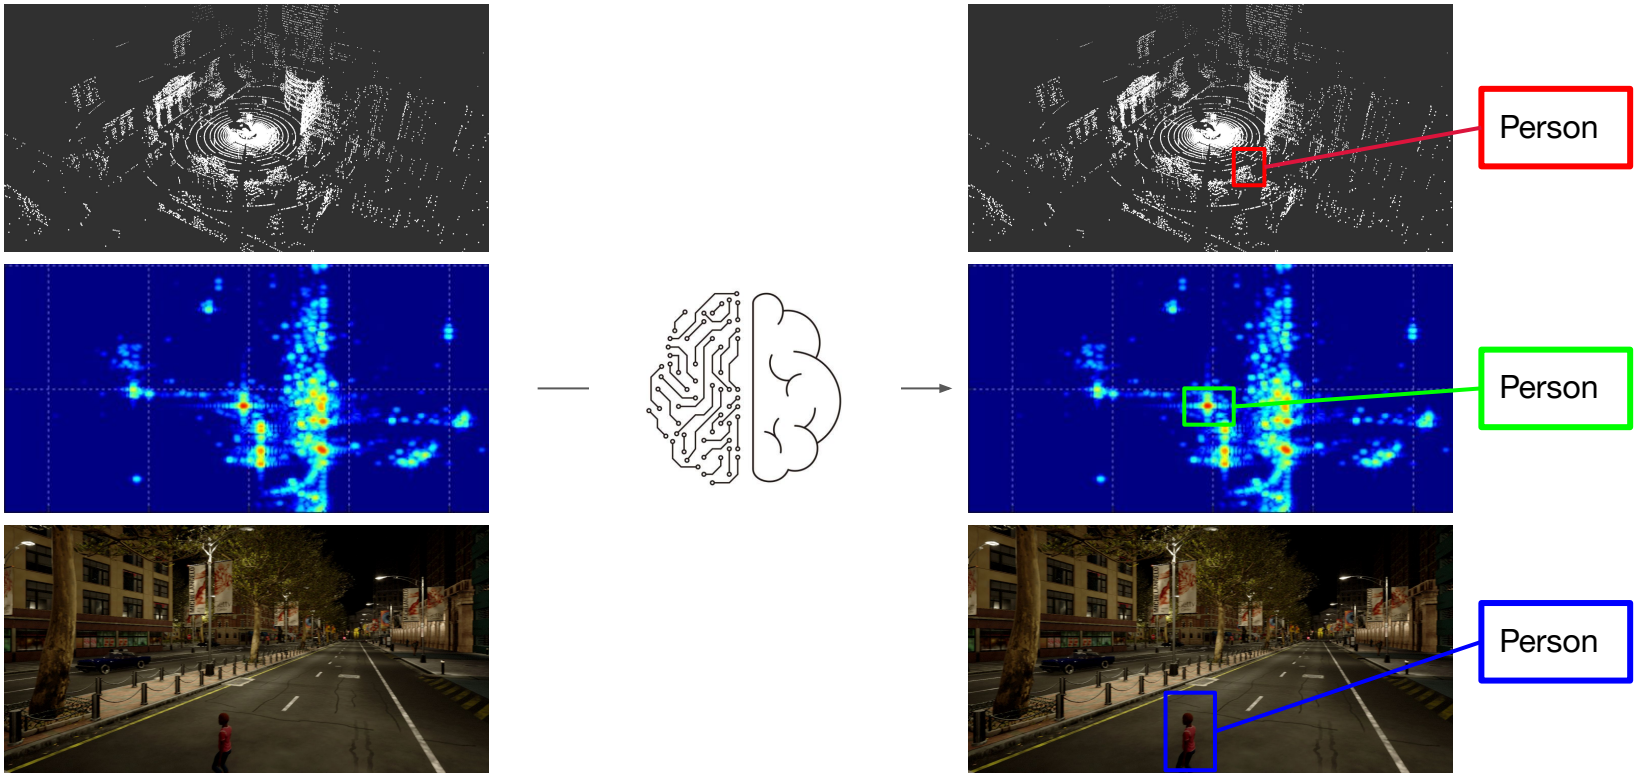

Making Sense of Pixels and Points

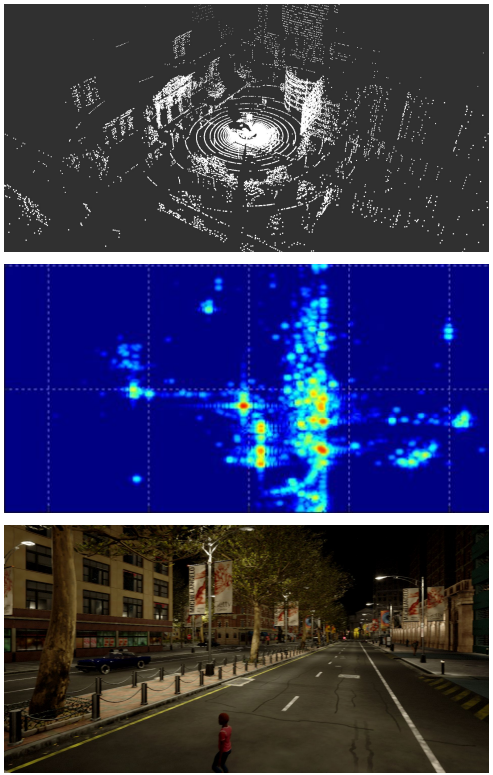

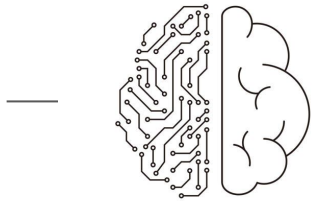

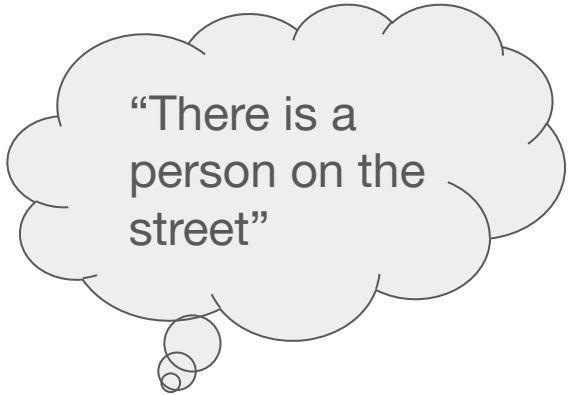

Computer vision (CV)

Classification

CAT

Object Detection

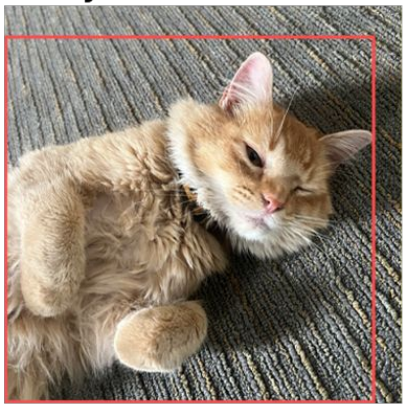

CAT

Semantic Segmentation

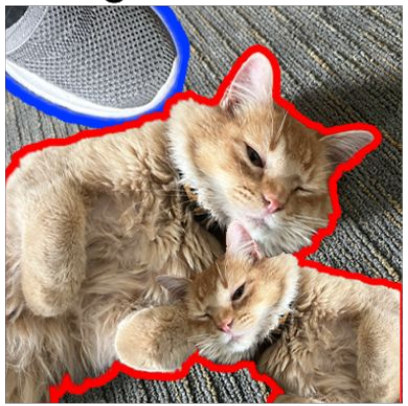

SHOE, CAT

Instance Segmentation

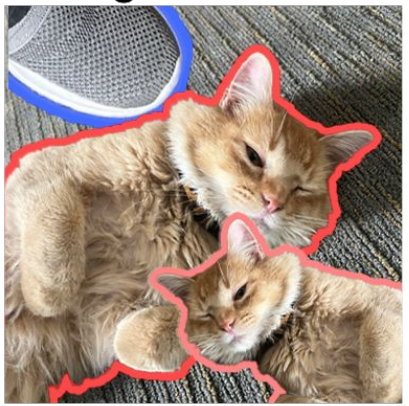

SHOE, CAT1, CAT2

single object

multiple objects

Semantic Segmentation

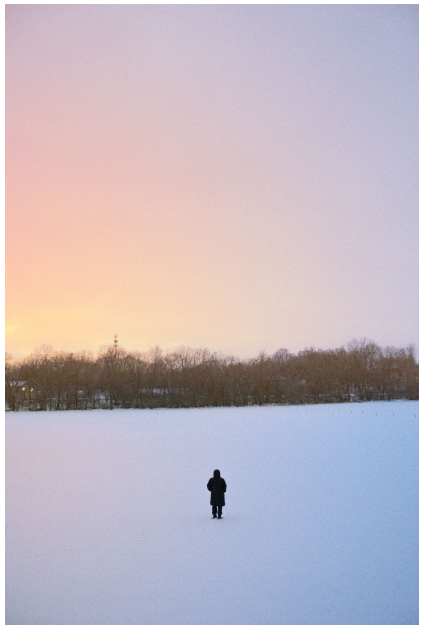

Input

Segmented
1: Person
2: Snow
3: Sky
4: Trees

3	3	3	3	3	3	3
3	3	3	3	3	3	3
3	3	3	3	3	3	3
3	3	3	3	3	3	3
3	3	3	3	3	3	3
3	3	3	3	3	3	3
4	4	4	4	4	4	4
2	2	2	2	2	2	2
2	2	2	1	2	2	2
2	2	2	1	2	2	2
2	2	2	2	2	2	2

Semantic Labels

Instance Segmentation

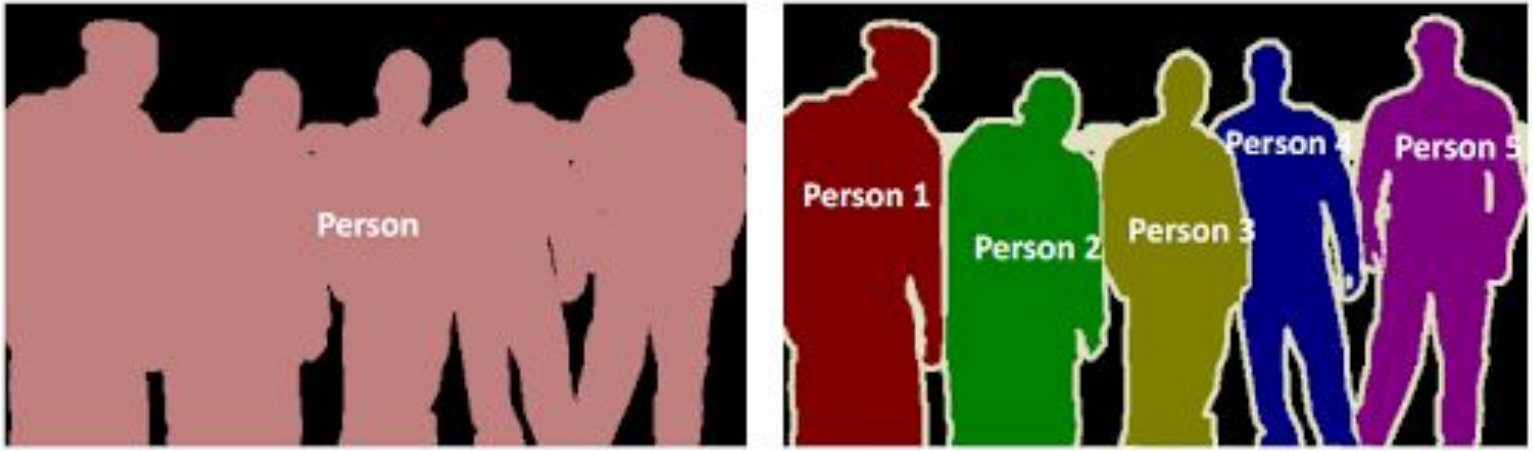

Semantic Segmentation

Instance Segmentation

Video Instance Segmentation

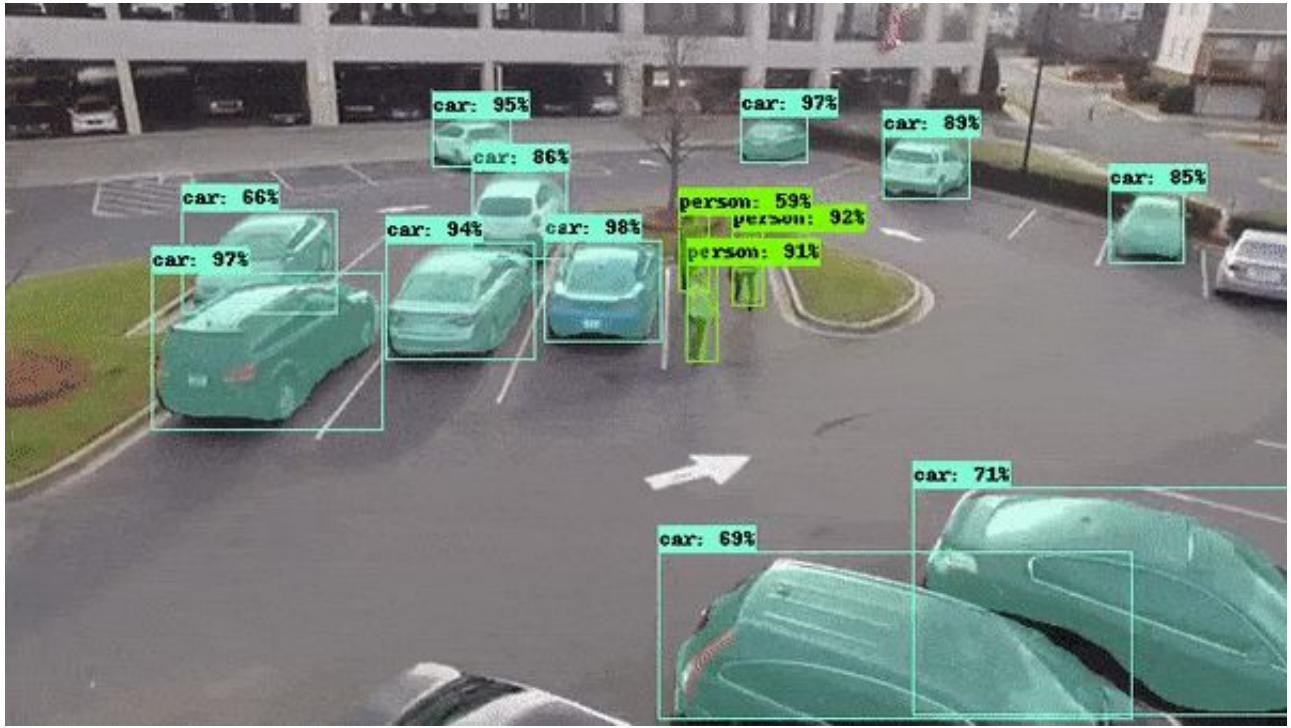

“How much does weather affect the performance of computer vision? And what can we do about it?”

Team Organization

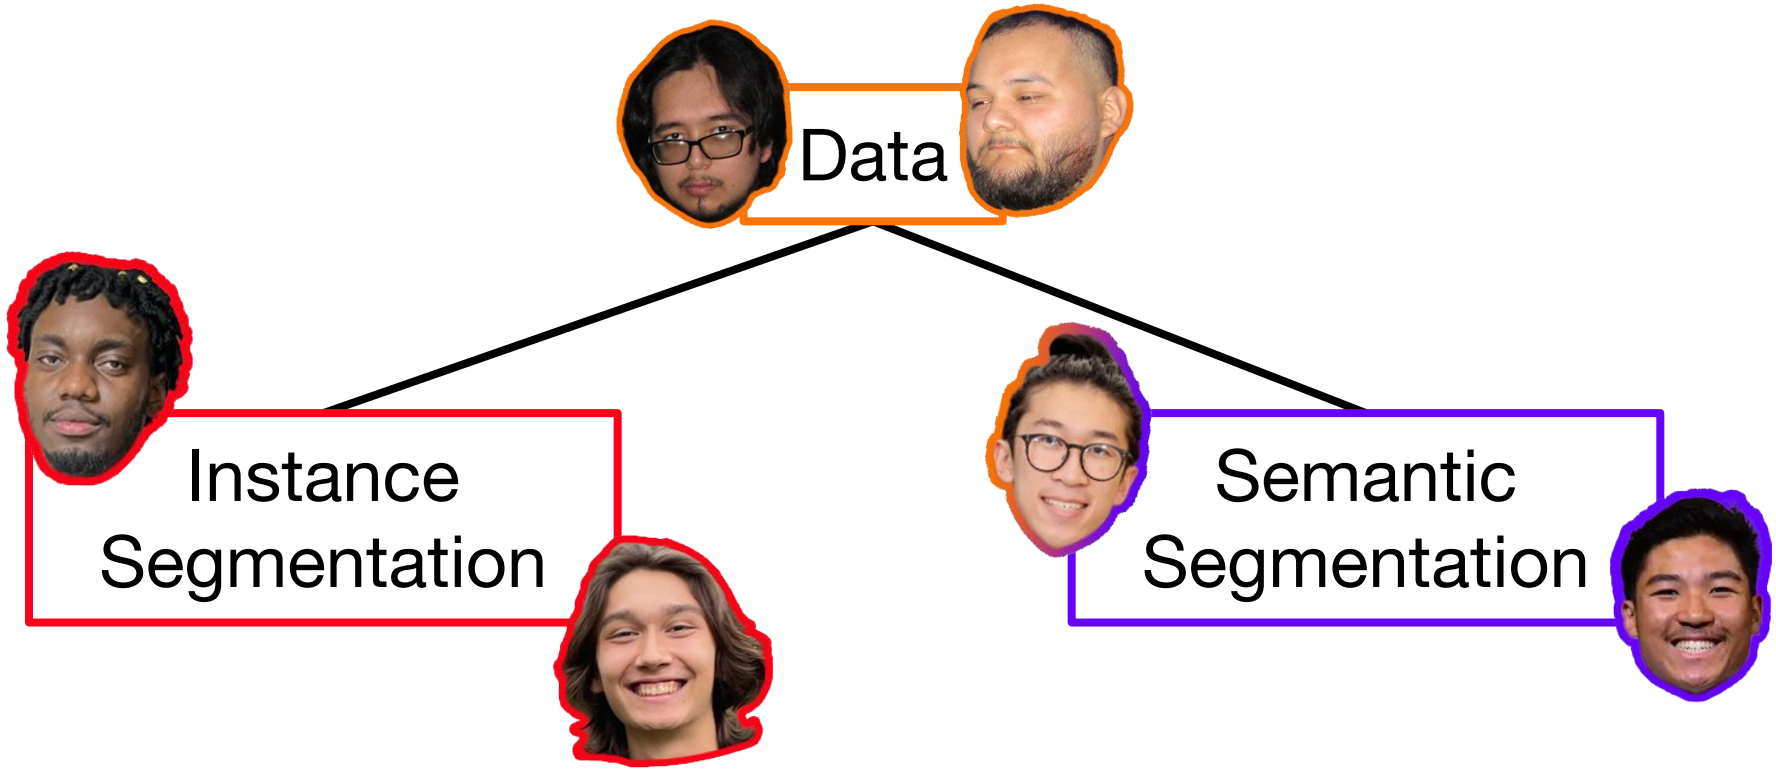

1. Defining computer vision + background

2. Measuring the effect of weather

3. Mitigating the effect of weather

4. Future work + conclusions

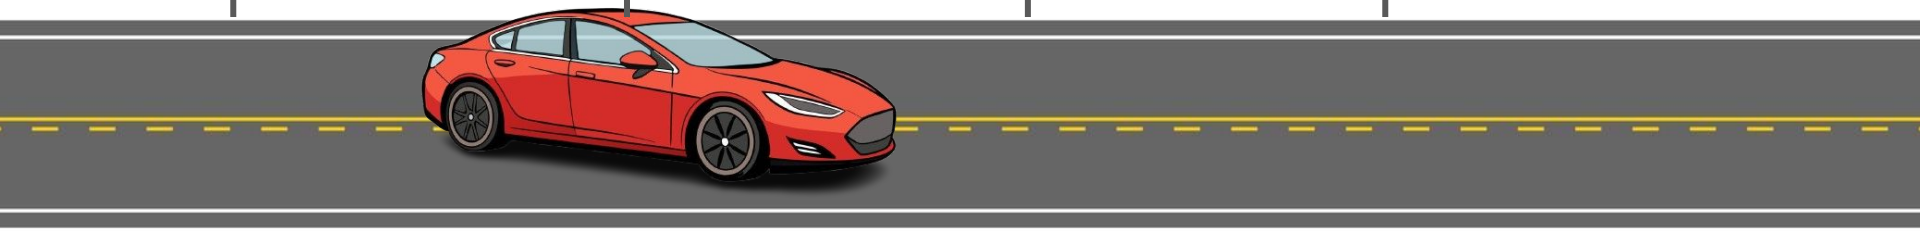

Part 2 Overview: Setup

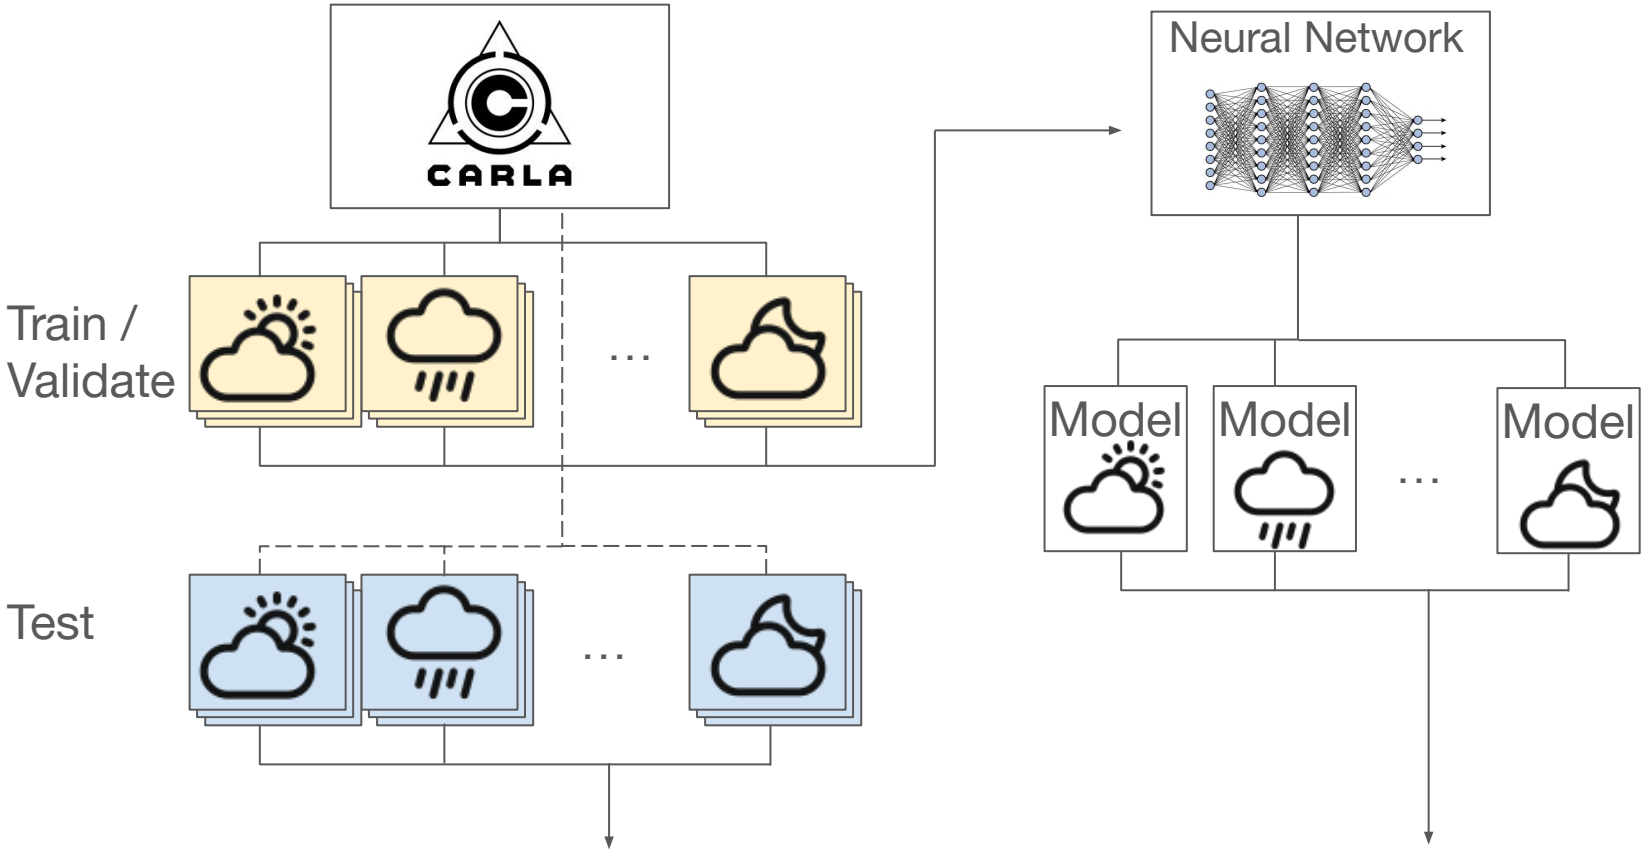

Part 2 Overview: Evaluation

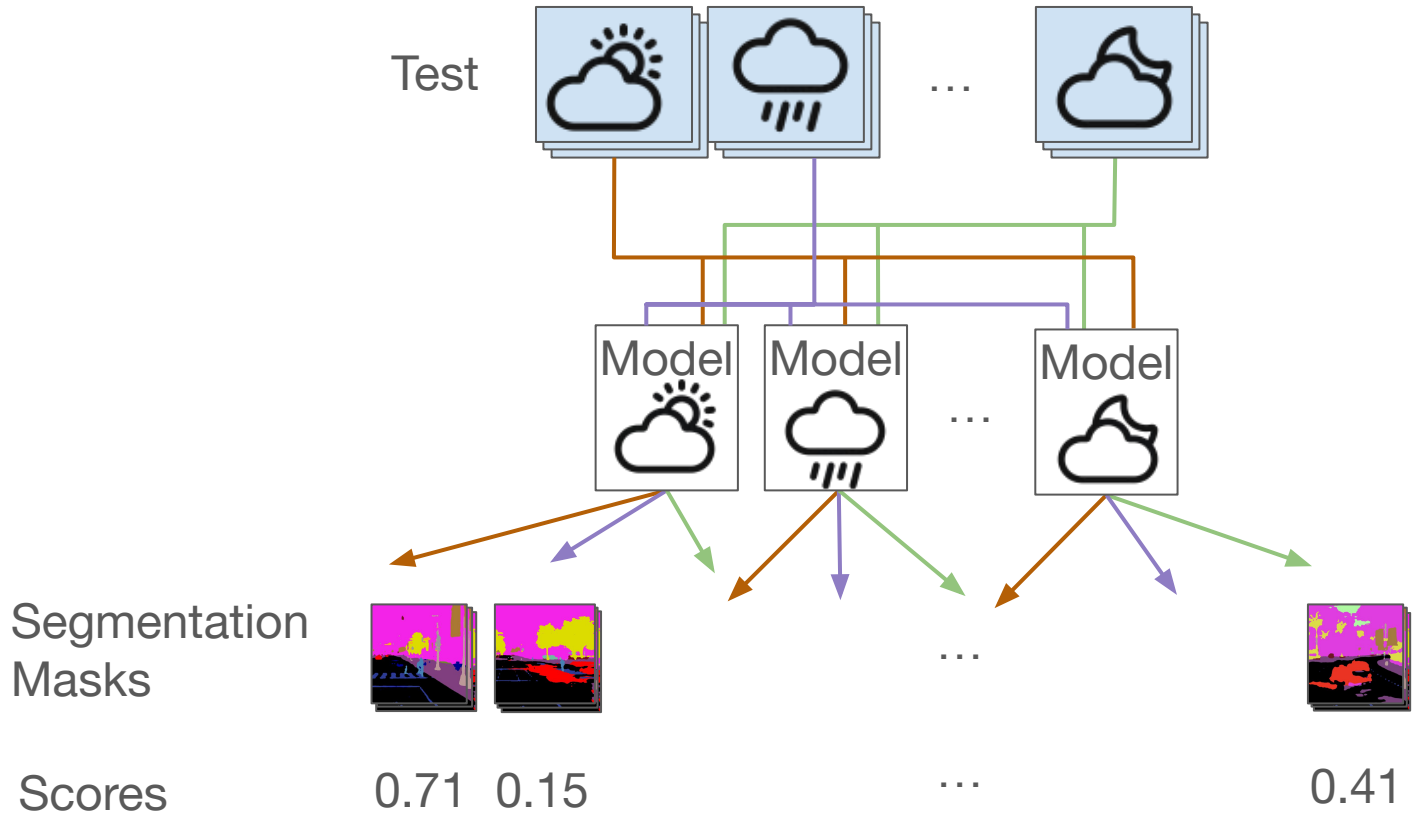

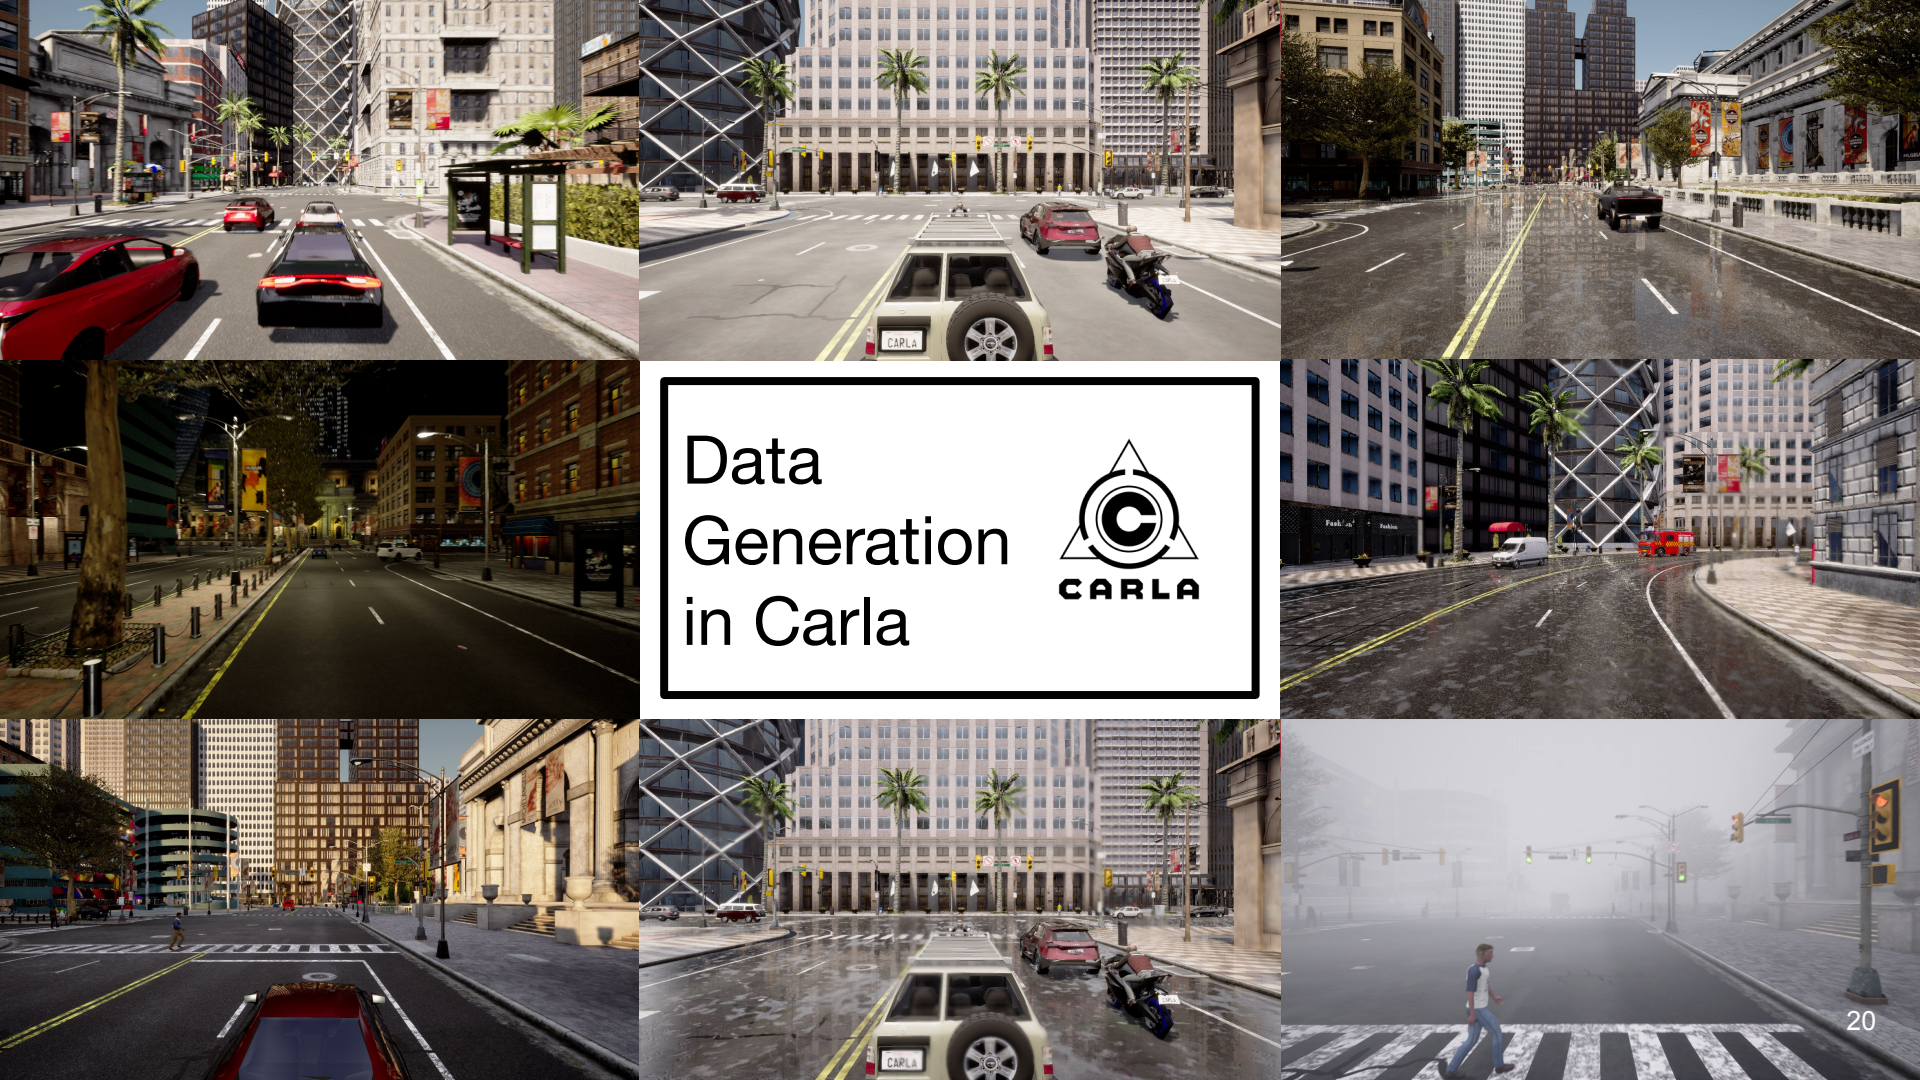
Data
Generation
in Carla
CARLA

What is Ground Truth?


ONE WAY
MUSEUM
MUSEUM


MUSEUM
MUSEUM
ONE WAY


Getty
Center

Convolutional Neural Network: Encoder

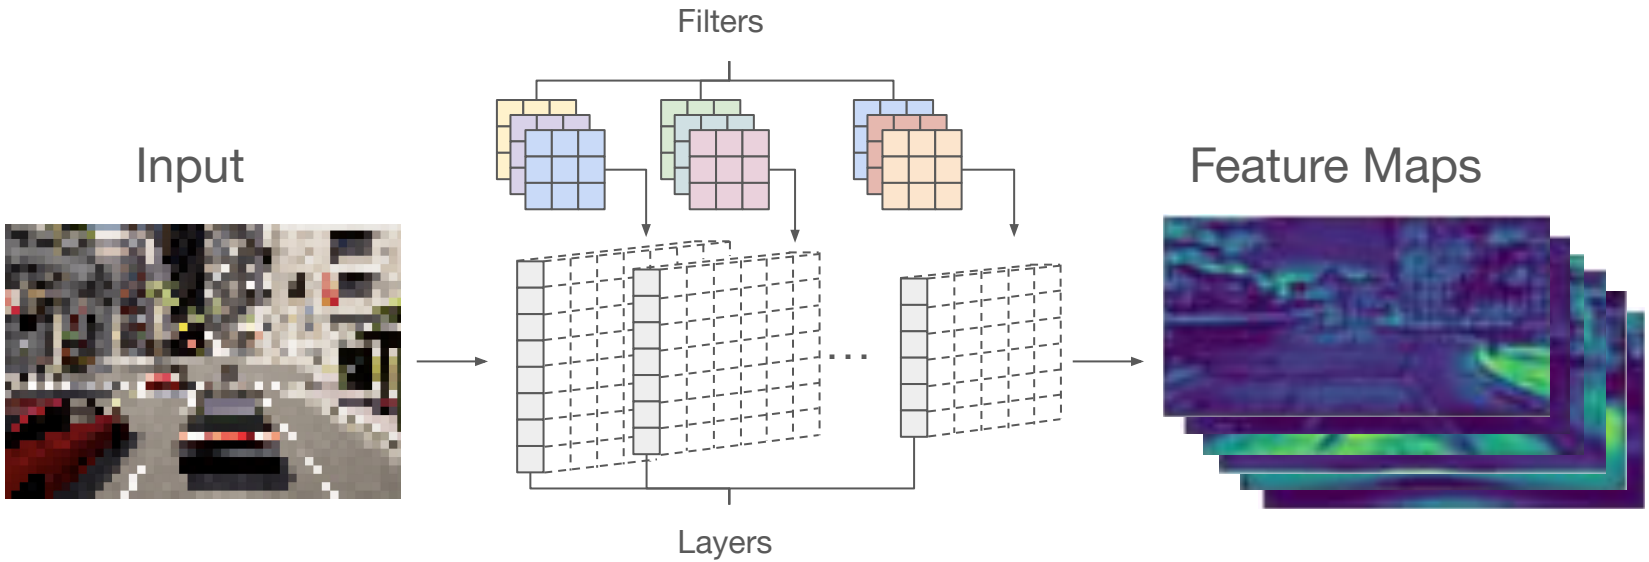

Convolutional Neural Network: Kernel / Filter

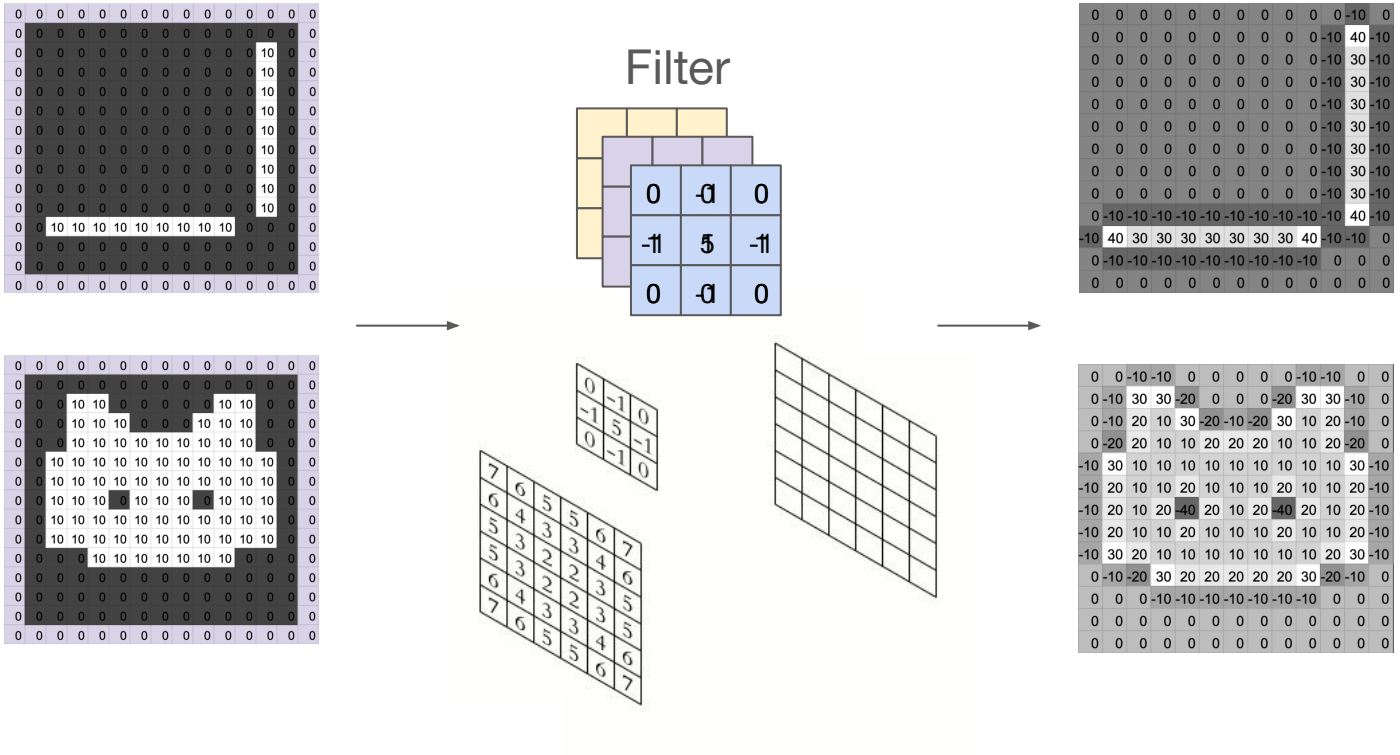

Convolutional Neural Network: Feature Maps

Convolutional Neural Network: Feature Maps

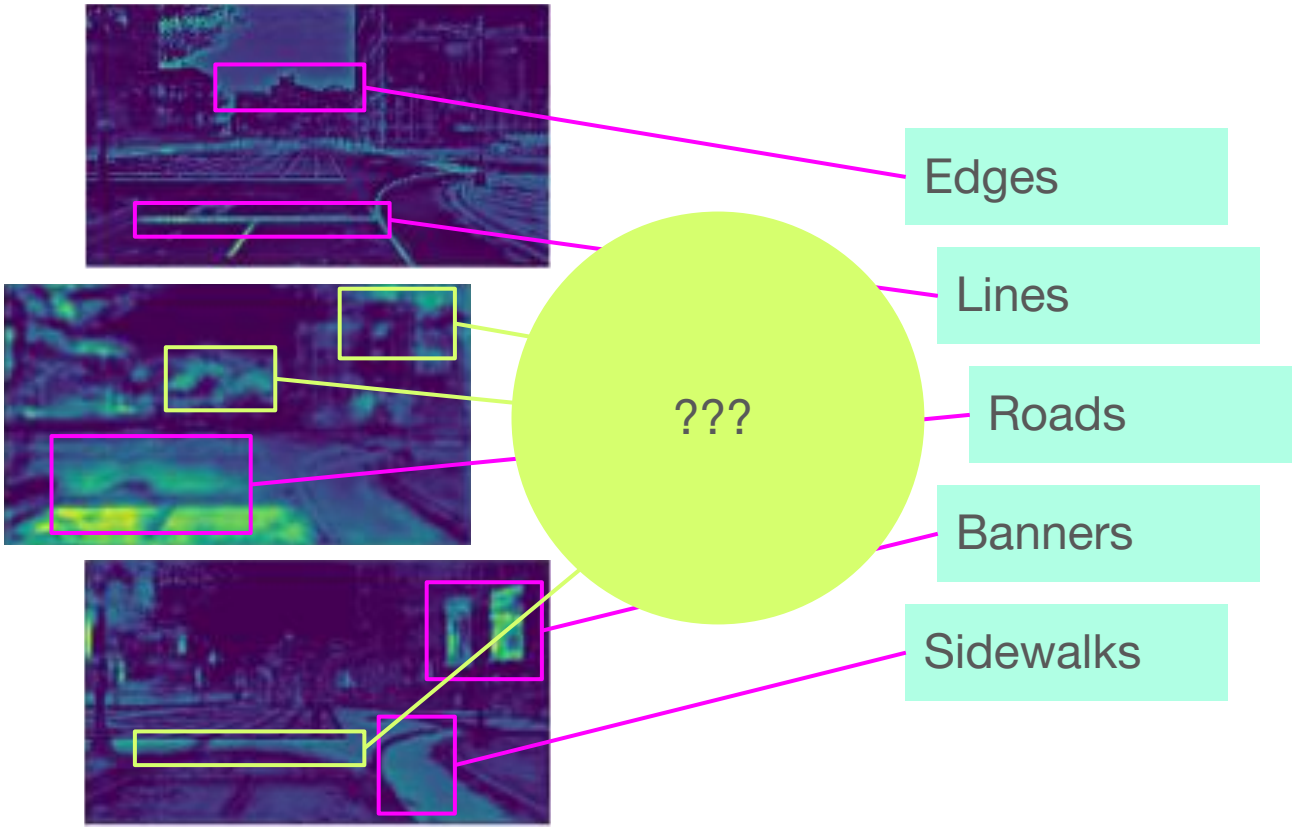

Convolutional Neural Network: Feature Maps

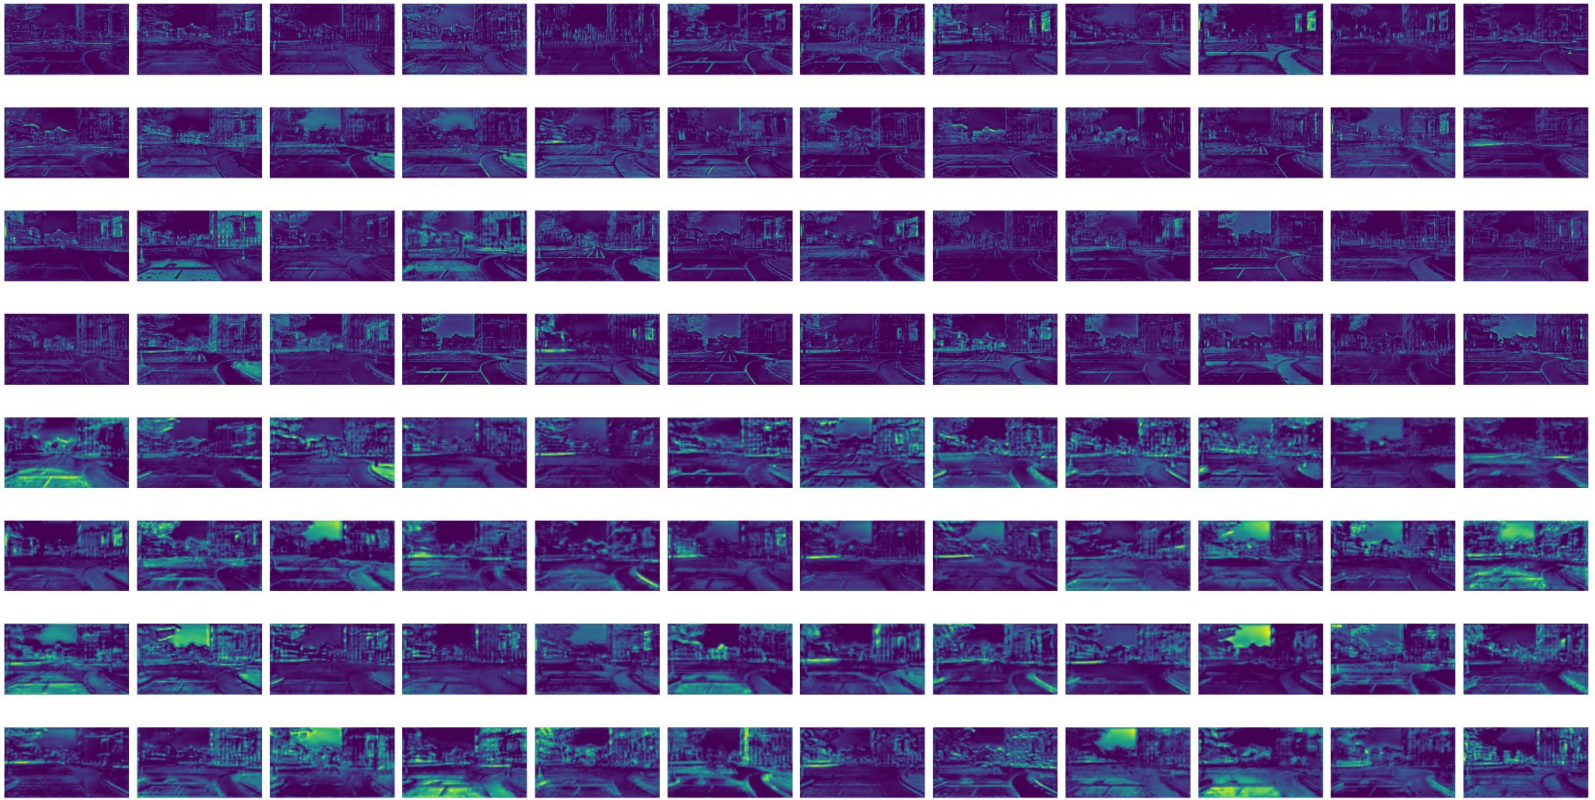

Convolutional Neural Network: Decoder

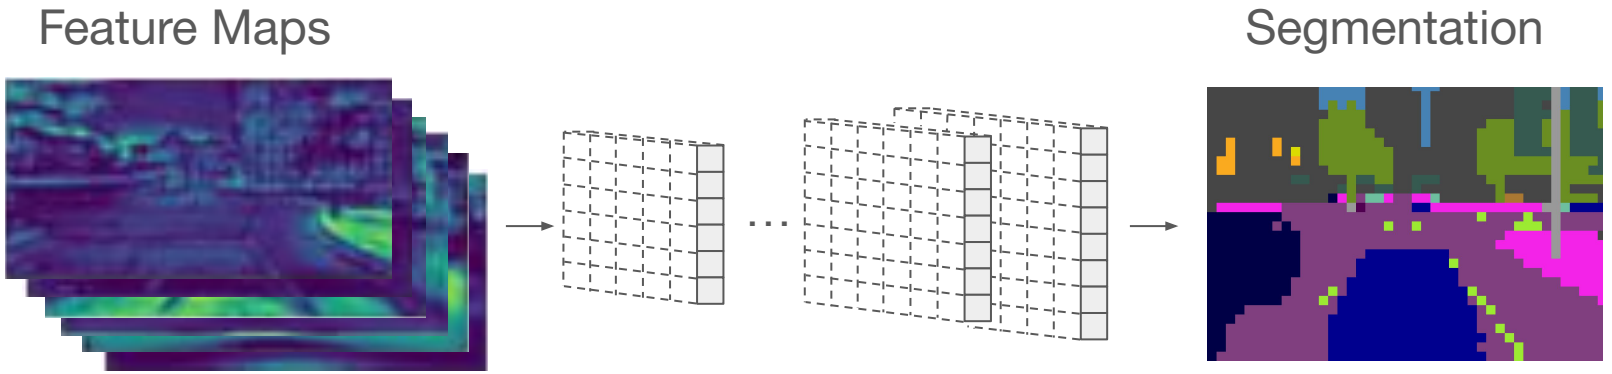


MICHELANGELO

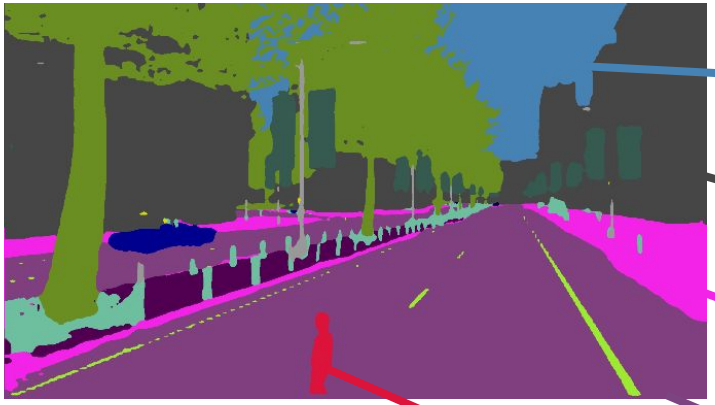

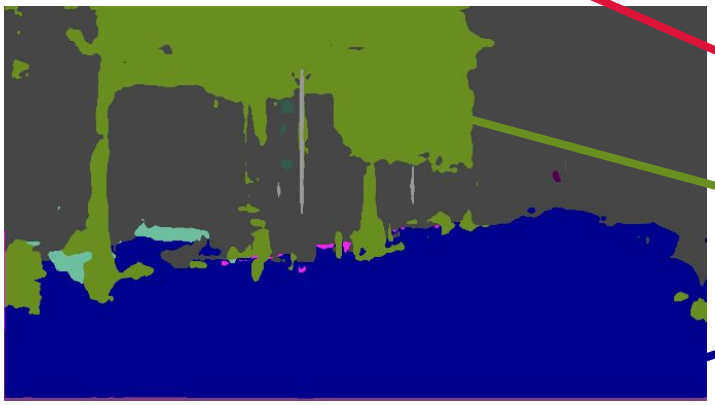

Sky

Building

Sidewalks

Road

Pedestrian

Vegetation

Car

Metrics: Mean Intersection over Union (MIOU)

$$IoU = \frac{\text{Area of Intersection}}{\text{Area of Union}}$$

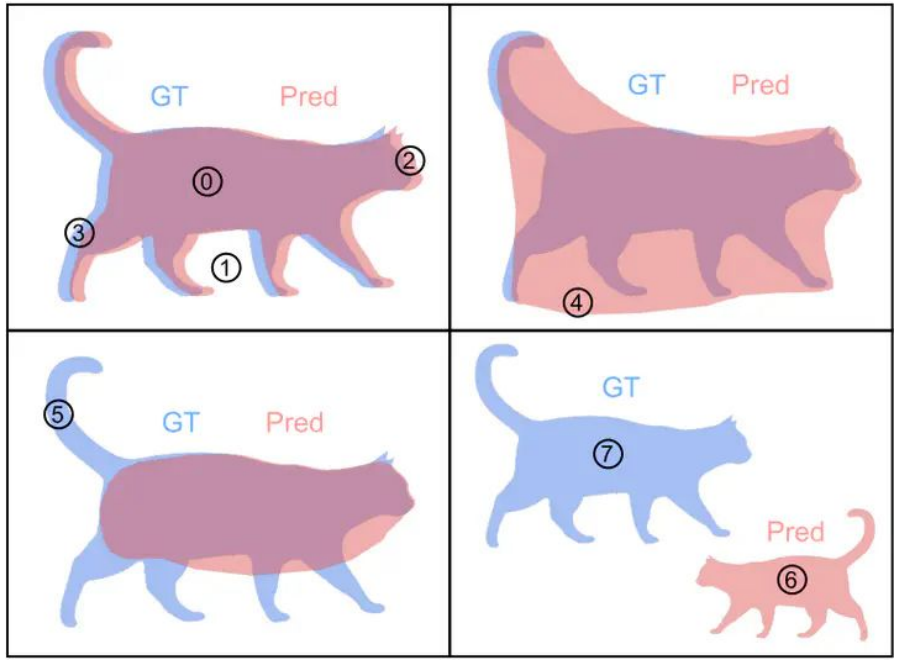

Semantic Segmentation Results

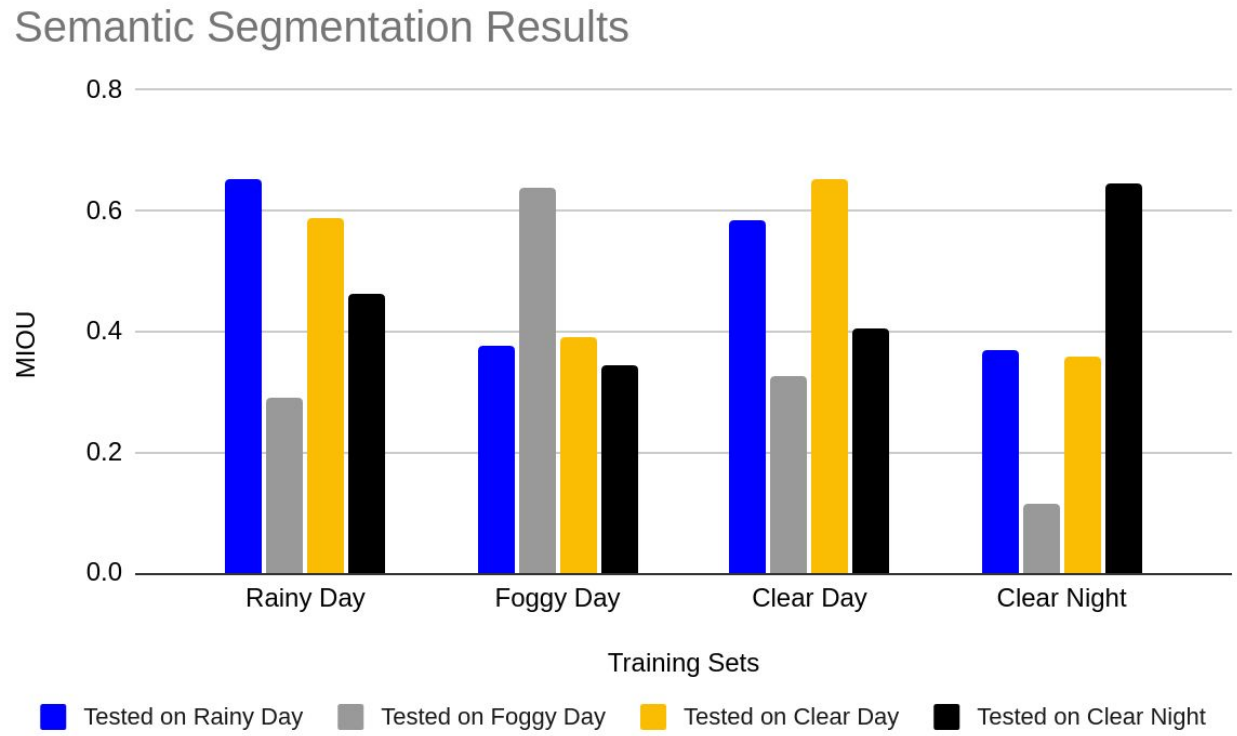

Semantic Segmentation Results

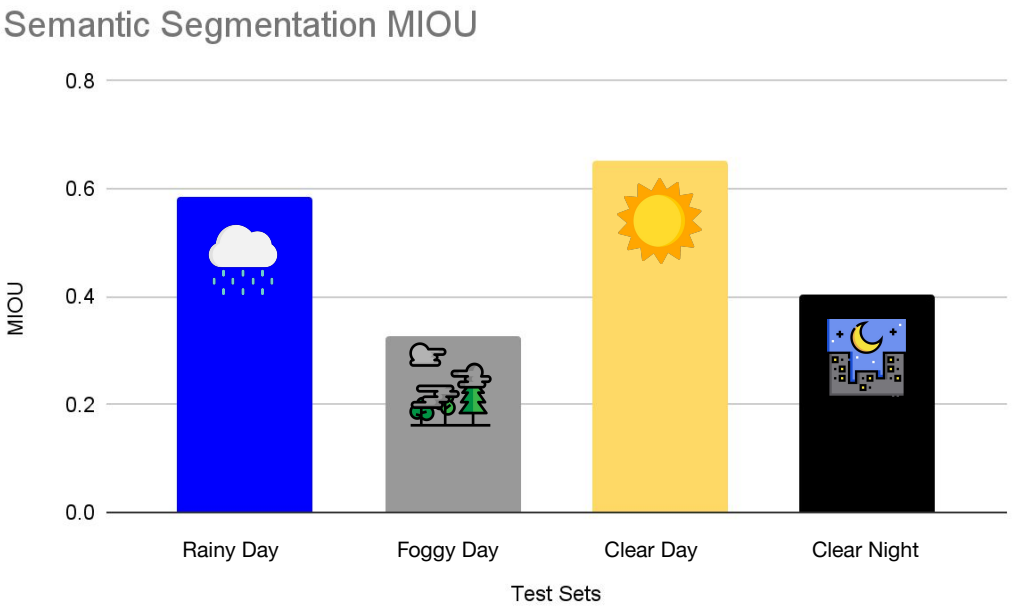

Training Set

Semantic Segmentation Results

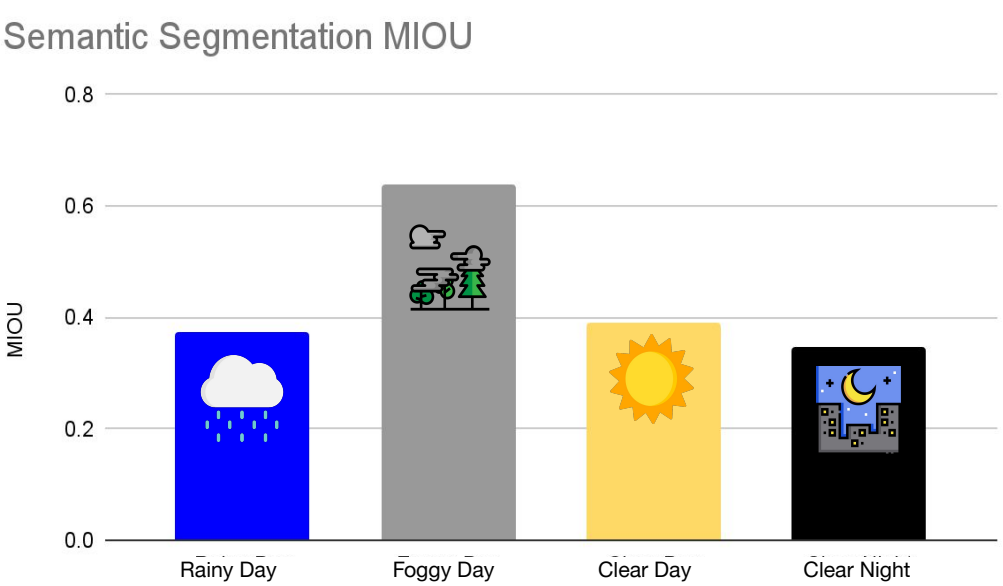

Interpreting MIOU

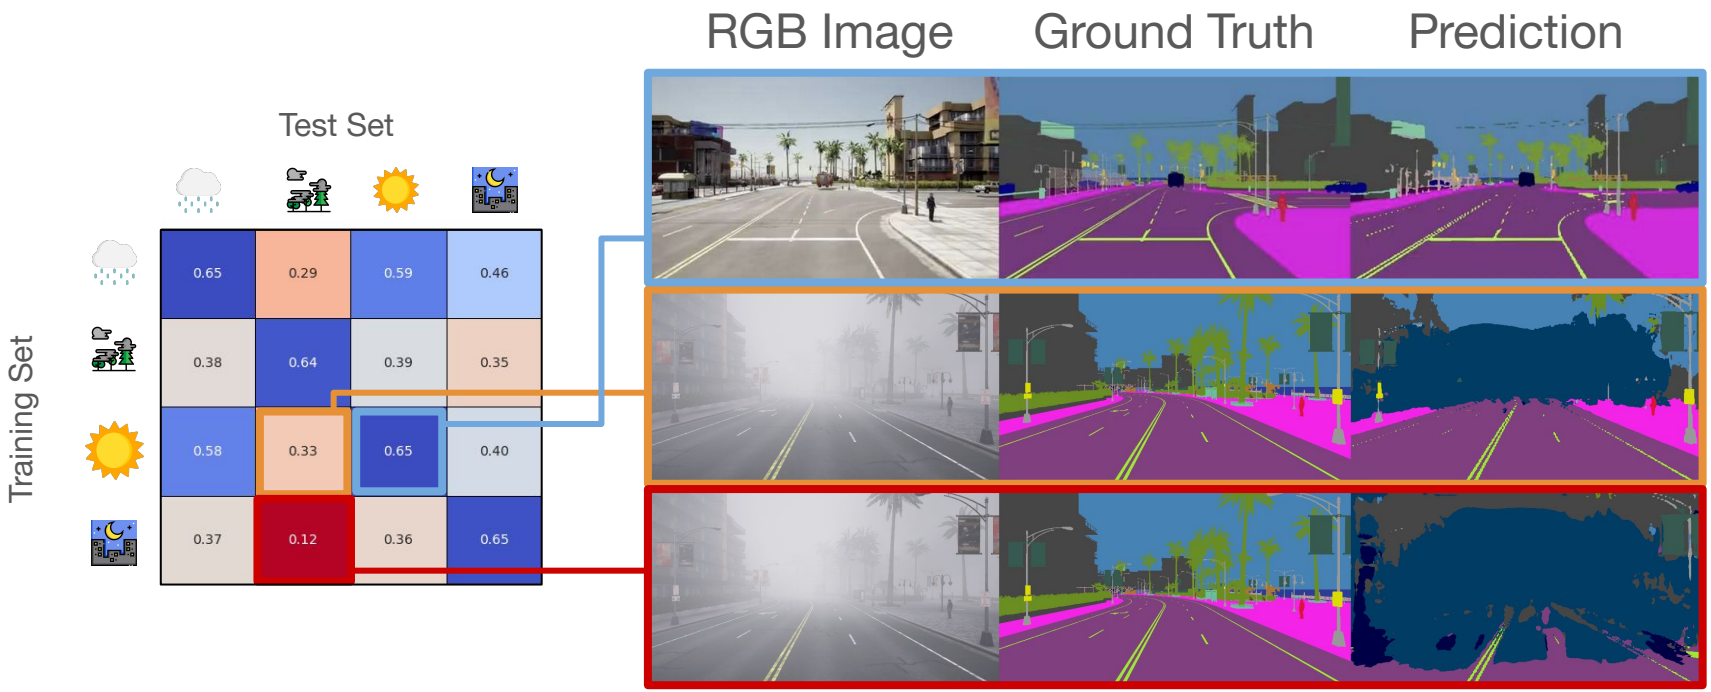

Segmentation Results: Closer Look

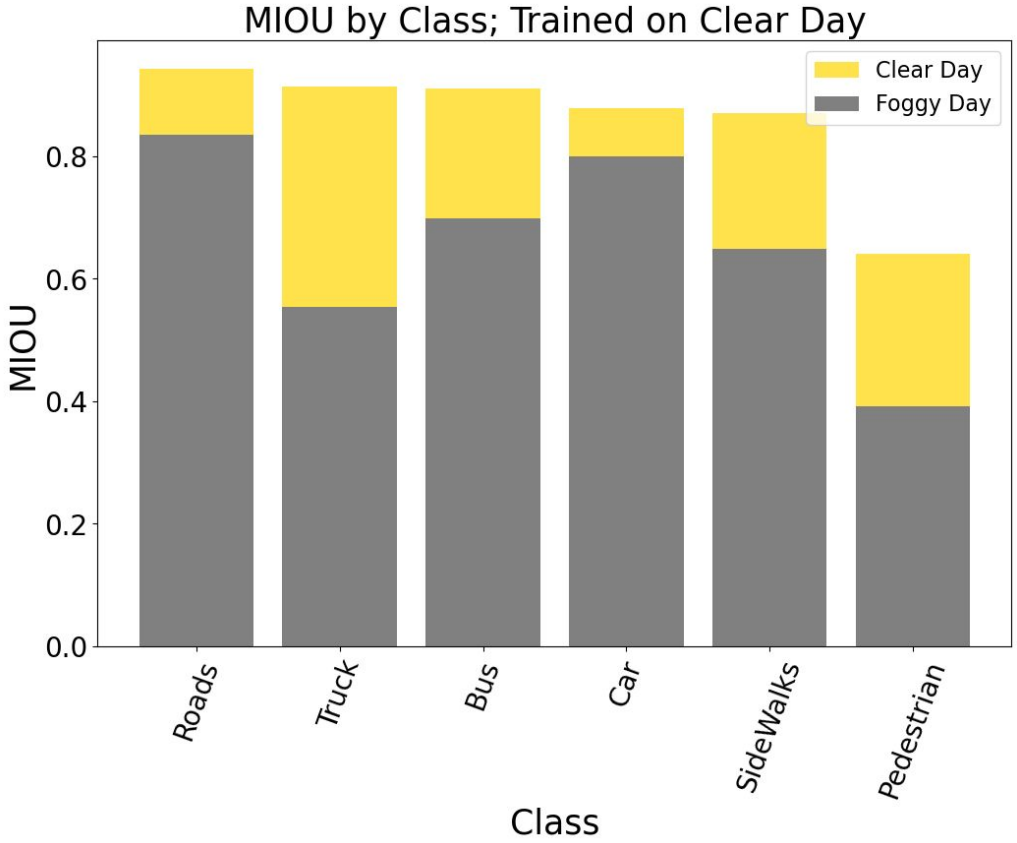

Video Instance Segmentation Reintroduced

Metrics: Mean Average Precision (MAP)

If IOU threshold = 0.5…

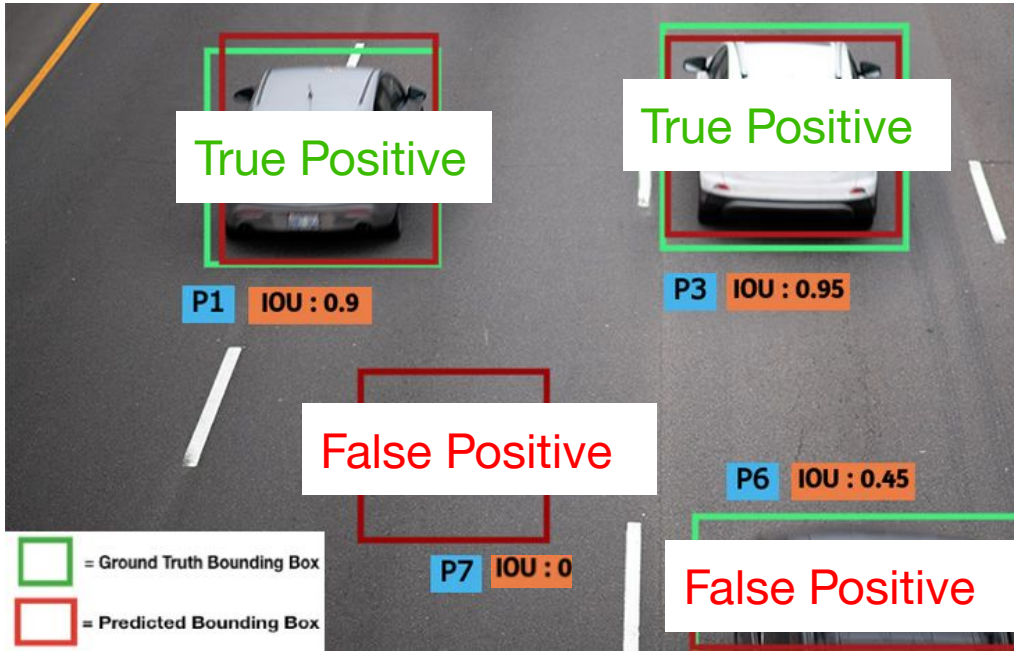

If IOU of prediction > threshold:

True Positive: +1 +1

else:

False Positive: +1 +1

$$\text{Precision} = \frac{\text{True Positive}}{\text{True Positive} + \text{False Positive}}$$

Instance Segmentation Results

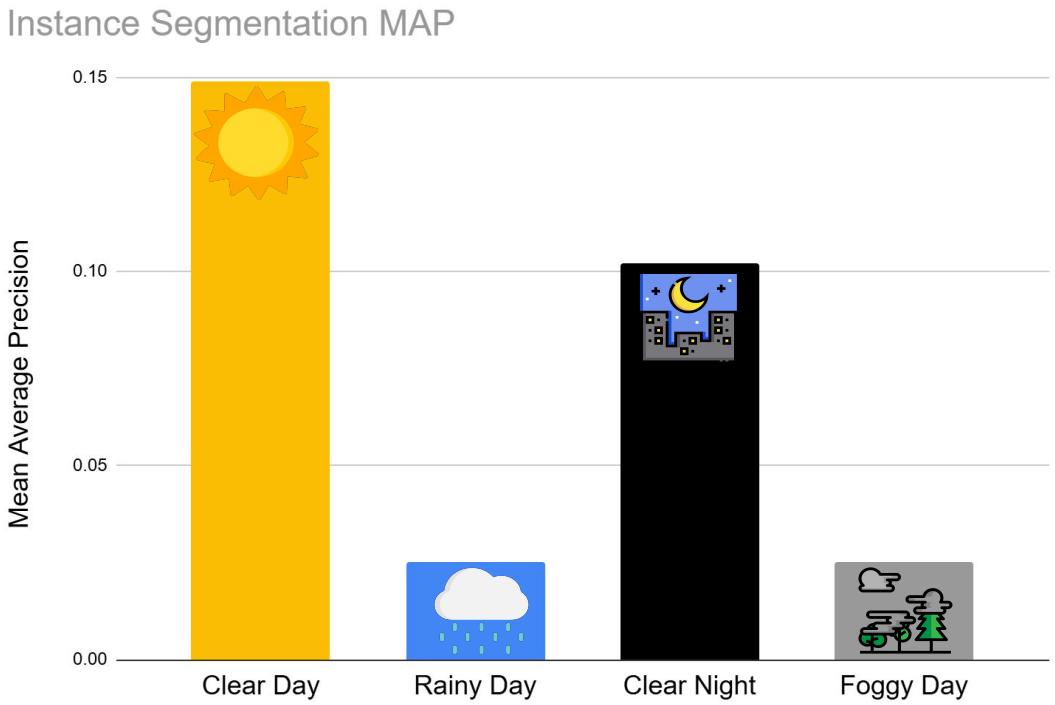

Training Set

Instance Segmentation Results

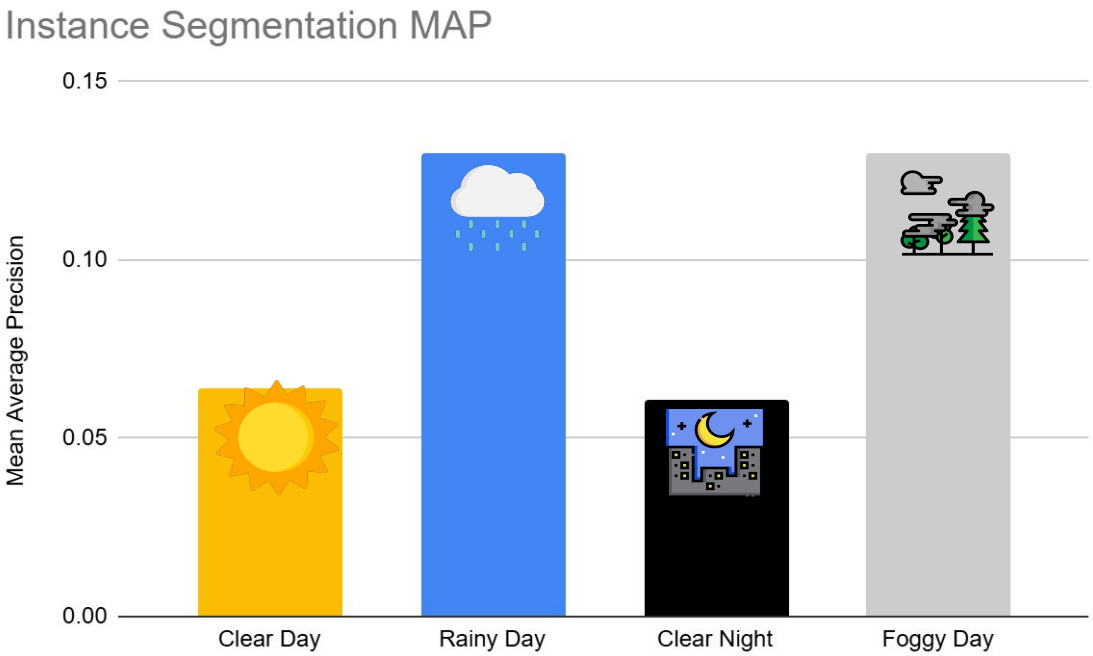

Training Set

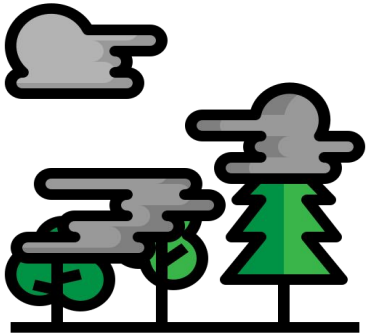

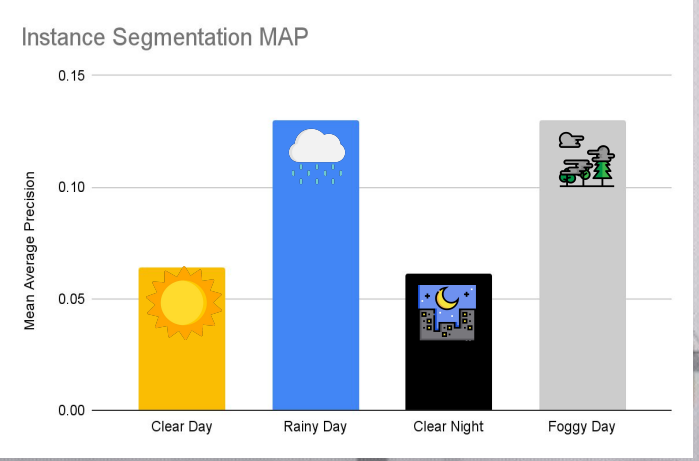
Instance Segmentation MAP
Mean Average Precision
0.15
0.10
0.05
0.00
Clear Day
Rainy Day
Clear Night
Foggy Day

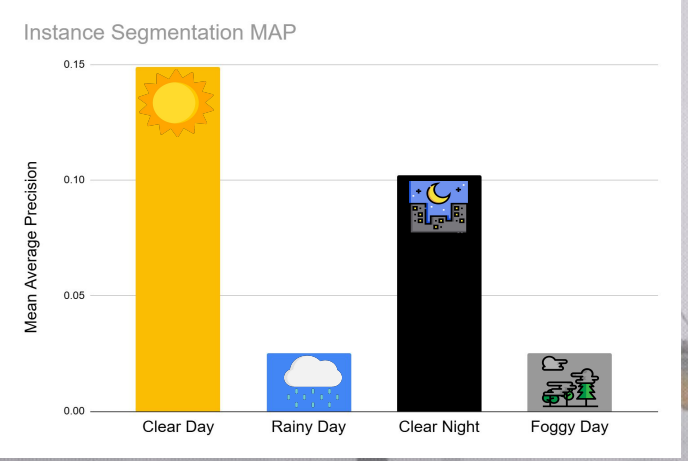
Instance Segmentation MAP
Mean Average Precision
0.15
0.10
0.05
0.00
Clear Day
Rainy Day
Clear Night
Foggy Day

1. Defining computer vision + background

2. Measuring the effect of weather

3. Mitigating the effect of weather

4. Future work + conclusions

Weather Mitigation Strategies

Domain Adaptation

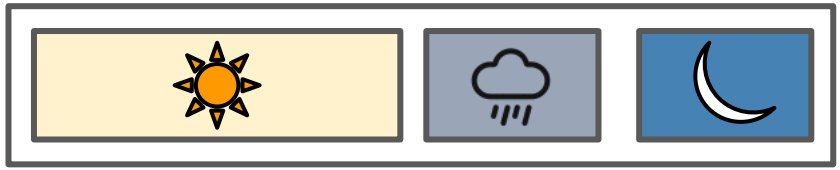

Deweathering

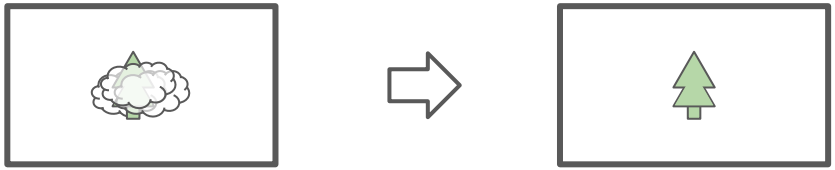

Sensor Fusion

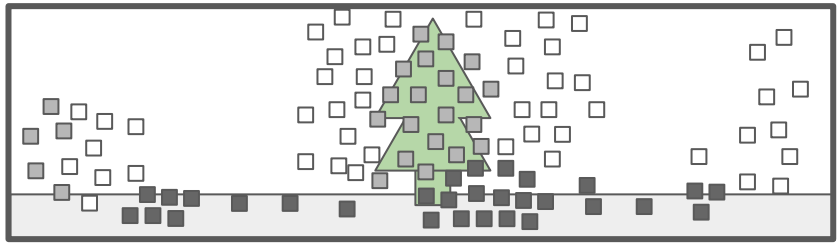

Domain Adaptation

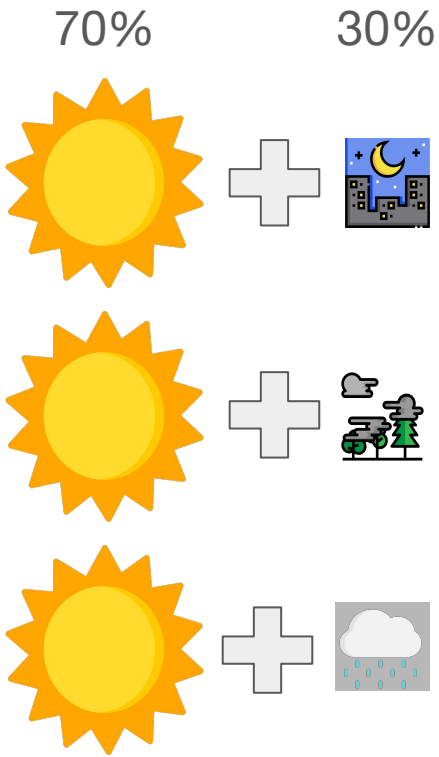

Domain Adaptation Results

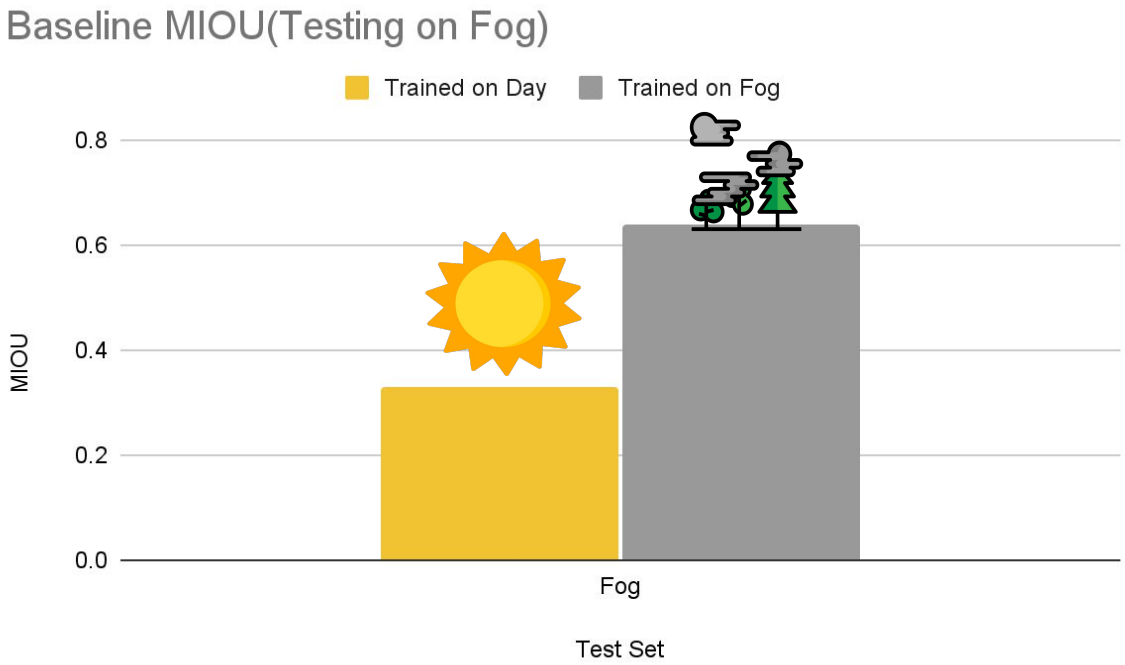

Domain Adaptation Results

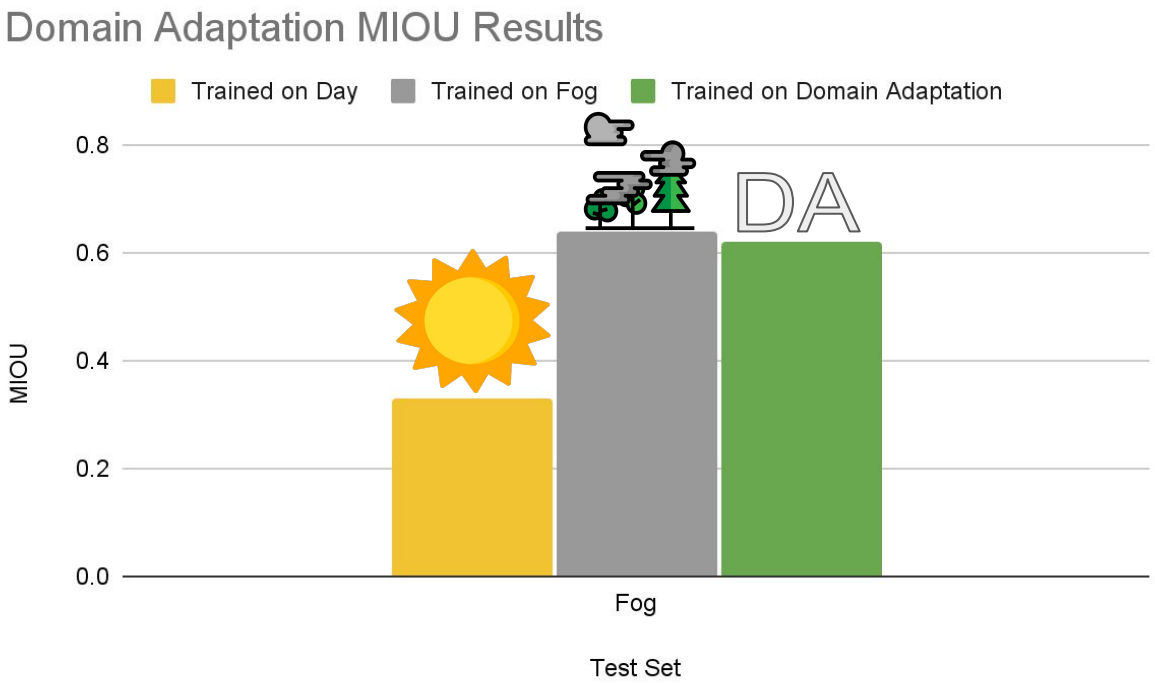

Domain Adaptation Results

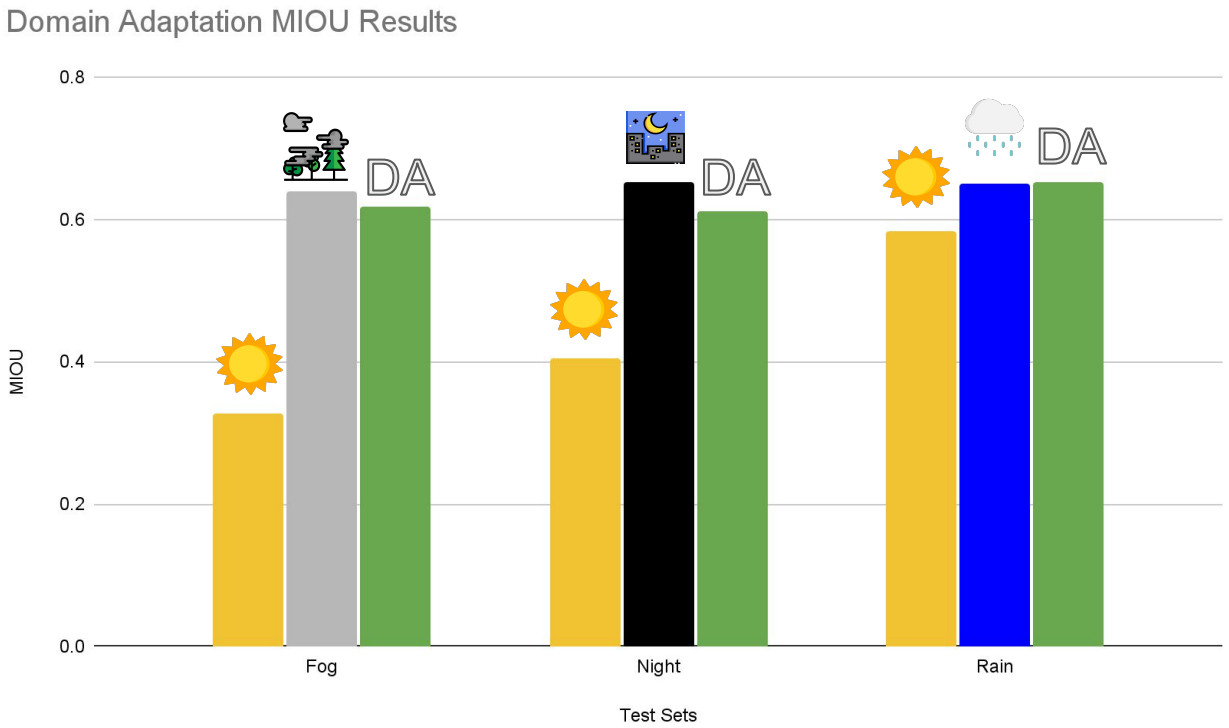

Domain Adaptation Results

Foggy Day(without Domain Adaptation): MIOU .328

Foggy Day(with Domain Adaptation): MIOU .619

Domain Adaptation Results, Instance Segmentation

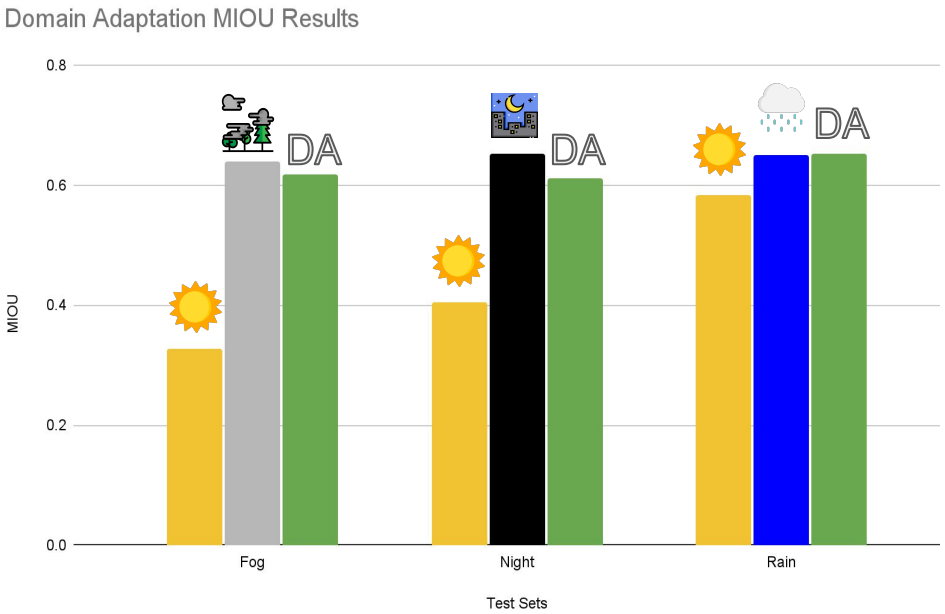

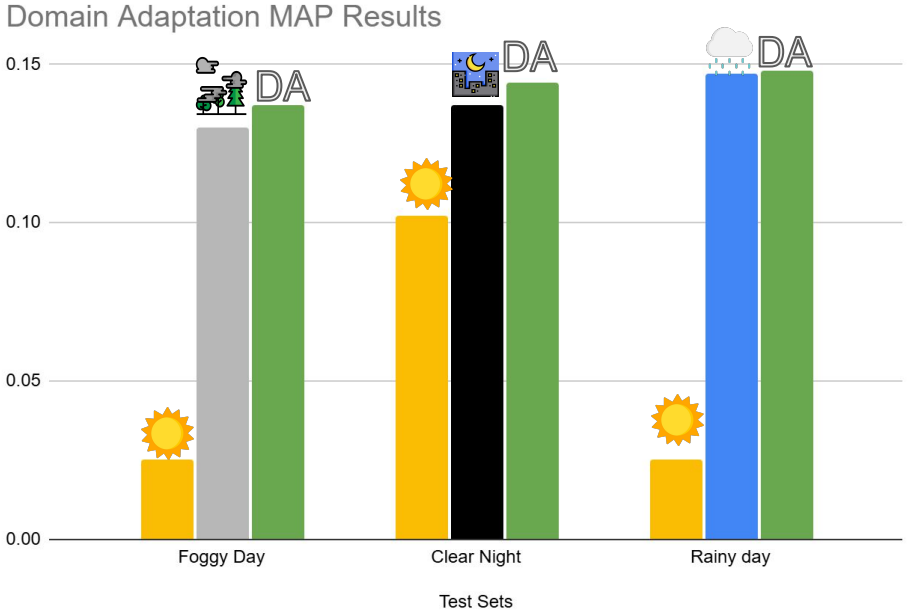

Domain Adaptation Results

Trained on only Clear Day

Domain Adaptation: Trained on clear/foggy split

Deweathering: Image Processing

Look Up Table

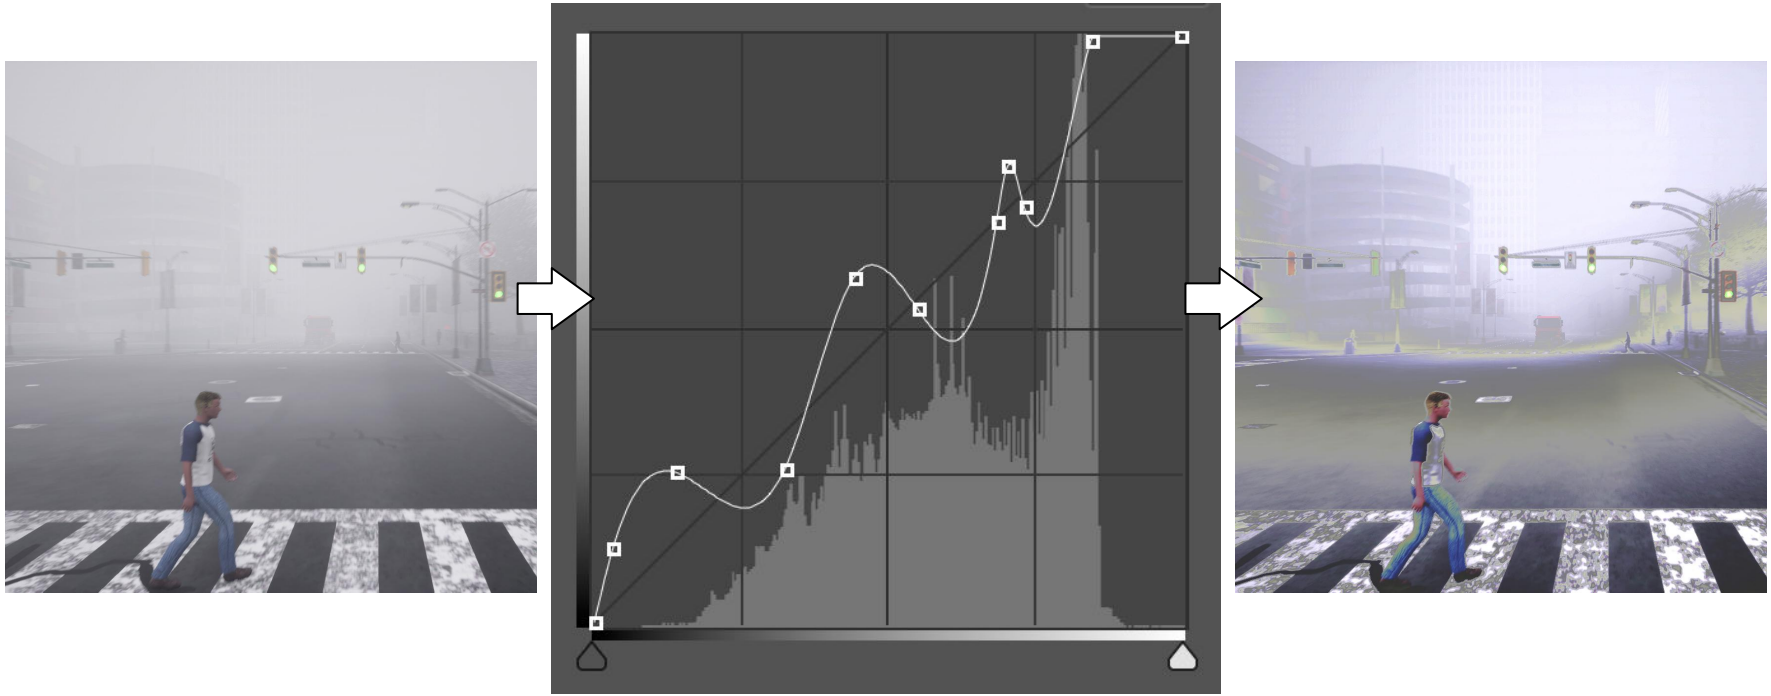

Deweathering: Domain Translation

Ground Truth

Foggy Image On Clear Day Trained Model

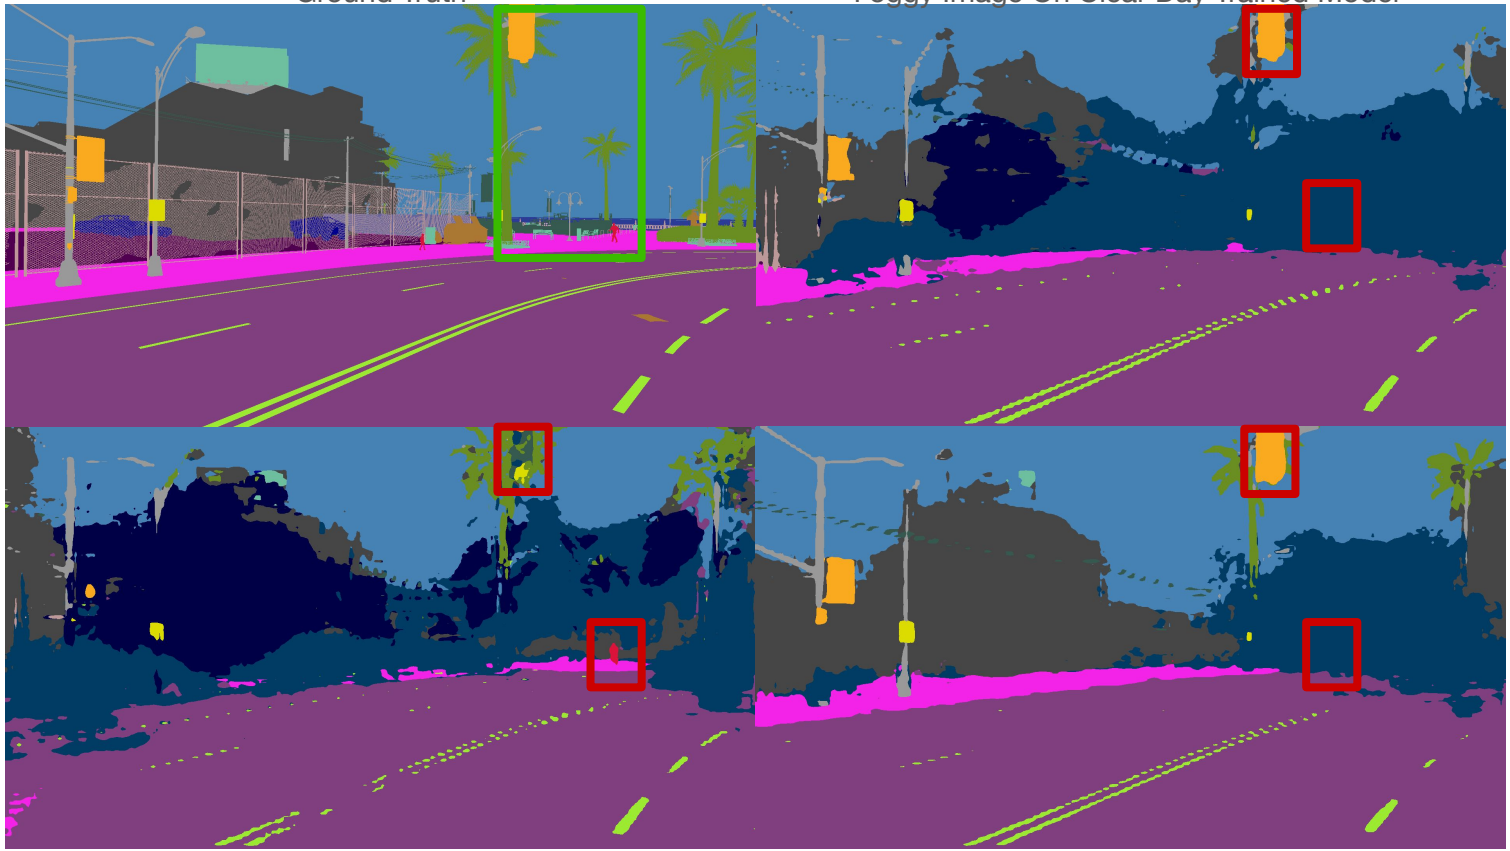

Deweathered (IP) Foggy Image

Deweathered (DT) Foggy Image

Deweathering: Quantitative Results

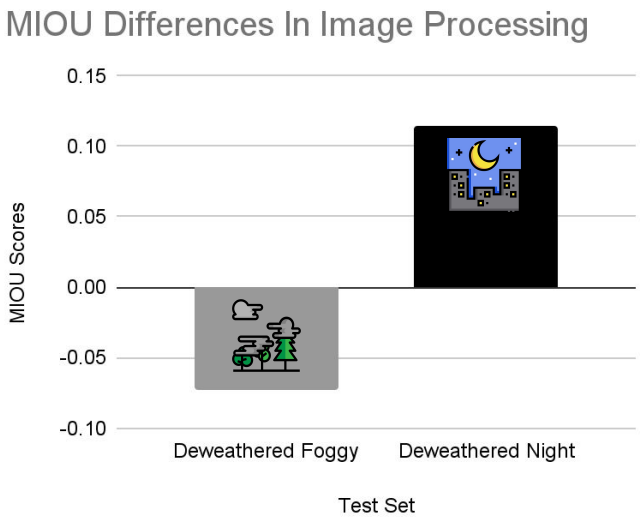

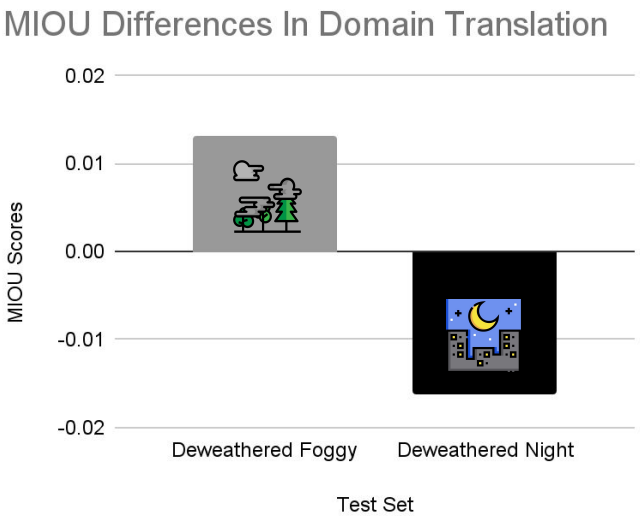

Training Set

* 0 being our baseline MIOU score

Image Processing: Quantitative Results

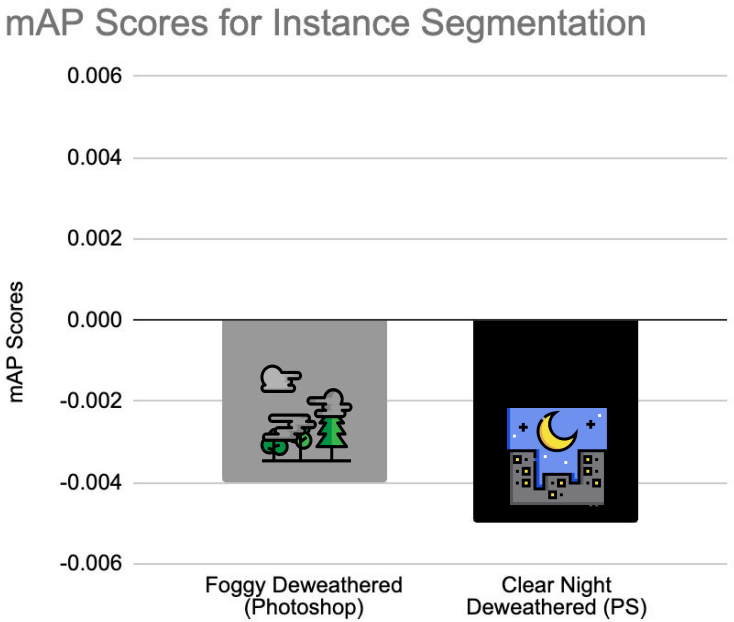

Training Set

* 0 being our baseline mAP score

Sensor Fusion

LIDAR

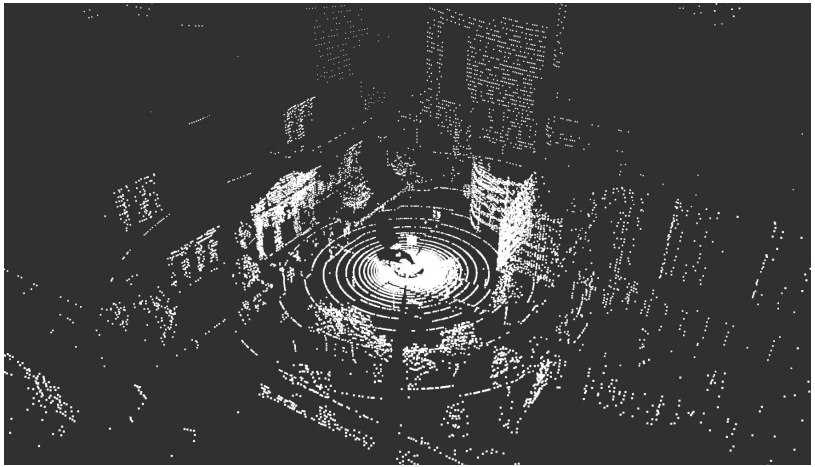

Camera

Data Transformation

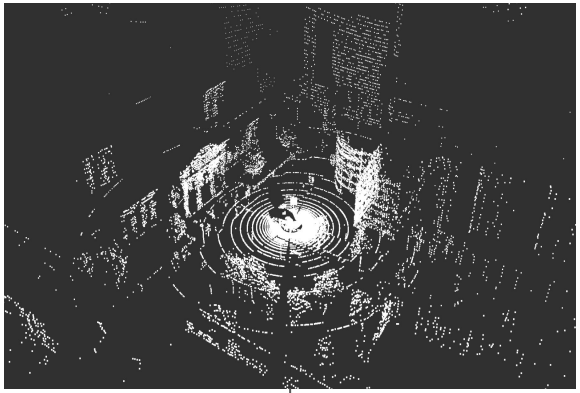

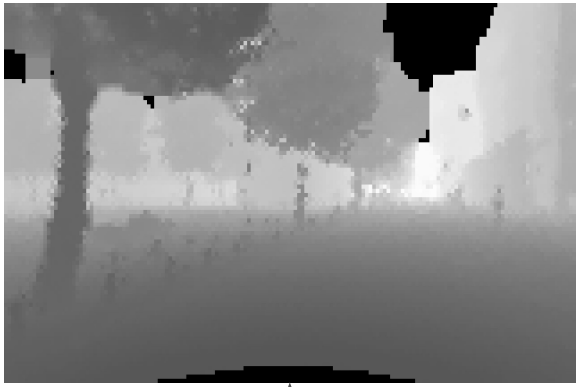

Projection

Max Pooling +
Interpolation

Mid Fusion Strategy

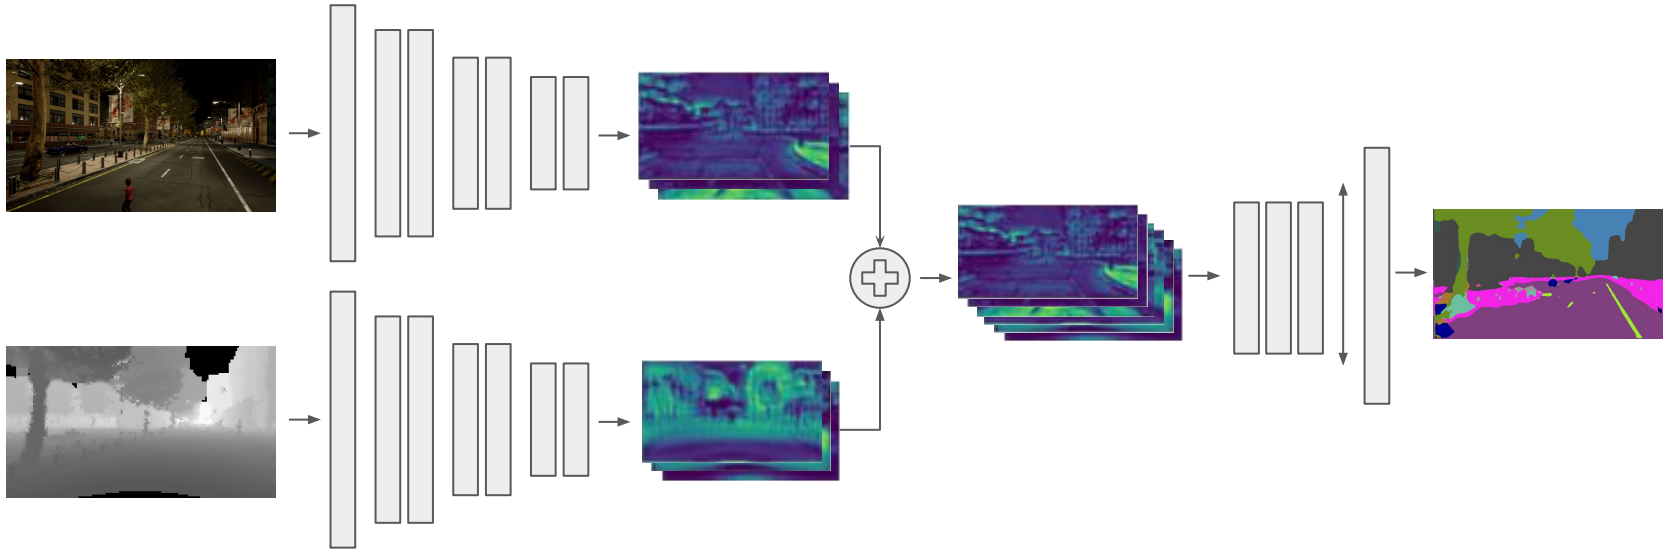

Sensor Fusion Baseline

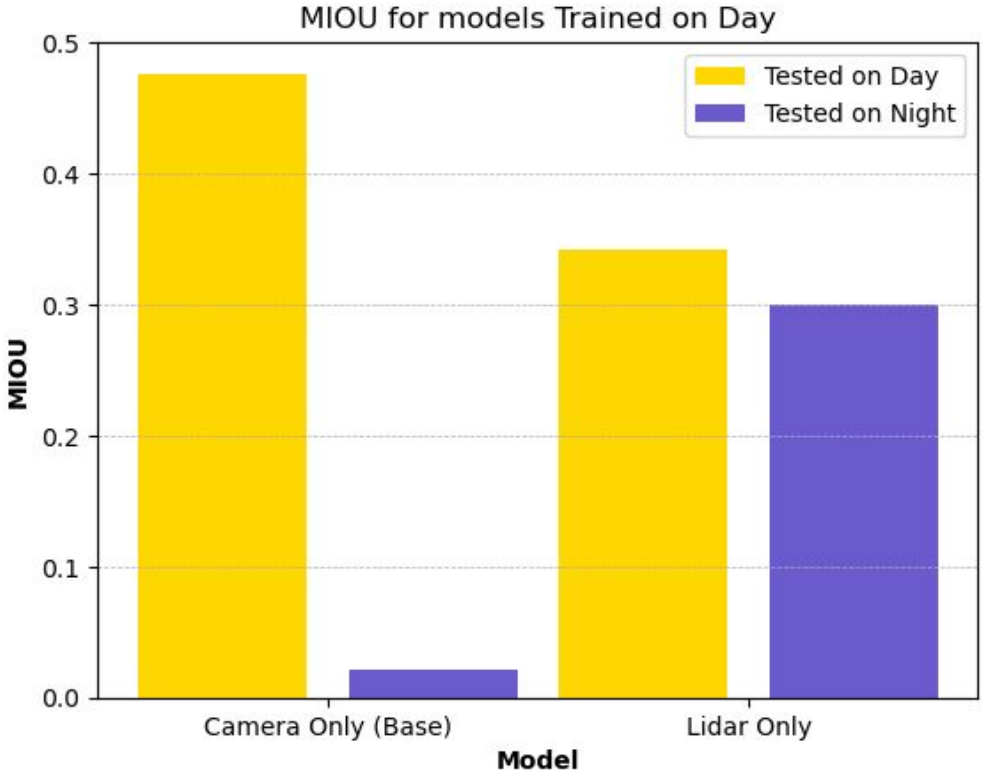

Sensor Fusion Results

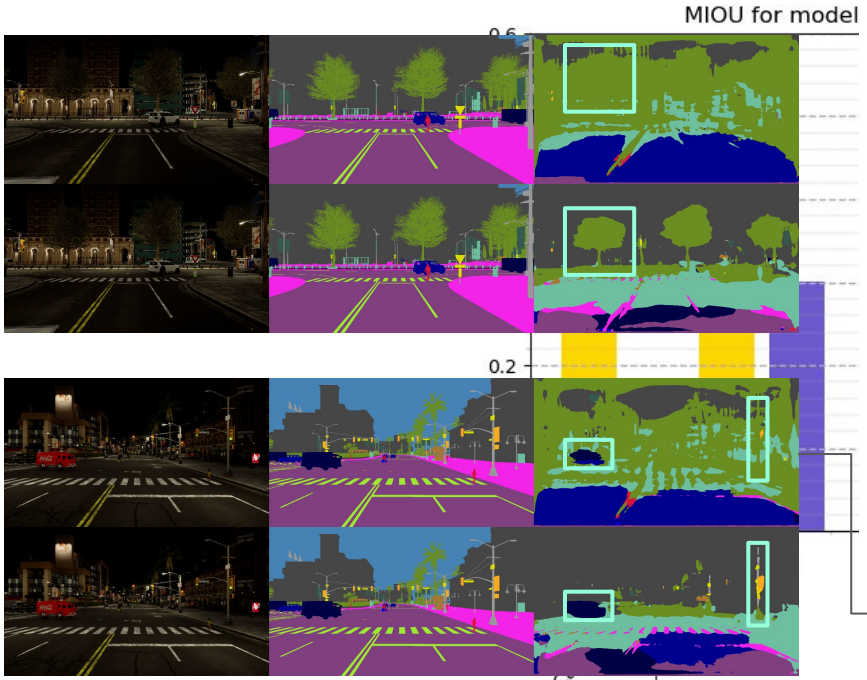

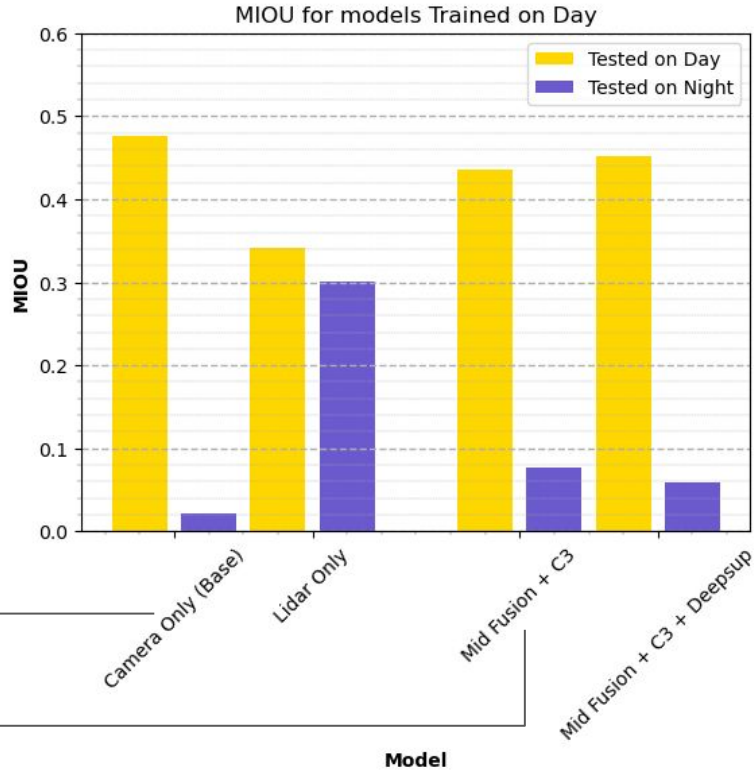

1. Defining computer vision + background

2. Measuring the effect of weather

3. Mitigating the effect of weather

4. Future work + conclusions

Improving MaskTrackRCNN's Tracking Branch

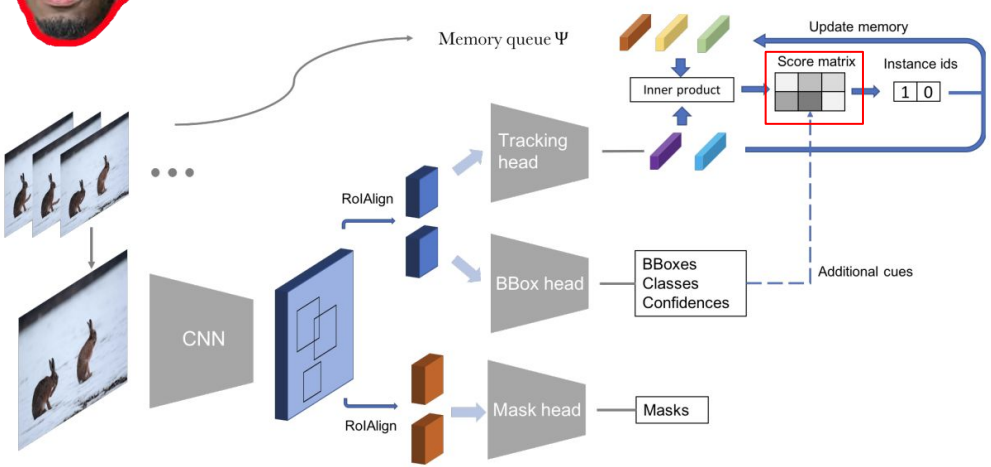

1. Greedy vs Globally Optimal Association

	1	2	3
a	0.9	0.8	0.3
b	0.7	0.6	0.9
c	0.4	0.6	0.8

	1	2	3
a	0.9	0.8	0.3
b	0.7	0.6	0.9
c	0.4	0.6	0.8

2. Incorporating Motion Cues in Score Calculation

$$v_i(n) = \log p_i(n) + \alpha \log s_i + \beta \mathrm{IoU}(b_i, b_n) + \gamma \delta(c_i, c_n) \quad (3)$$

+ ?

Future Work: Adding Snow to Carla

Future Work: Train On LUTs

https://pixabay.com/images/search/nvidia/

Future Work: Sim to Real

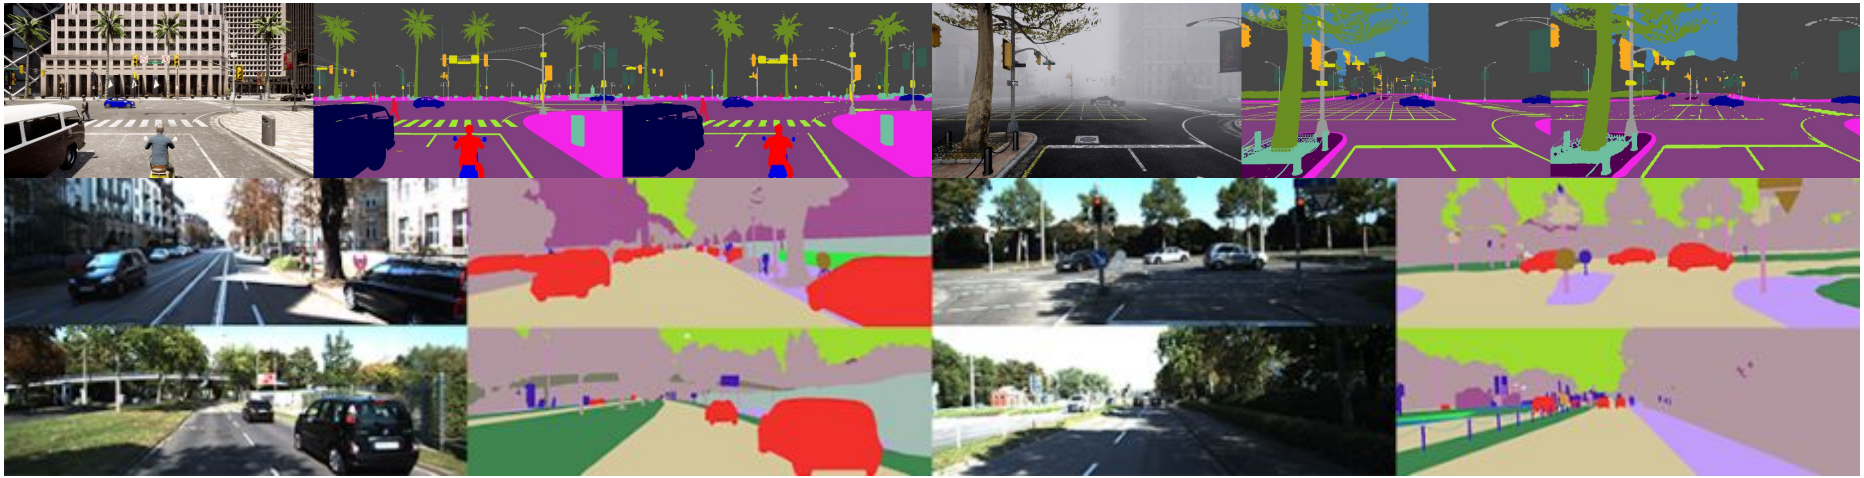

Future Work: Real Time Performance

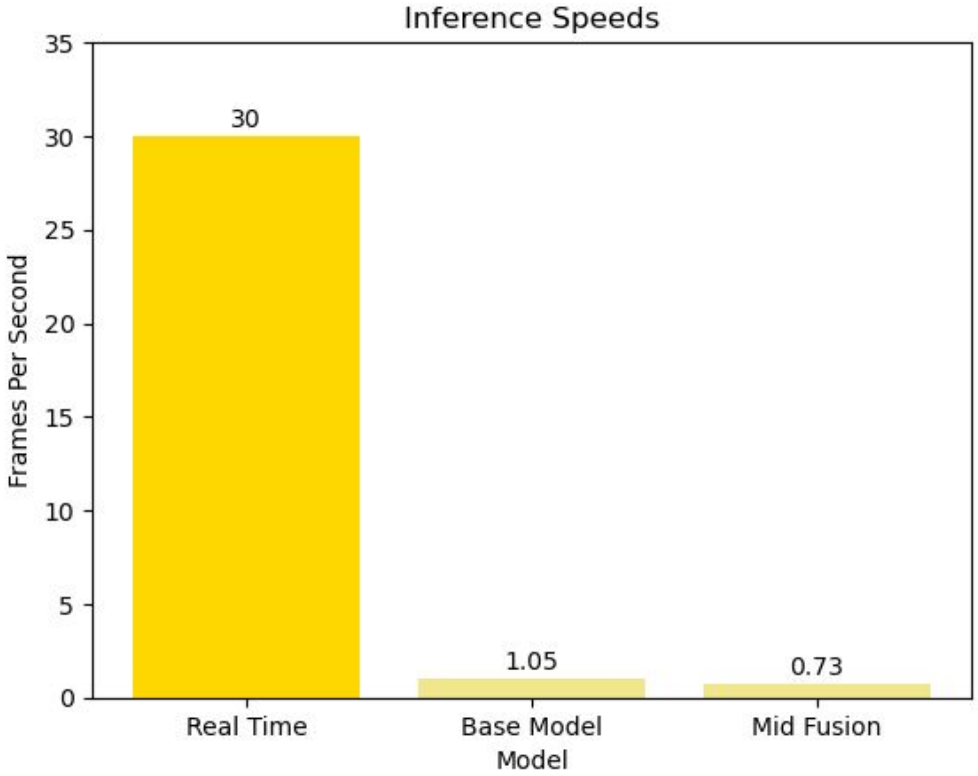

Future Work: Training with all weathers

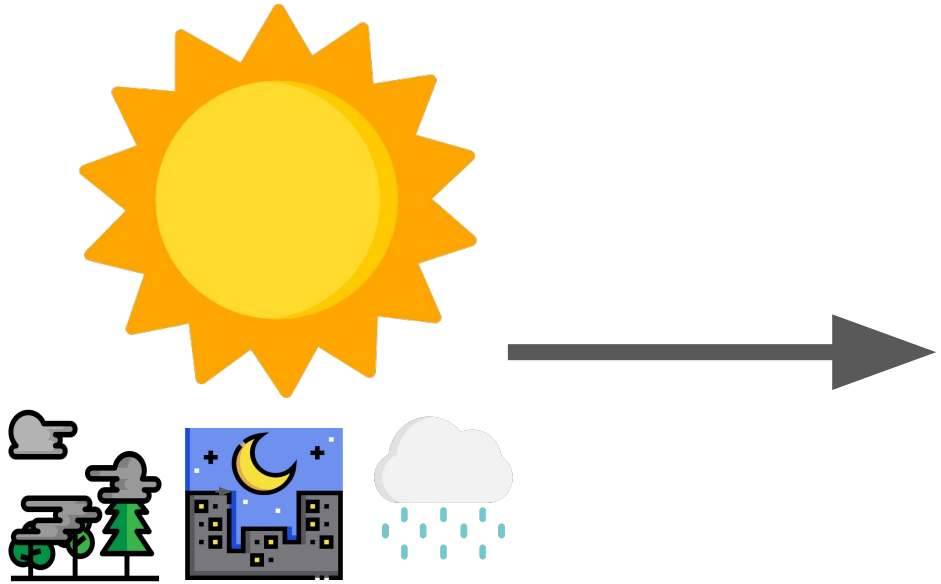

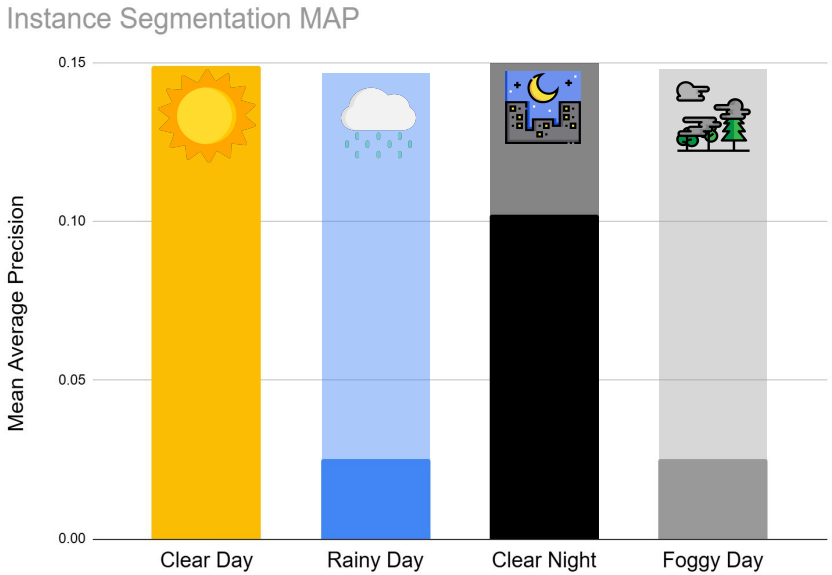

What have we talked about?

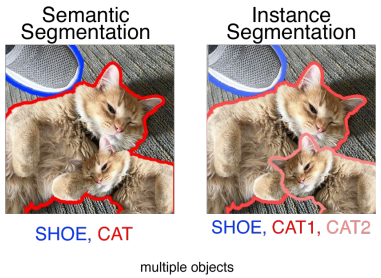

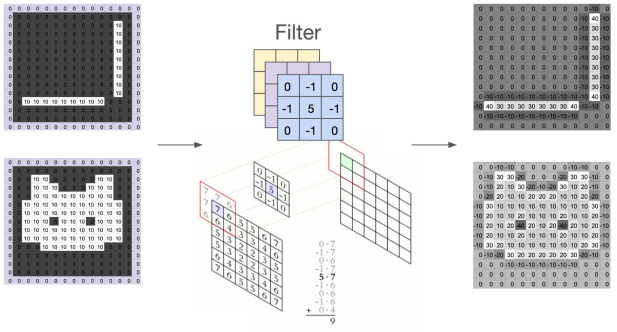

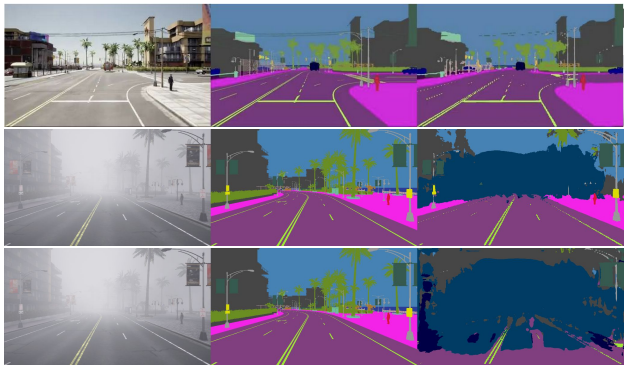

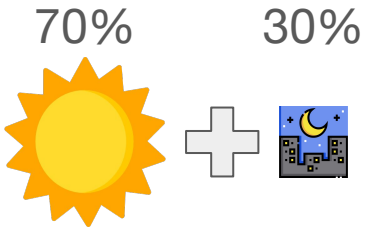

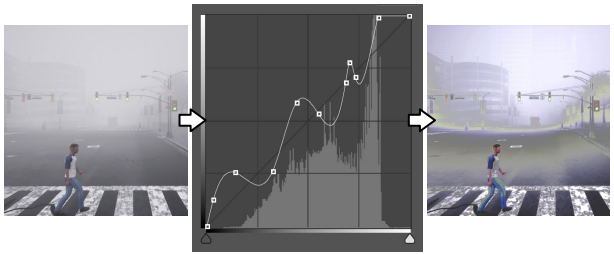

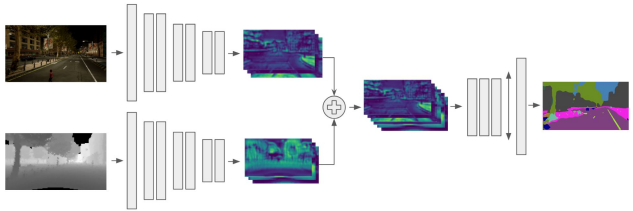

Acknowledgments (the GOATs)

Professor Tanya Amert

Mike Tie

Image Attribution

- https://giphy.com/gifs/supercompressor-dallas-best-driving-roads-5xtDarLIo2eSlHf2egE
- https://www.weather.gov/lot/2024dec04
- https://giphy.com/gifs/clsHRed0PxHhmGx1Gn
- https://avikkundu.medium.com/performing-instance-segmentation-on-x-ray-images-with-mask-r-cnn-761dbca23511
- https://www.flaticon.com/free-icons/sun
- https://www.flaticon.com/free-icons/rain
- https://www.flaticon.com/free-icons/haze
- https://www.flaticon.com/free-icons/night
- https://www.flaticon.com/free-icons/document
- https://www.shutterstock.com/image-vector/auto-parts-car-spare-set-600nw-2269761985.jpg
- https://www.automotivedive.com/news/waymo-6th-generation-driver-autonomous-driving-hardware-robotaxi-lidar-ai/725519/
- https://www.google.com/url?sa=i&url=https%3A%2F%2Fsemiwiki.com%2Feda%2F297752-bouncing-off-the-walls-how-real-time-radar-is-accelerating-the-development-of-autonomous-vehicles%2F&psig=AOvVaw1IFGJzE1PPZbYr_jMxdjzI&ust=1740785506671000&source=images&cd=vfe&opi=89978449&ved=0CBcQjhxqFwoTCIj1rcKB5YsDFQAAAAAdAAAAABAE
- https://www.foxbusiness.com/lifestyle/uber-waymo-expand-partnership-austin-atlanta
- https://www.automotivedive.com/news/waymo-6th-generation-driver-autonomous-driving-hardware-robotaxi-lidar-ai/725519/
- https://avikkundu.medium.com/performing-instance-segmentation-on-x-ray-images-with-mask-r-cnn-761dbca23511
- https://www.flaticon.com/free-icon/car_846338?term=car&page=1&position=12&origin=tag&related_id=846338
- https://www.cvlibs.net/datasets/kitti/eval_semseg.php?benchmark=semantics2015
- https://en.m.wikipedia.org/wiki/File:2D_Convolution_Animation.gif
- https://www.google.com/url?sa=i&url=https%3A%2F%2Fen.m.wikipedia.org%2Fwiki%2FFile%3A2D_Convolution_Animation.gif&psig=AOvVaw05cdvrgy8Xm95G8Z6_QLIS&ust=1740813440265000&source=images&cd=vfe&opi=89978449&ved=0CBQQjhxqFwoTCOjM2Mrp5YsDFQAAAAAdAAAAABAs
- https://en.m.wikipedia.org/wiki/File:2D_Convolution_Animation.gif
- https://static.vecteezy.com/system/resources/thumbnails/018/878/198/small_2x/empty-straight-road-with-marking-overhead-view-highway-seamless-pattern-roadway-horizontal-template-isolated-on-white-background-carriageway-element-of-city-map-vector.jpg
- https://avikkundu.medium.com/performing-instance-segmentation-on-x-ray-images-with-mask-r-cnn-761dbca23511
- https://www.robertoamoroso.it/publication/bernhard-2024-miou/featured_hu7326fcf2264b218ea2c76406ba8d7e41_53827_720x2500_fit_q75_h2_lanczos_3.webp

Bibliography

Dosovitskiy, A., Ros, G., Codevilla, F., López, A., and Koltun, A., “CARLA: An Open Urban Driving Simulator,” 2017. Available: https://proceedings.mlr.press/v78/dosovitskiy17a/dosovitskiy17a.pdf

Bijelic, M., et al, "Seeing Through Fog Without Seeing Fog: Deep Multimodal Sensor Fusion in Unseen Adverse Weather," in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 11679-11689.

Chen, K., et al. "MMDetection: Open MMLab Detection Toolbox and Benchmark," in arXiv preprint arXiv:1906.07155, 2019.

He, K., et al, "Mask R-CNN," 2018.

Yang, L., Fan, Y., and Xu N., ‘Video Instance Segmentation’, 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 5187–5196, 2019.

N. Aharon, R. Orfaig, B. Bobrovsky. "BoT-SORT: Robust Associations Multi-Pedestrian Tracking," in arXiv preprint arXiv:2206.14651, 2022.

Patil, P.W., Gupta, S., Rana, S., Venkatesh, S., and Murala, S., “Multi-weather Image Restoration via Domain Translation,” in Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), Oct. 2023, pp. 21696–21705.

Zhou, B., et al. "Semantic understanding of scenes through the ade20k dataset," in International Journal on Computer Vision, 2018.

Zhou, B., et al, "Scene Parsing through ADE20K Dataset," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017.

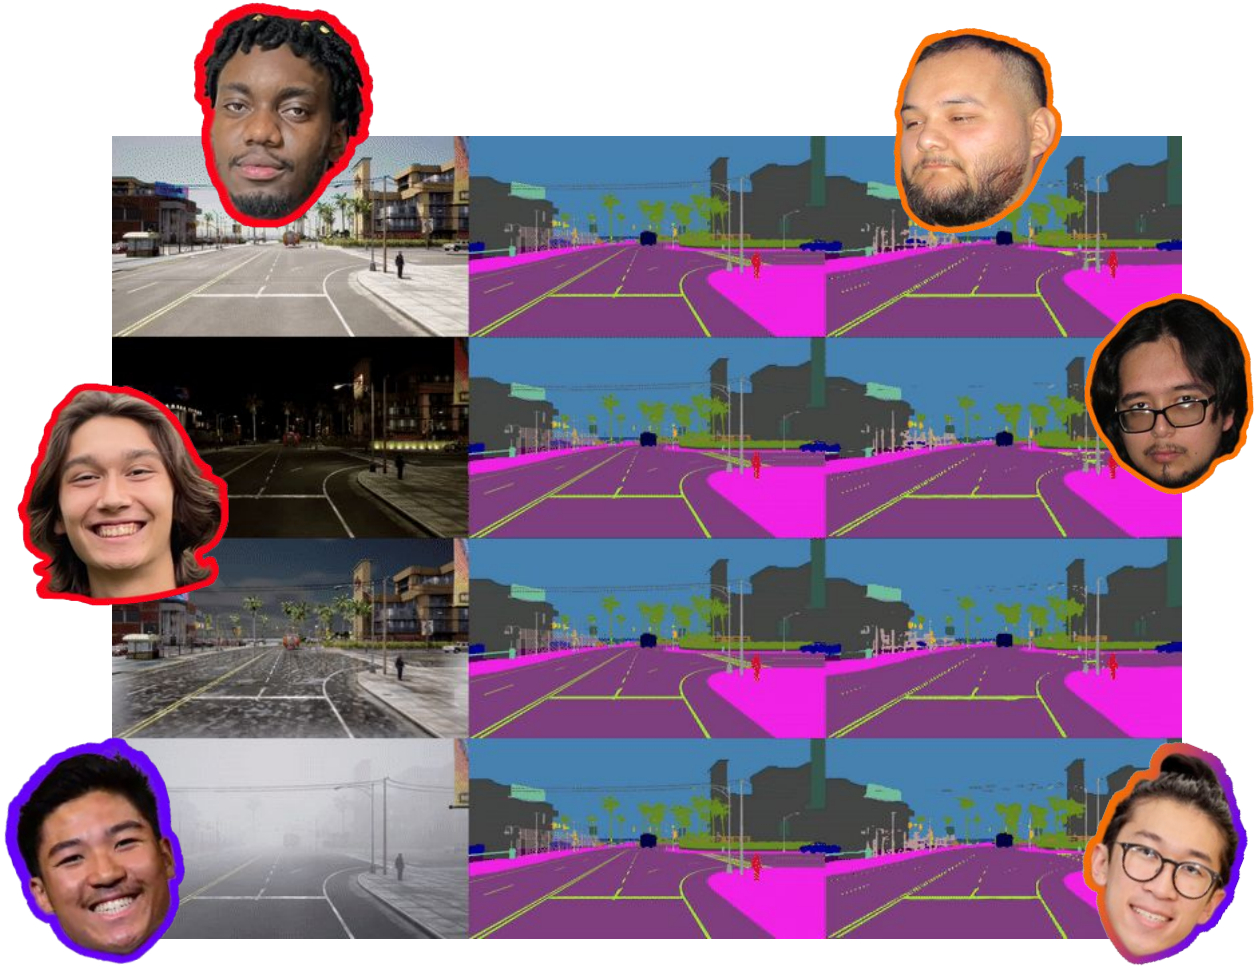

Supplementary Material

Metrics: Not all are equal

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

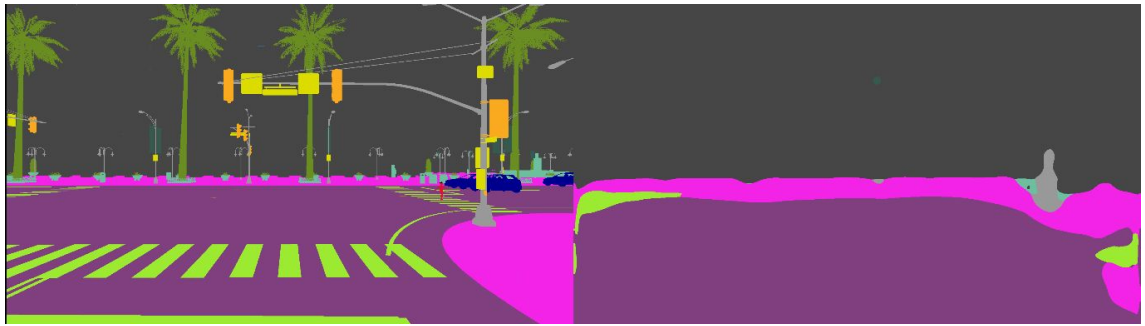

Early Fusion Architecture

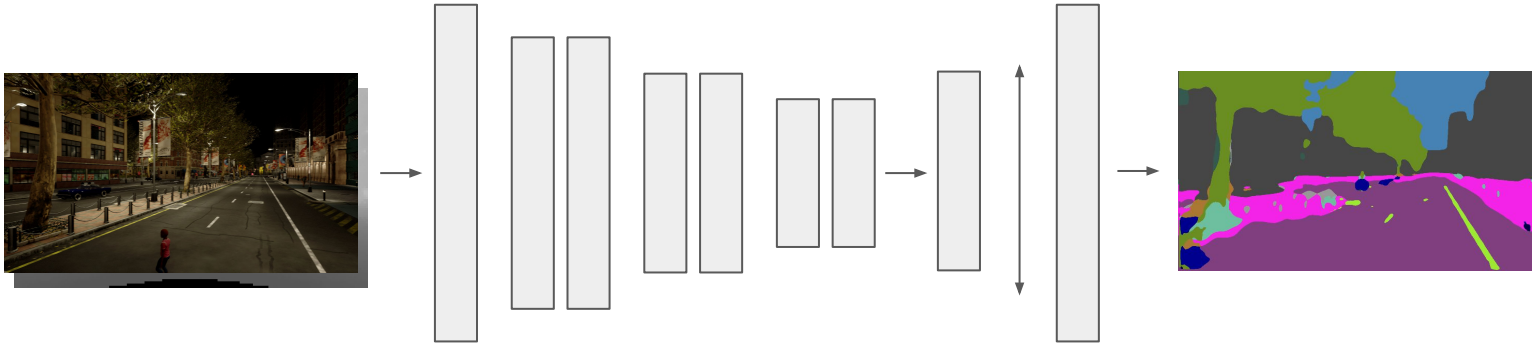

Deep Supervision Architecture

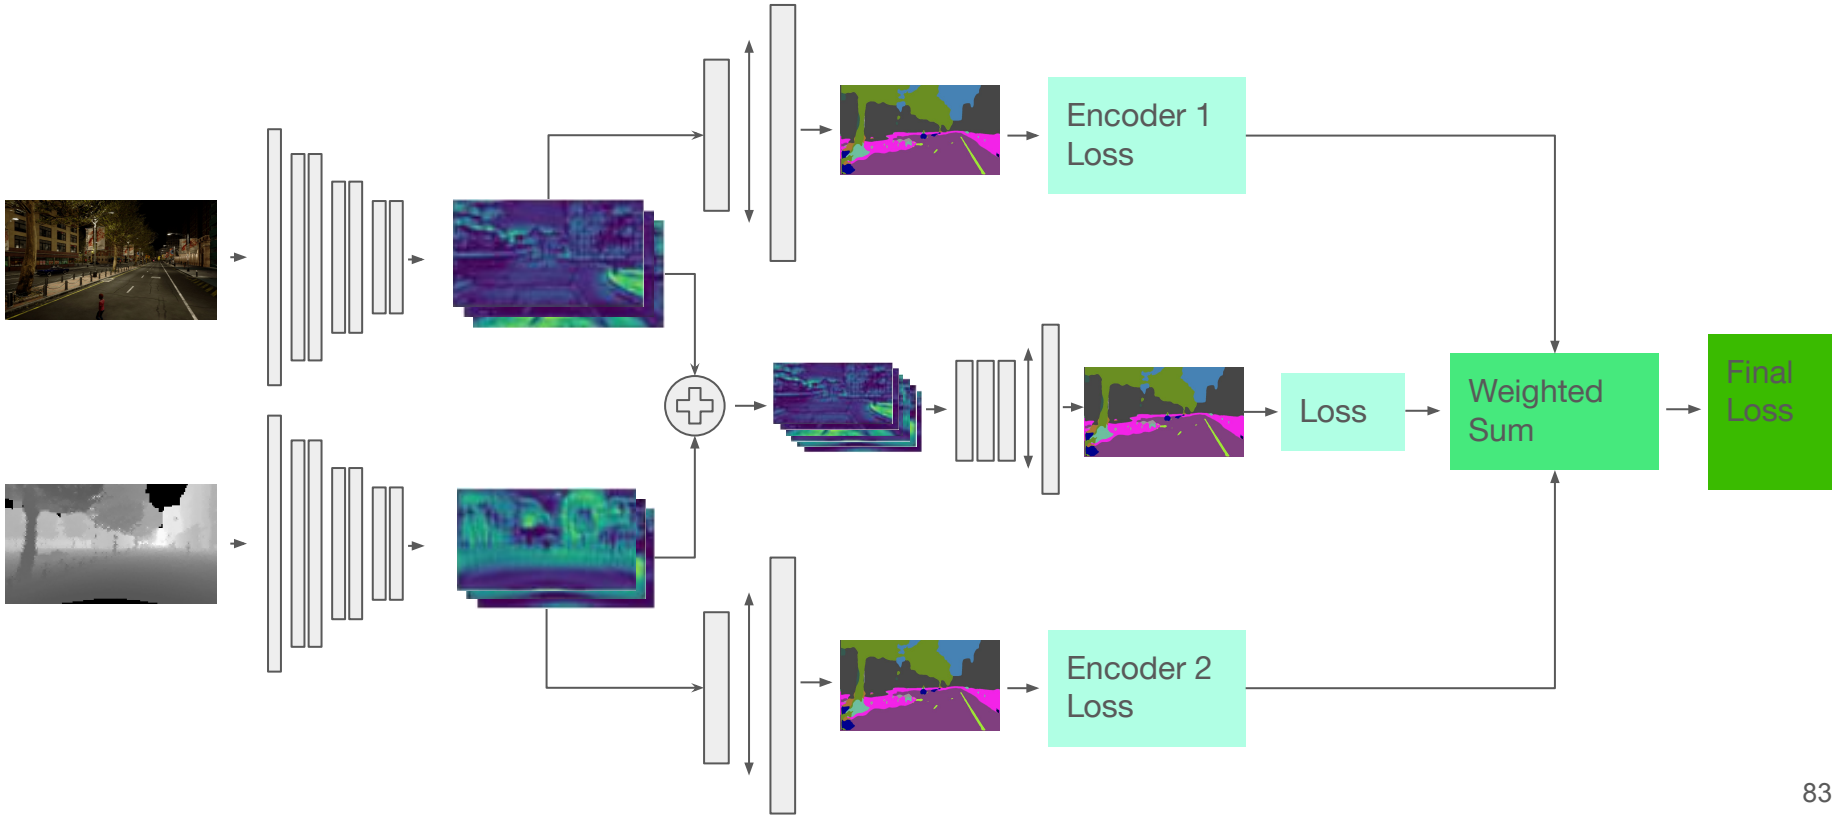

Entropy Fusion

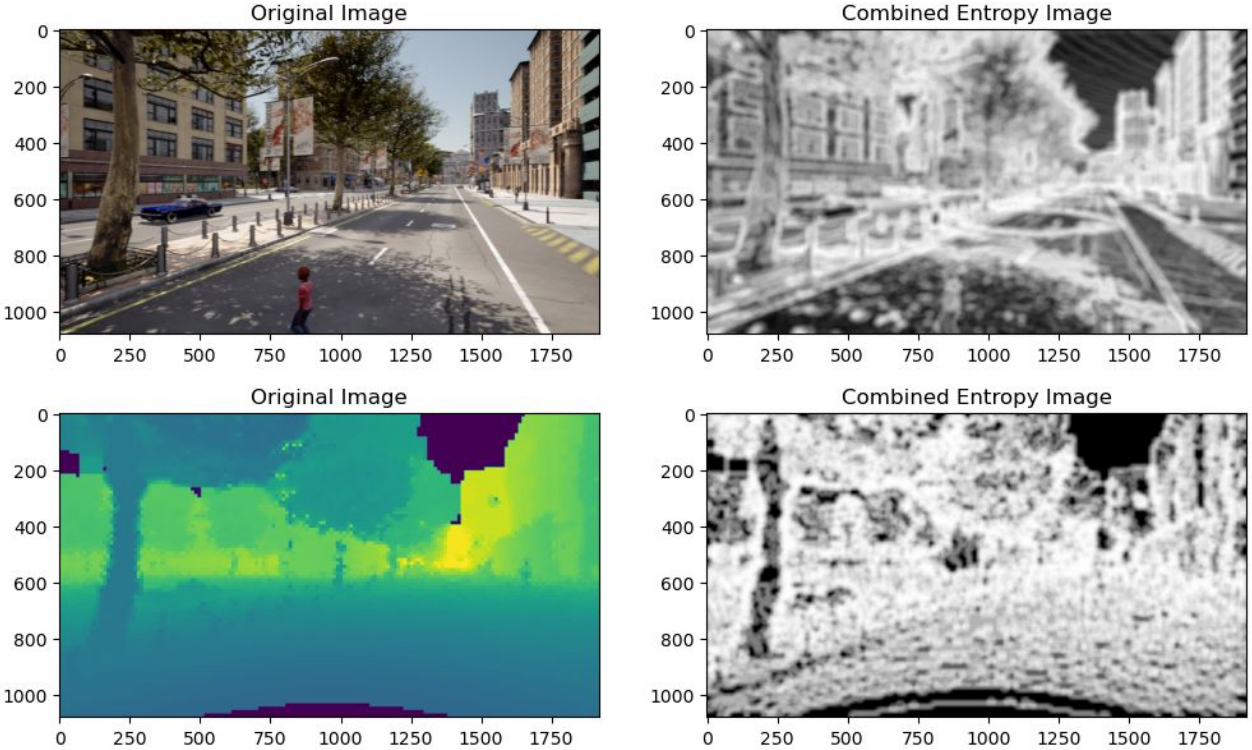

Mid Fusion: Surgery Qualitative Results

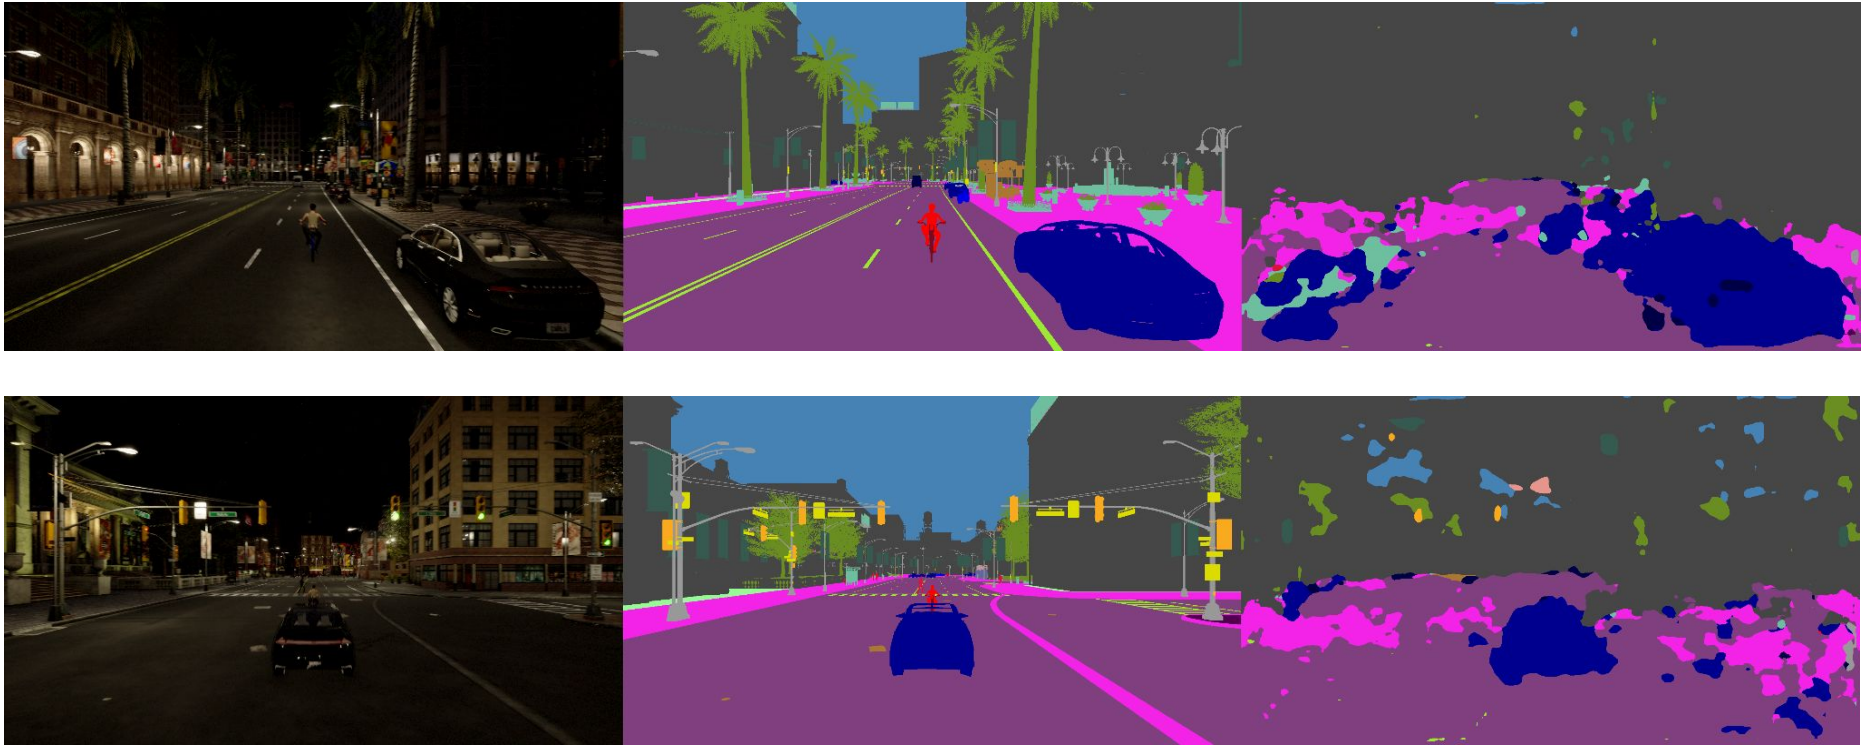

Lidar Only: Qualitative Results

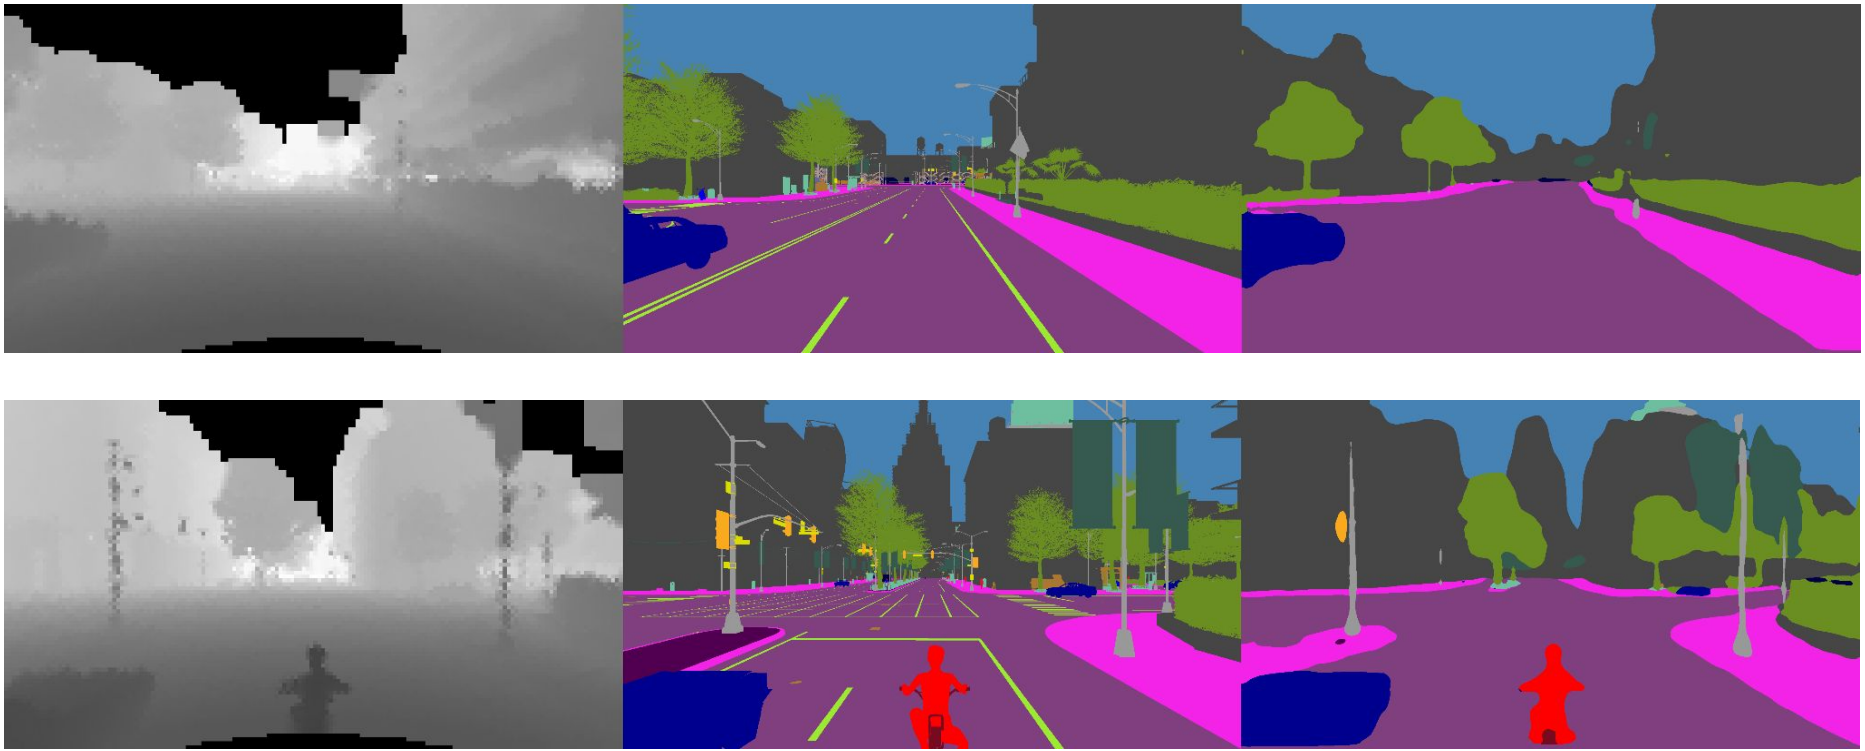

Sparse LIDAR: Qualitative Results

Depth Encoder Feature Maps

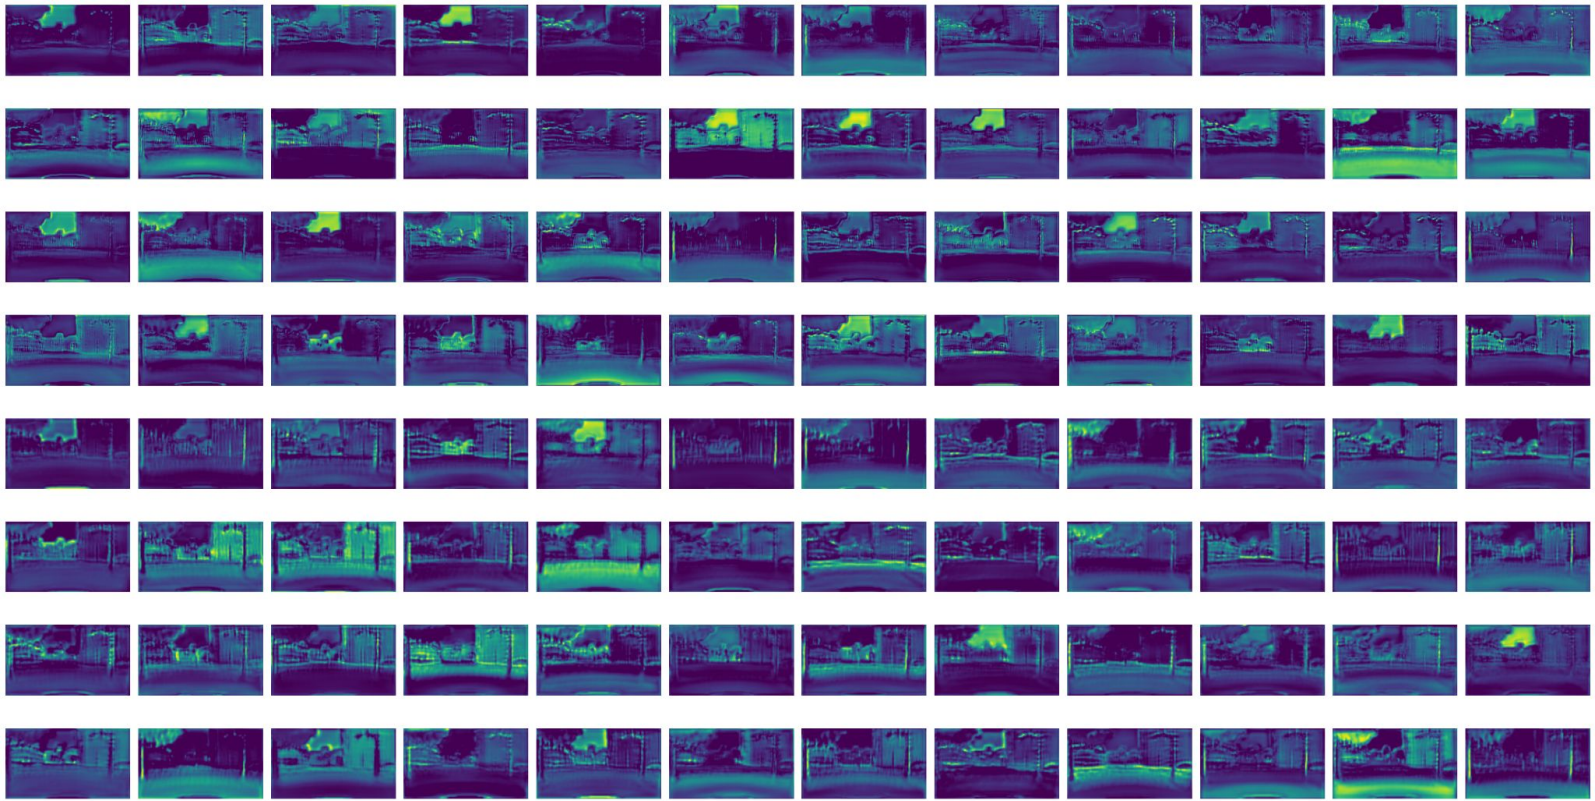

Instance segmentation metrics

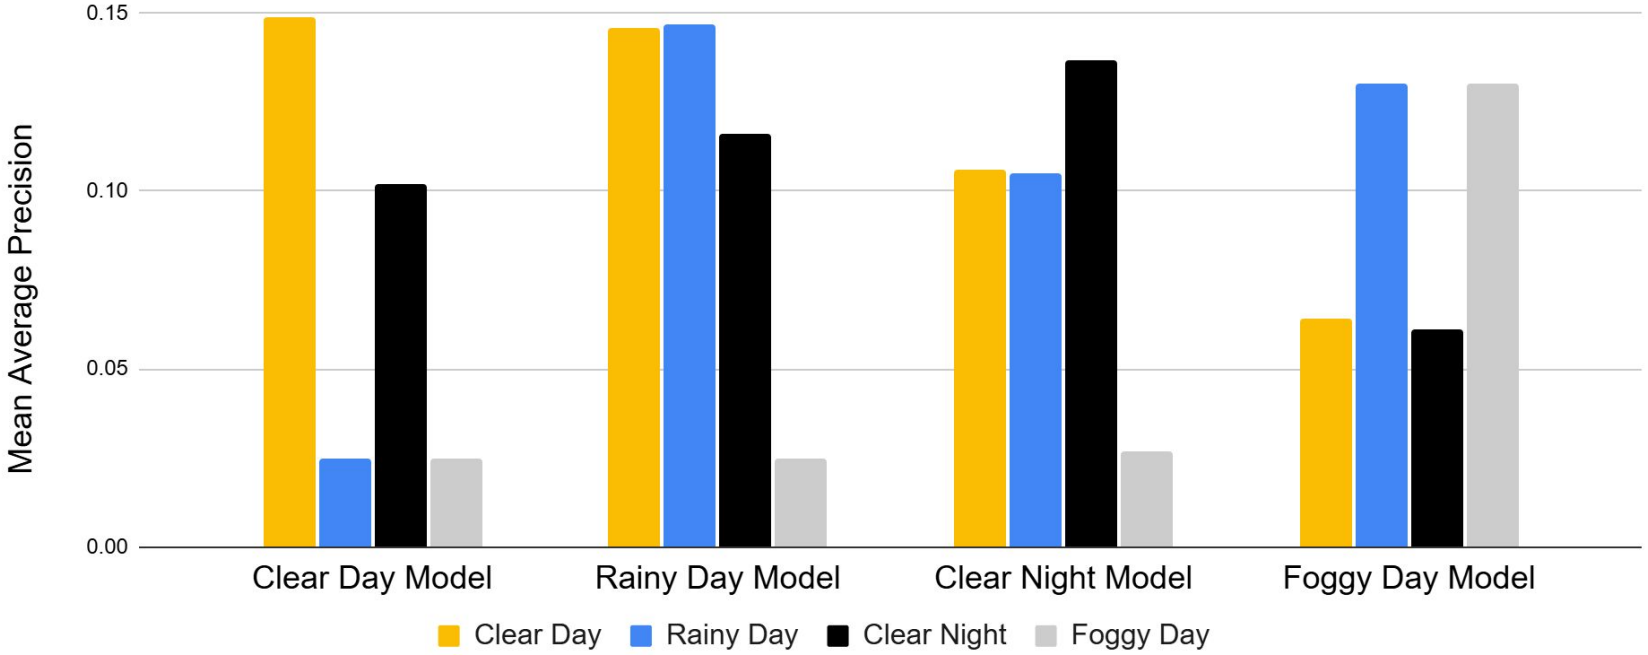

Training Set

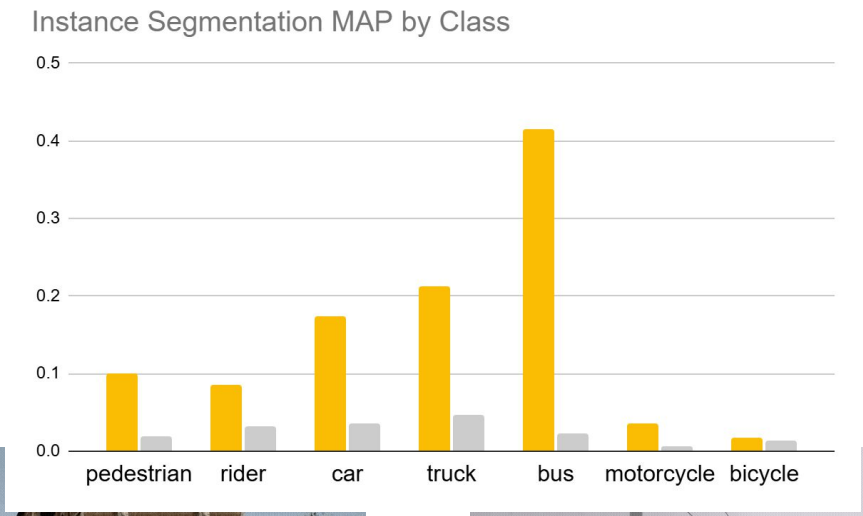

Training Set