ICRA 2020
31 May • 4 June 2020
Palais des Congrès de Paris
Paris • FRANCE

Learning to Optimally Segment Point Clouds

Peiyun Hu, David Held, Deva Ramanan

THE ROBOTICS INSTITUTE
Carnegie Mellon University

Paper ID: 2977

Raw LiDAR Scans

Today, most autonomous vehicles perceive the world through LiDAR point clouds.

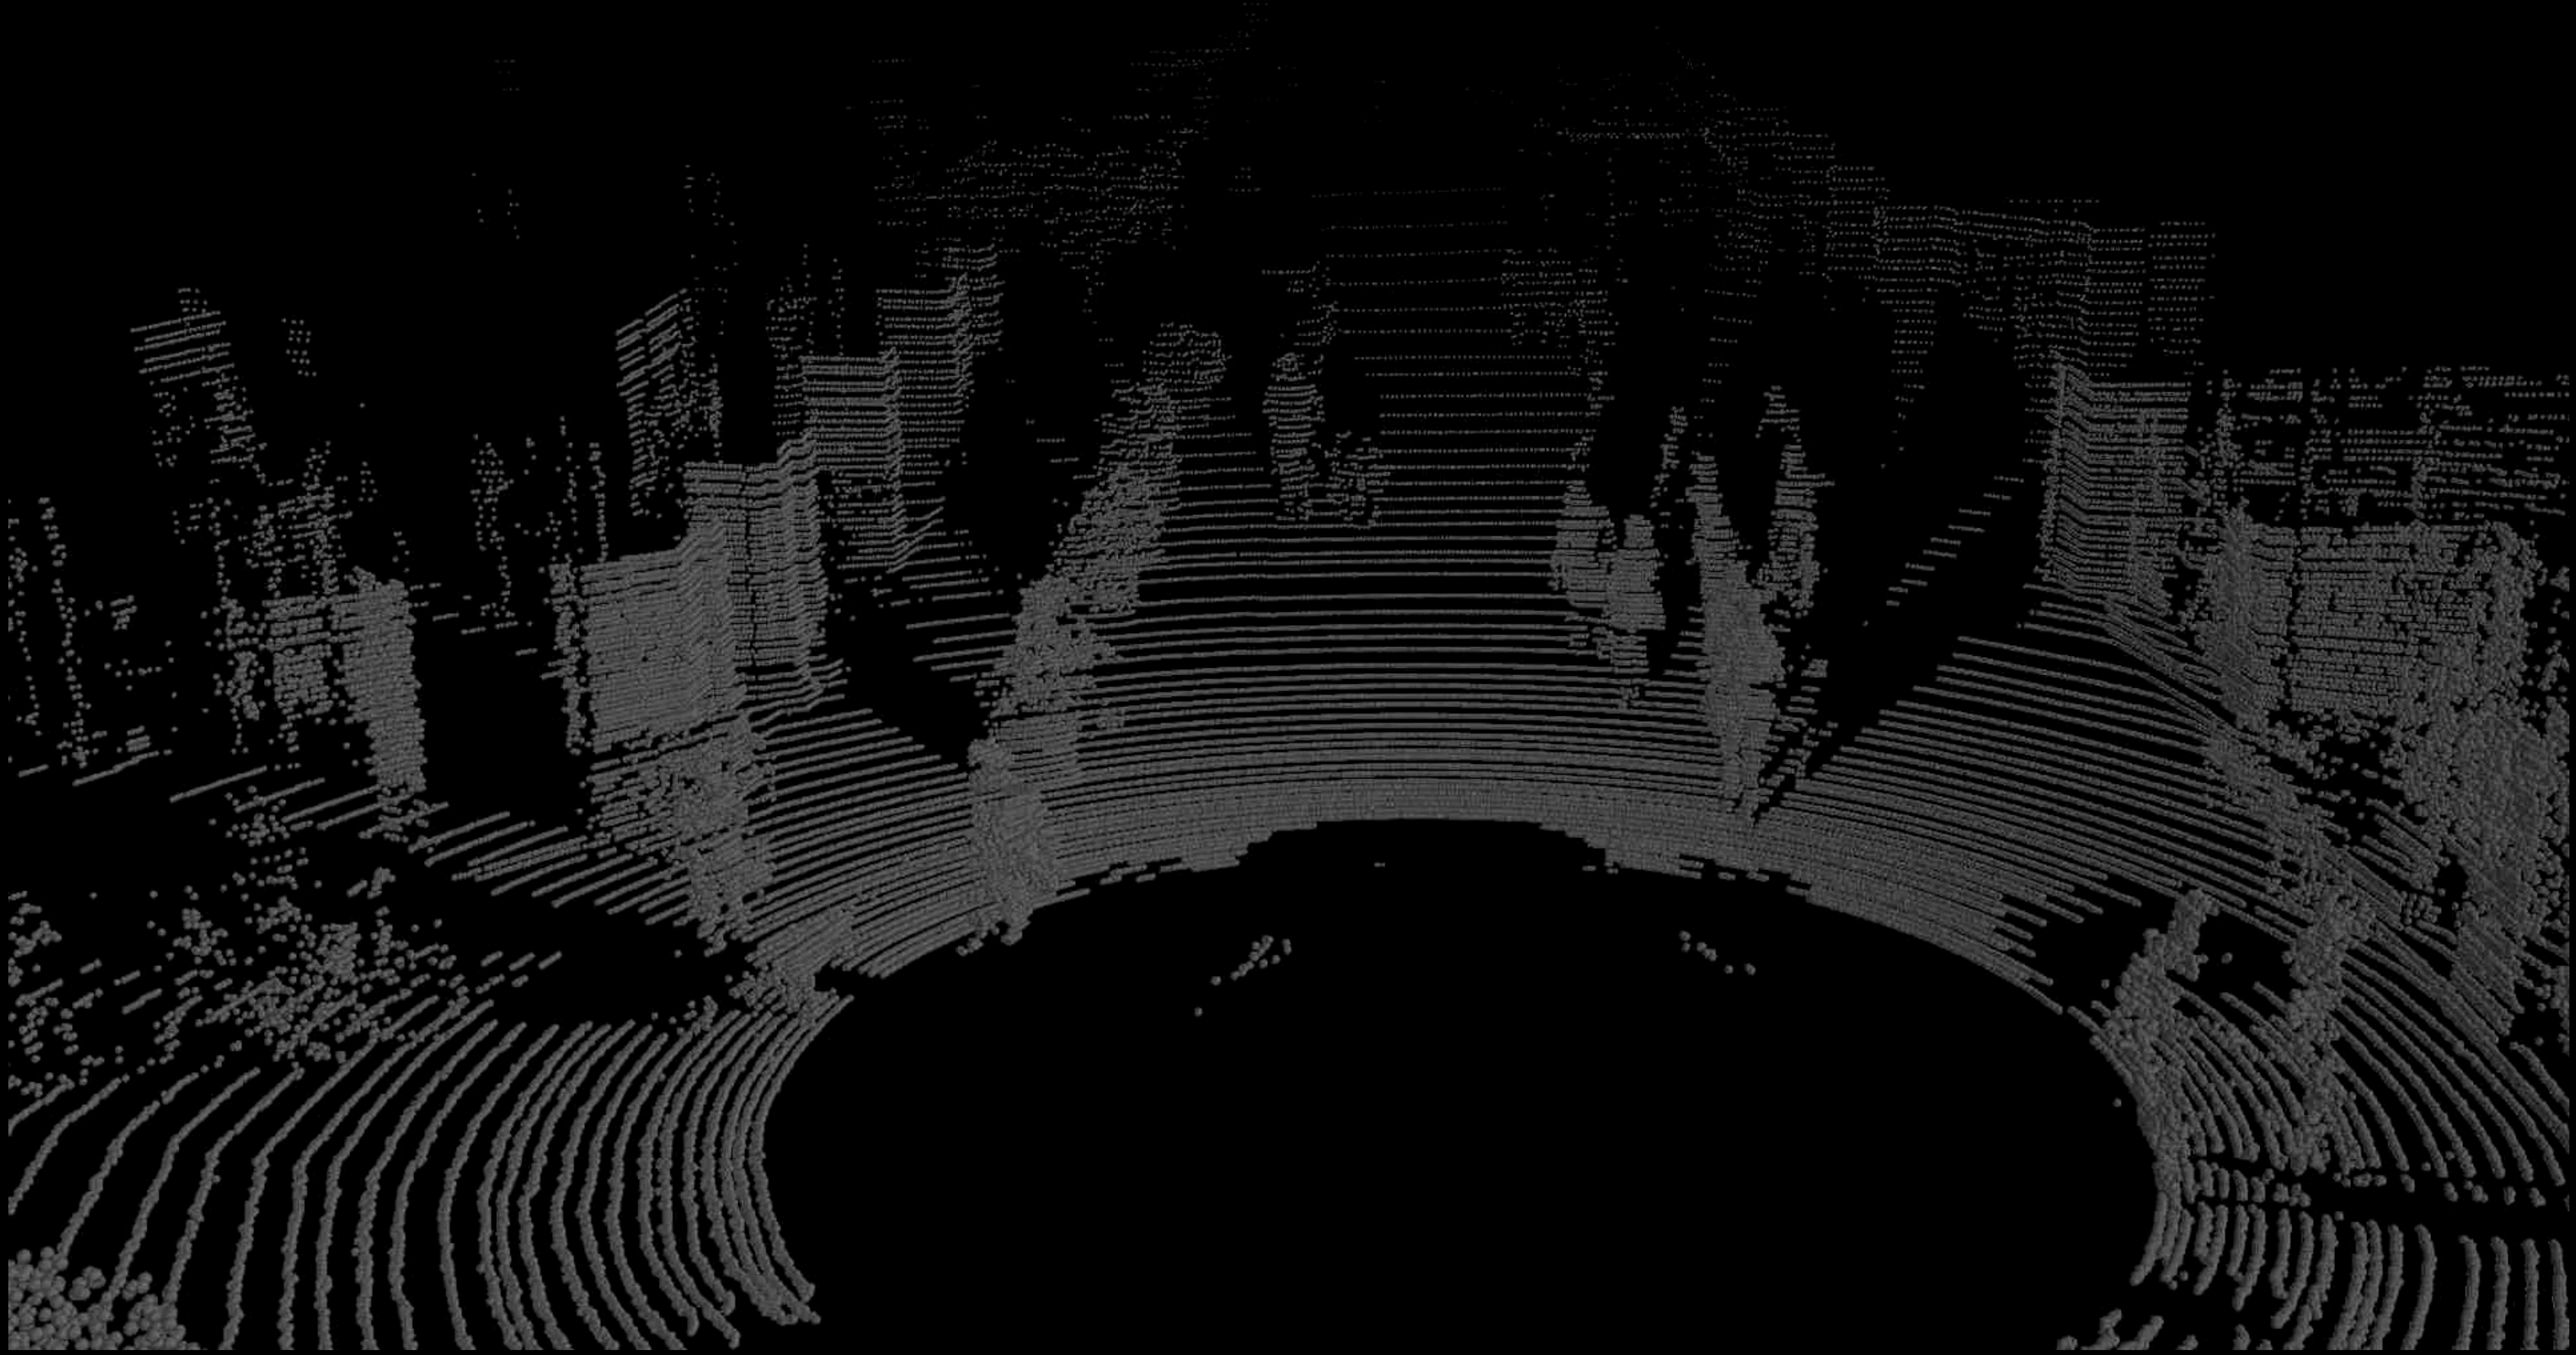

Map-based Preprocessing

They often use pre-built maps to first filter out points from the background, then run clustering on points from the foreground to obtain object-level perception.

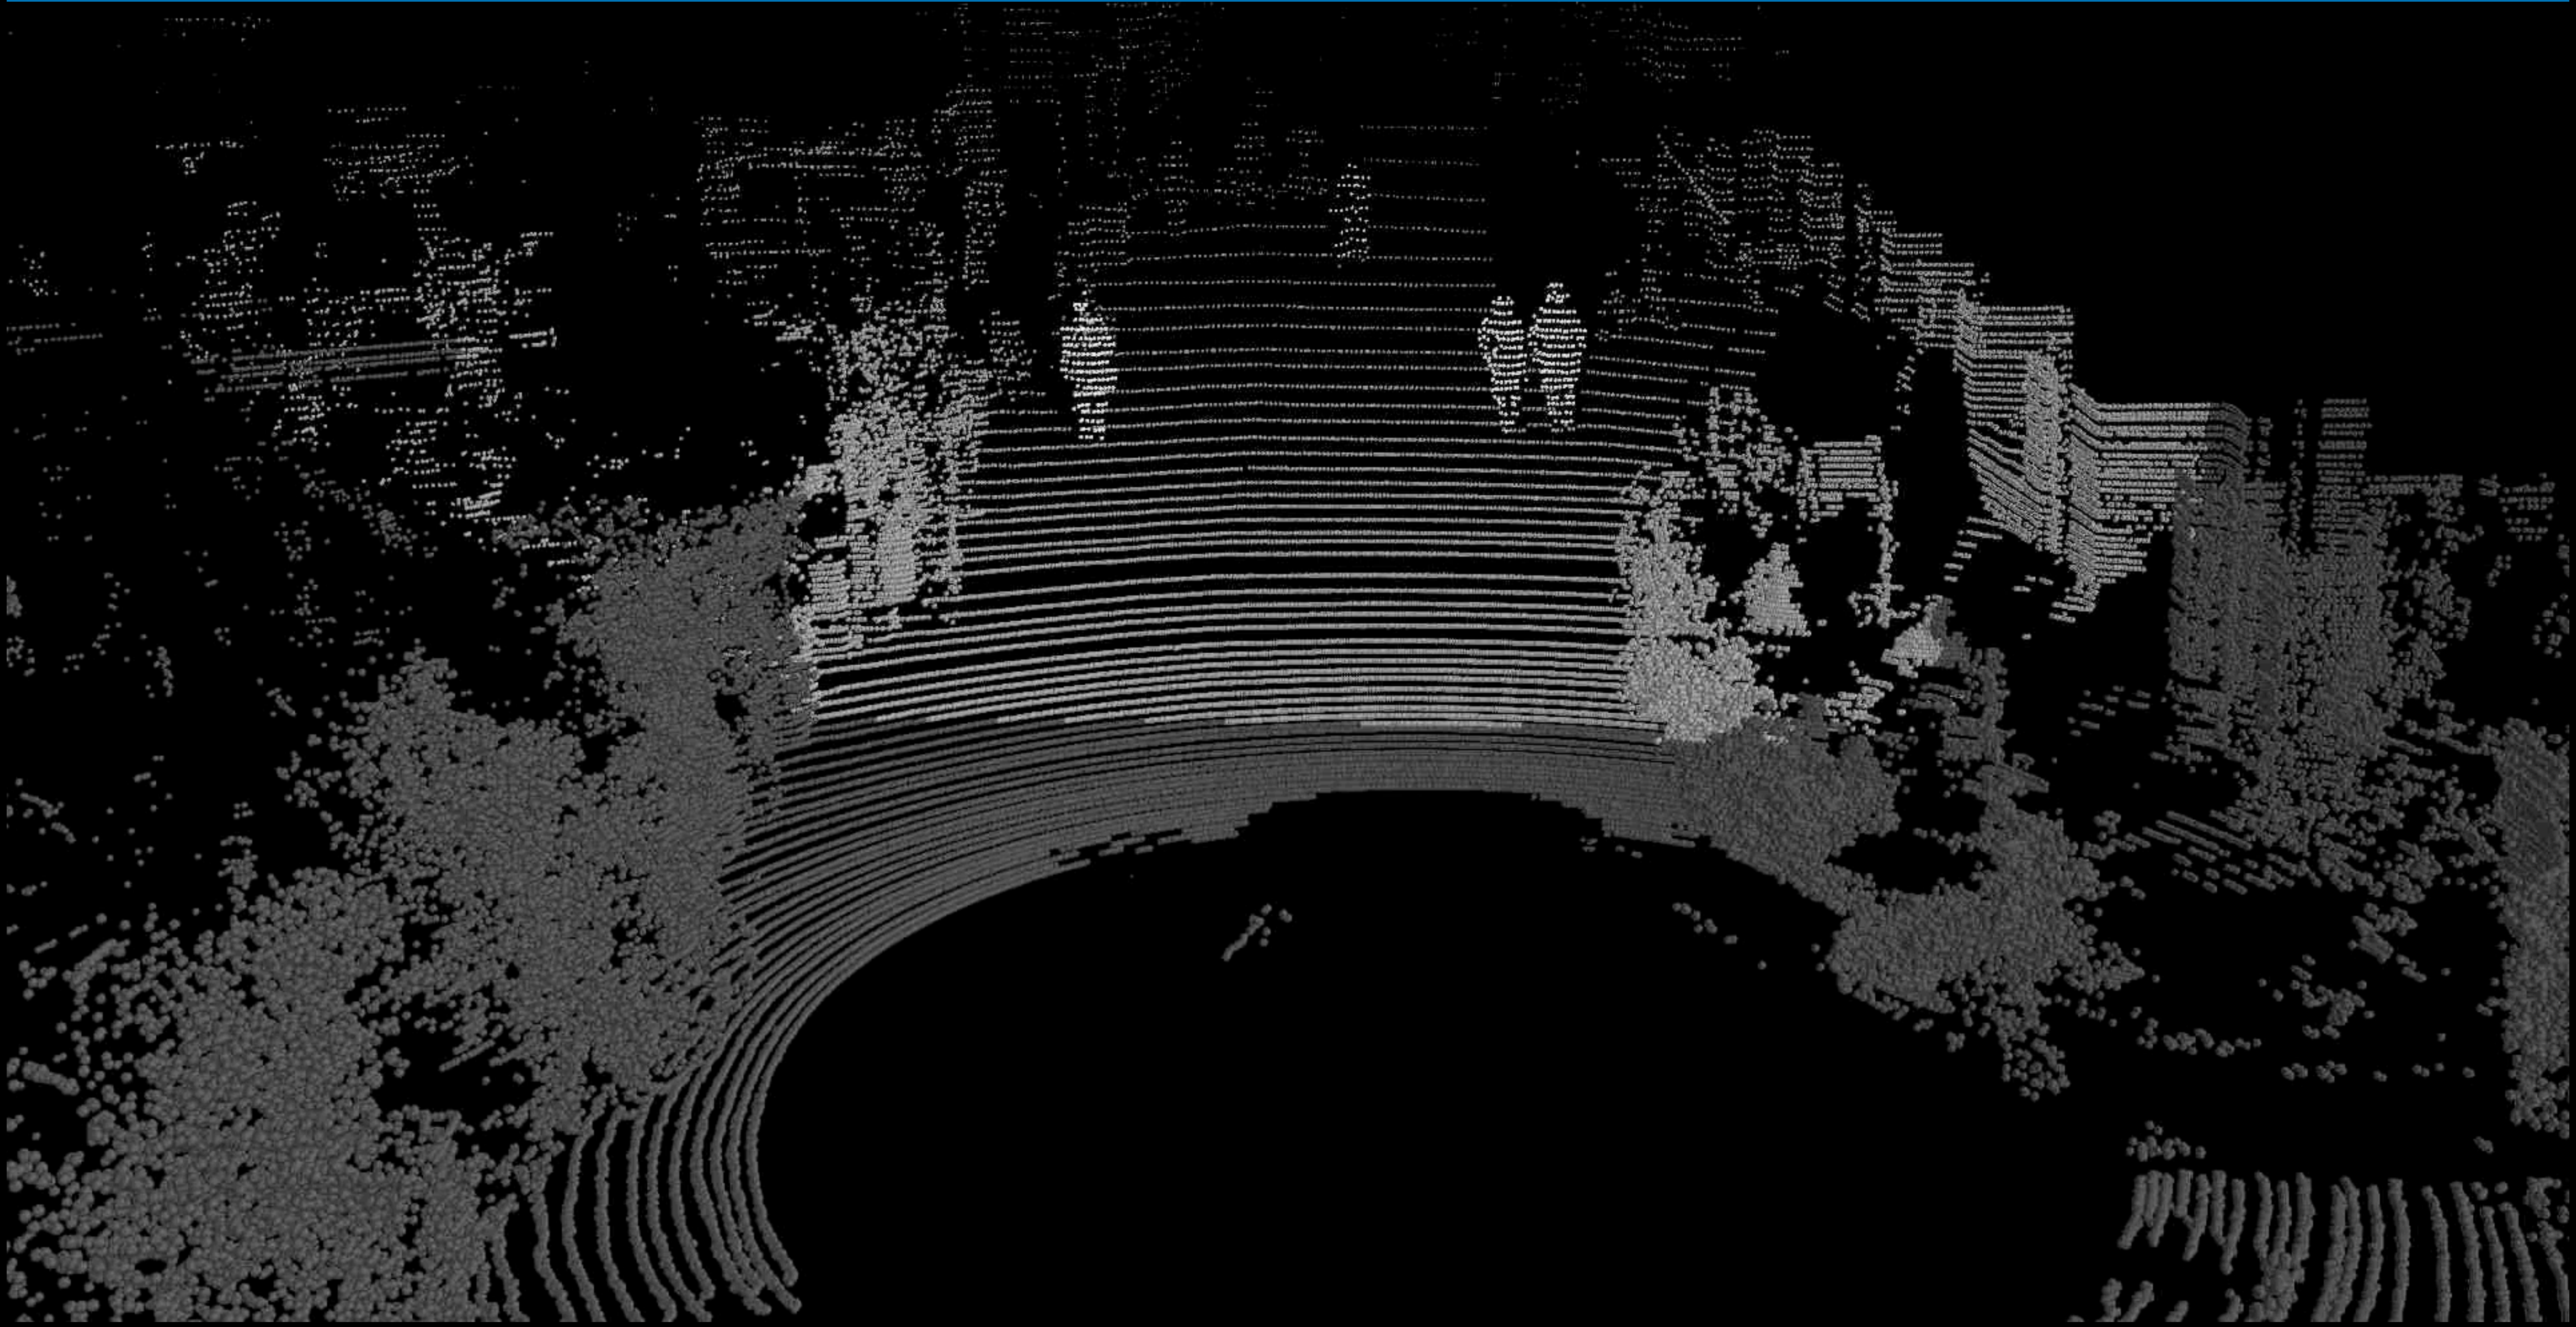

*We focus on a limited field of view in this work.

Object-level Perception

However, it is often hard to set the right hyper-parameters for clustering. For example, Euclidean Clustering with a large distance threshold tends to under-segments pedestrians.

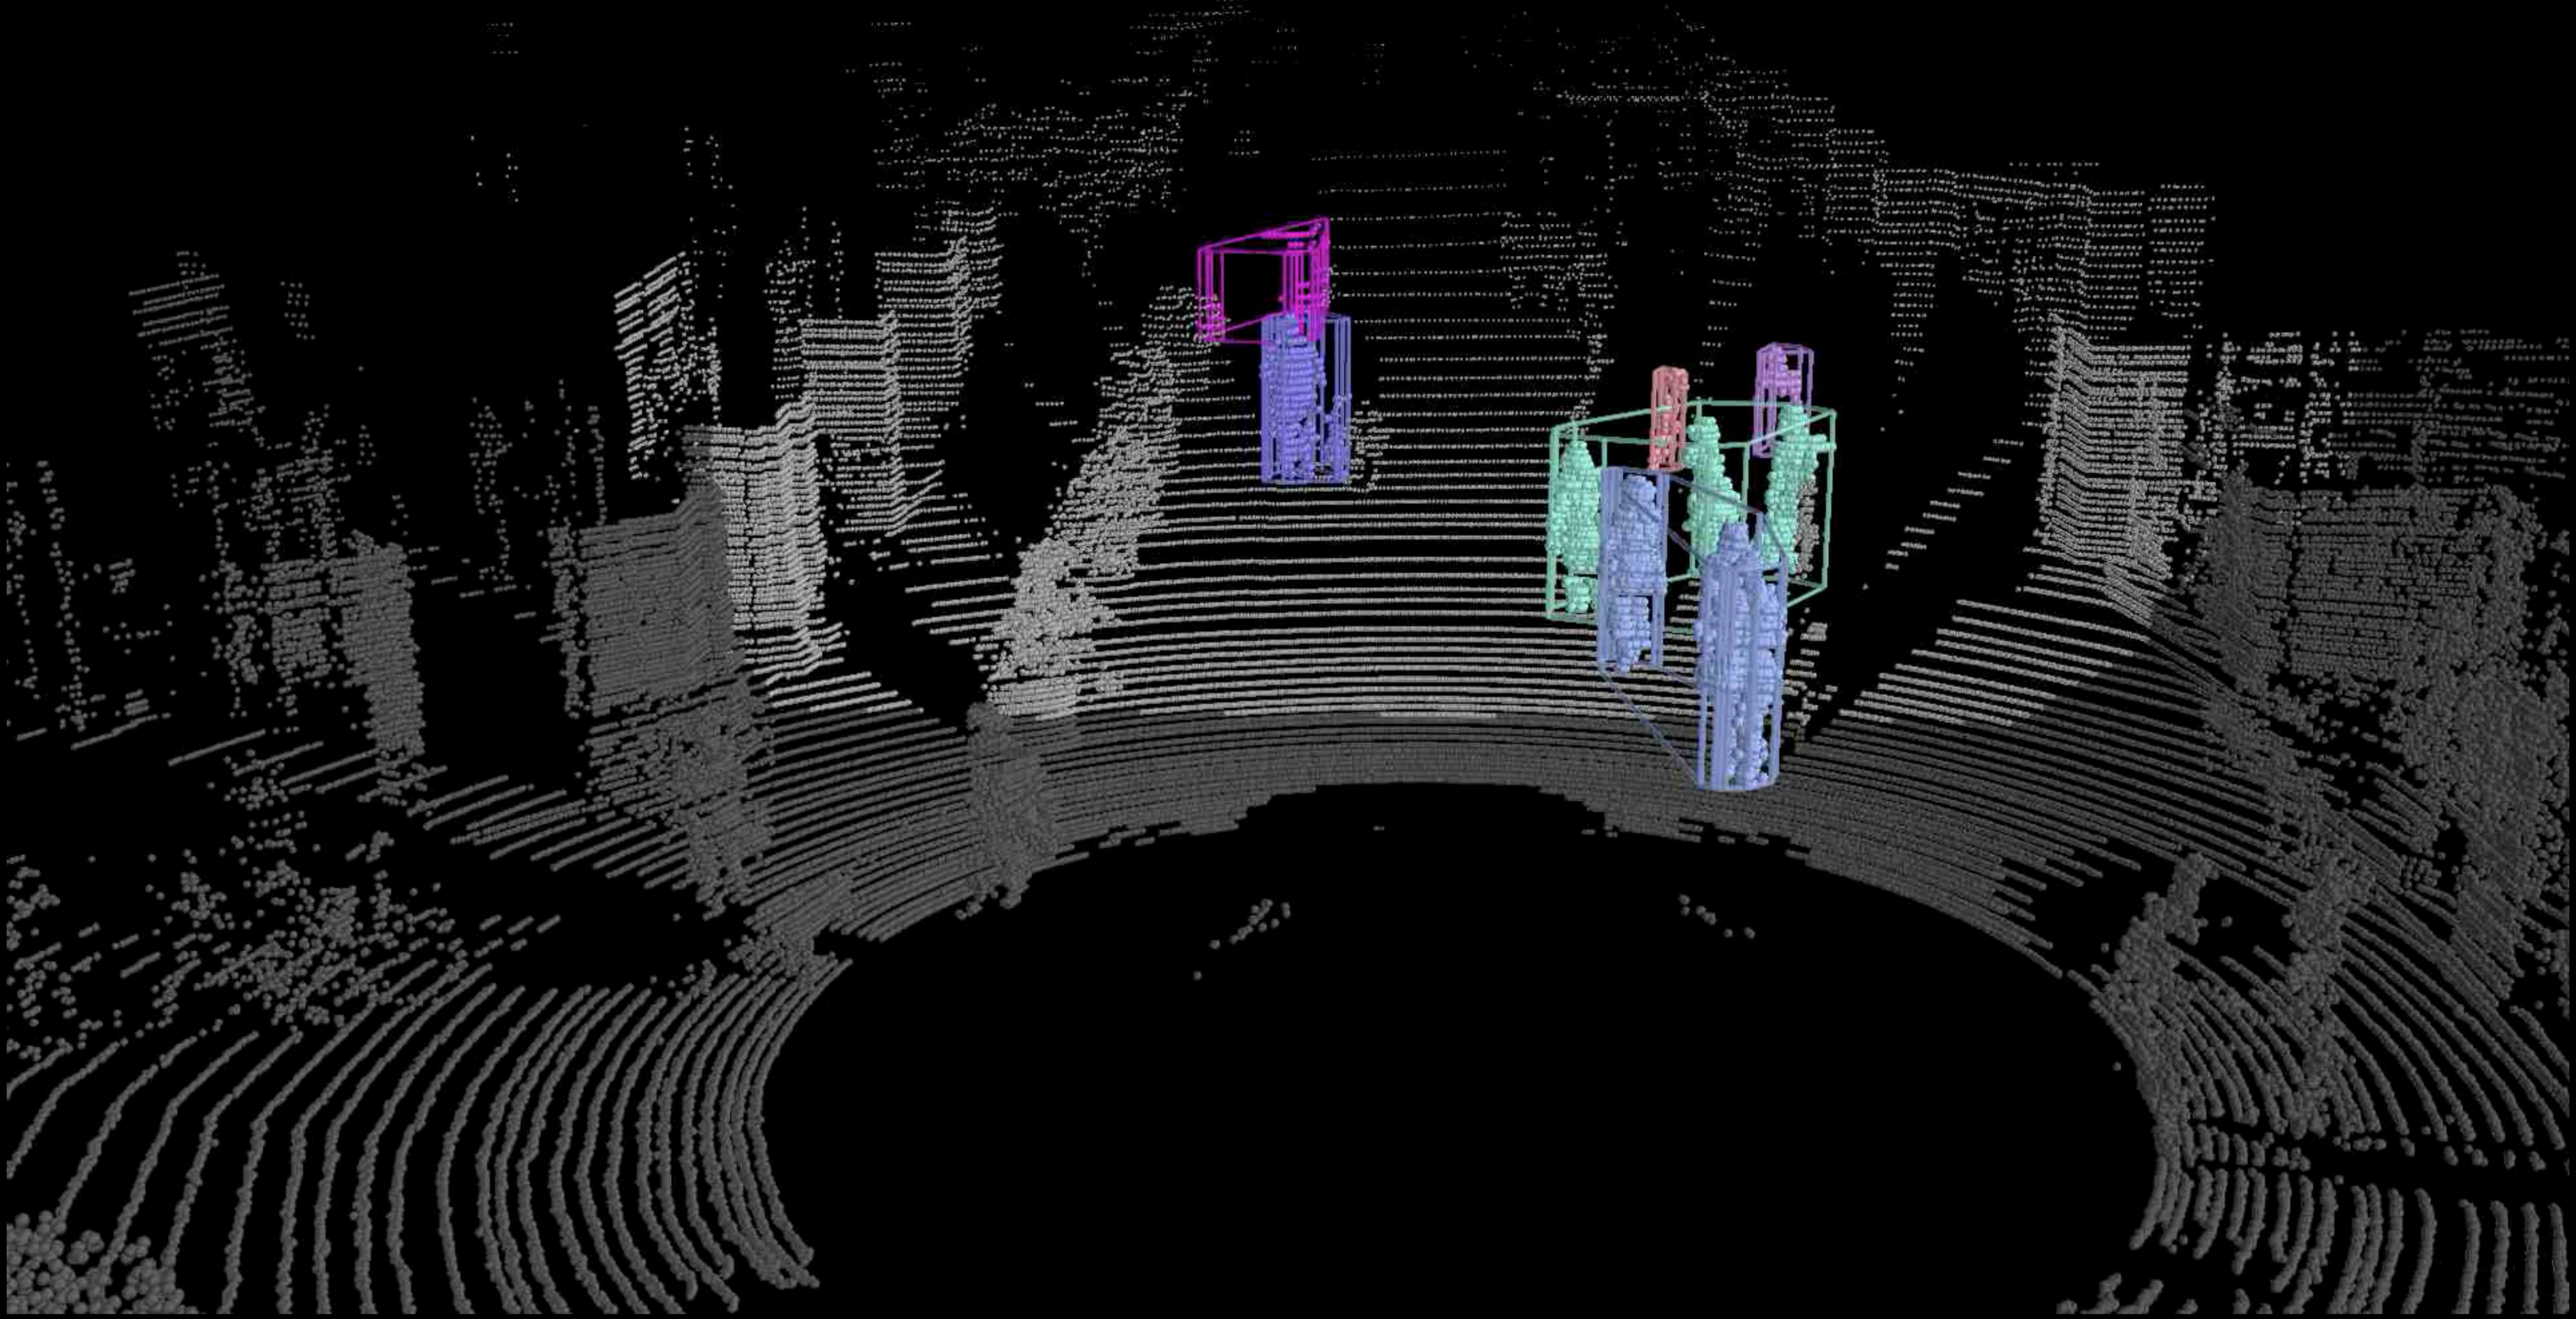

*Colors flicker because the algorithm does not track objects across time.

Object-level Perception

And Euclidean Clustering with a small distance threshold tends to over-segments vehicles.

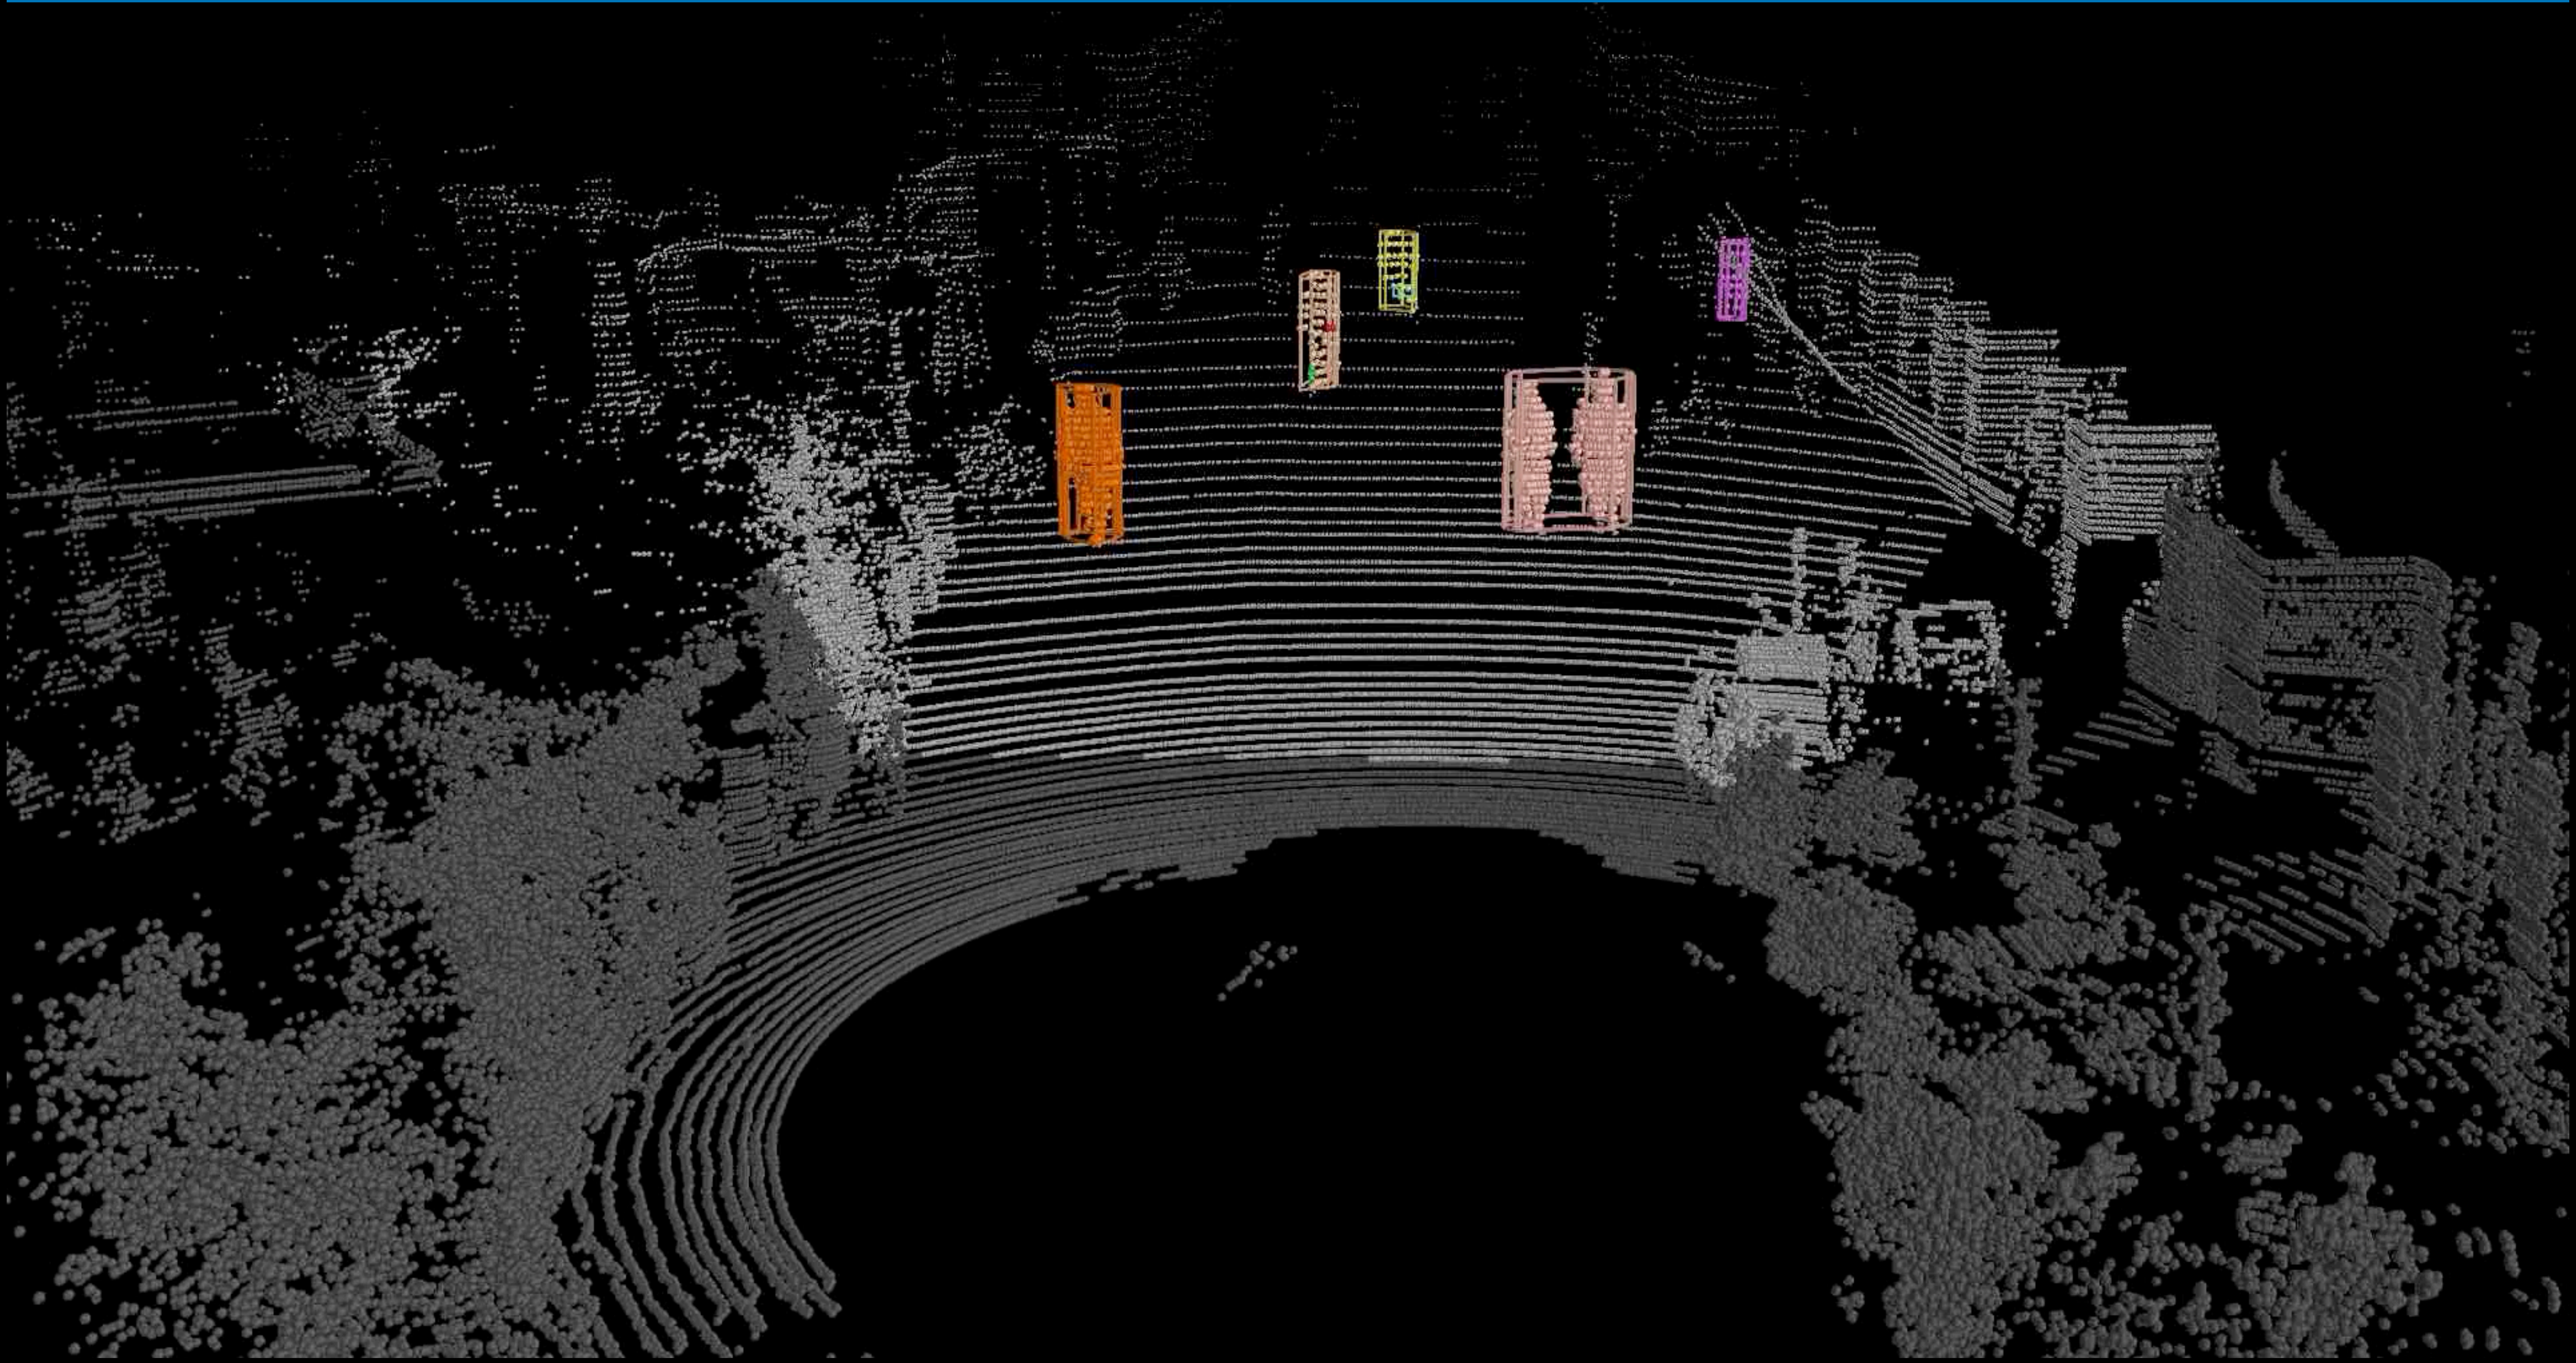

*Colors flicker because the algorithm does not track objects across time.

No One-Fits-All Solution

The best distance threshold often varies from scenario to scenario.

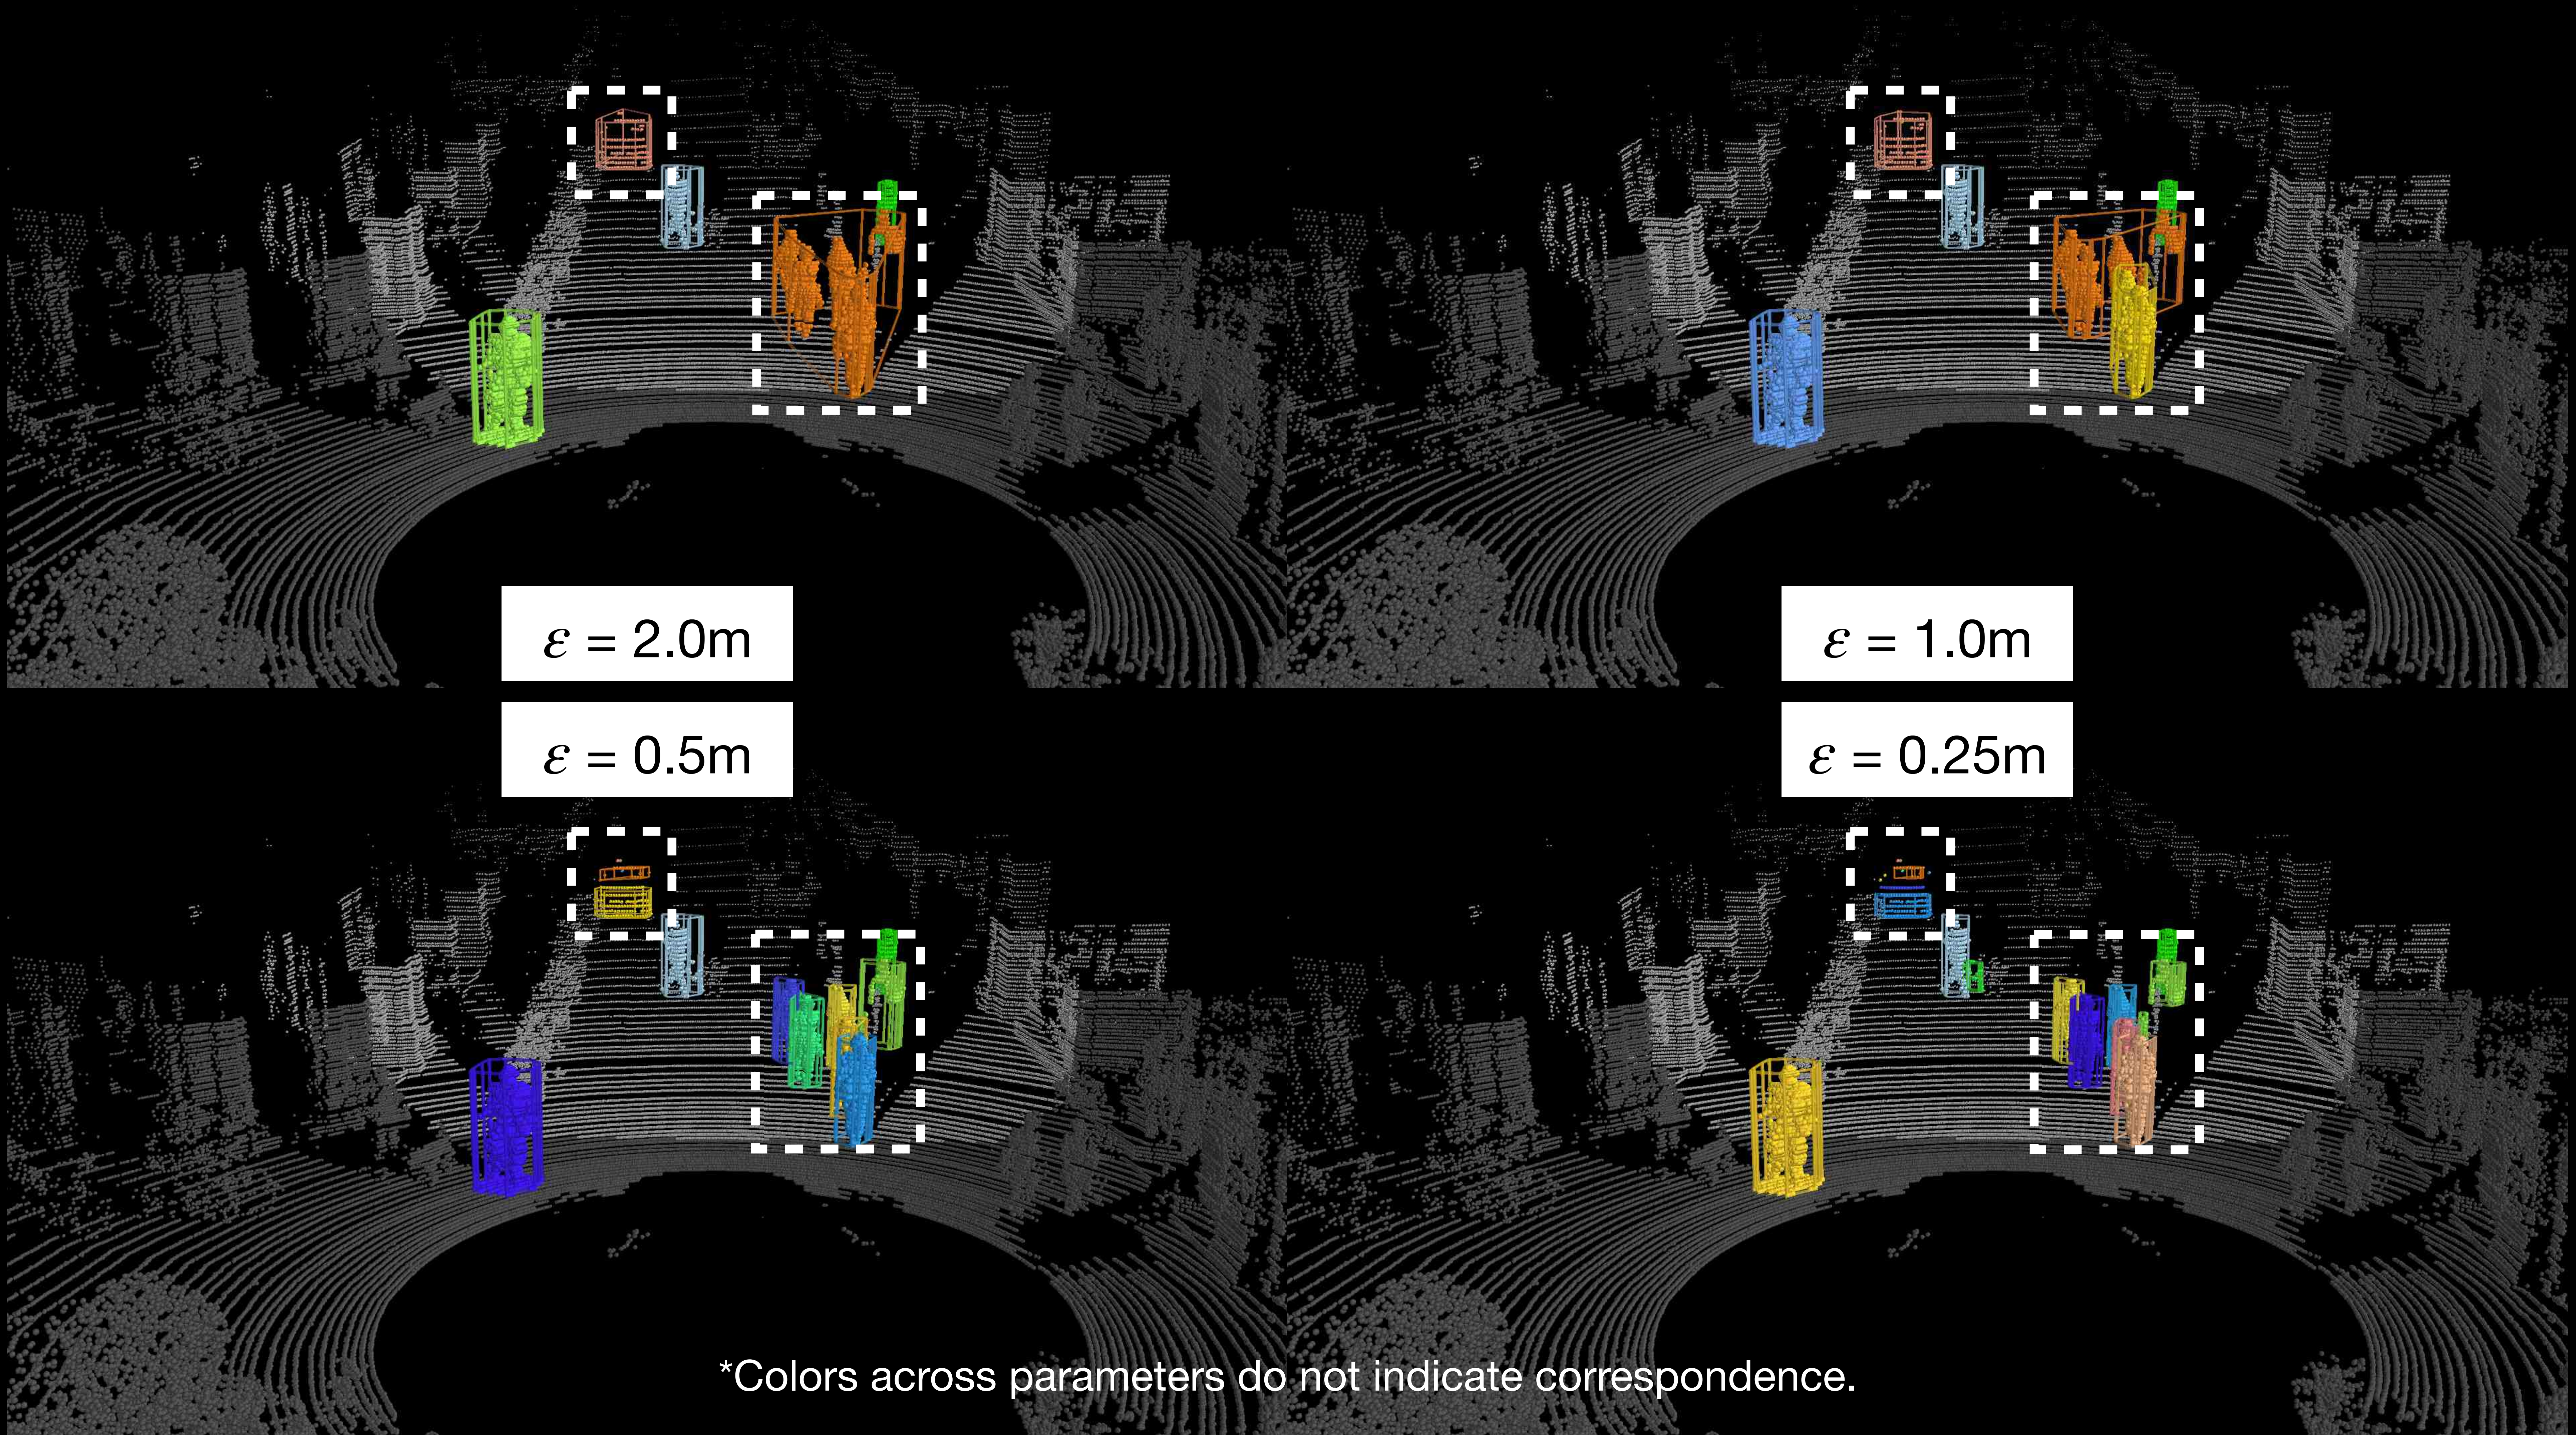

*Colors across parameters do not indicate correspondence.

A Hierarchical Perspective

Segmentations with different thresholds form a hierarchy, where nodes represent segments.

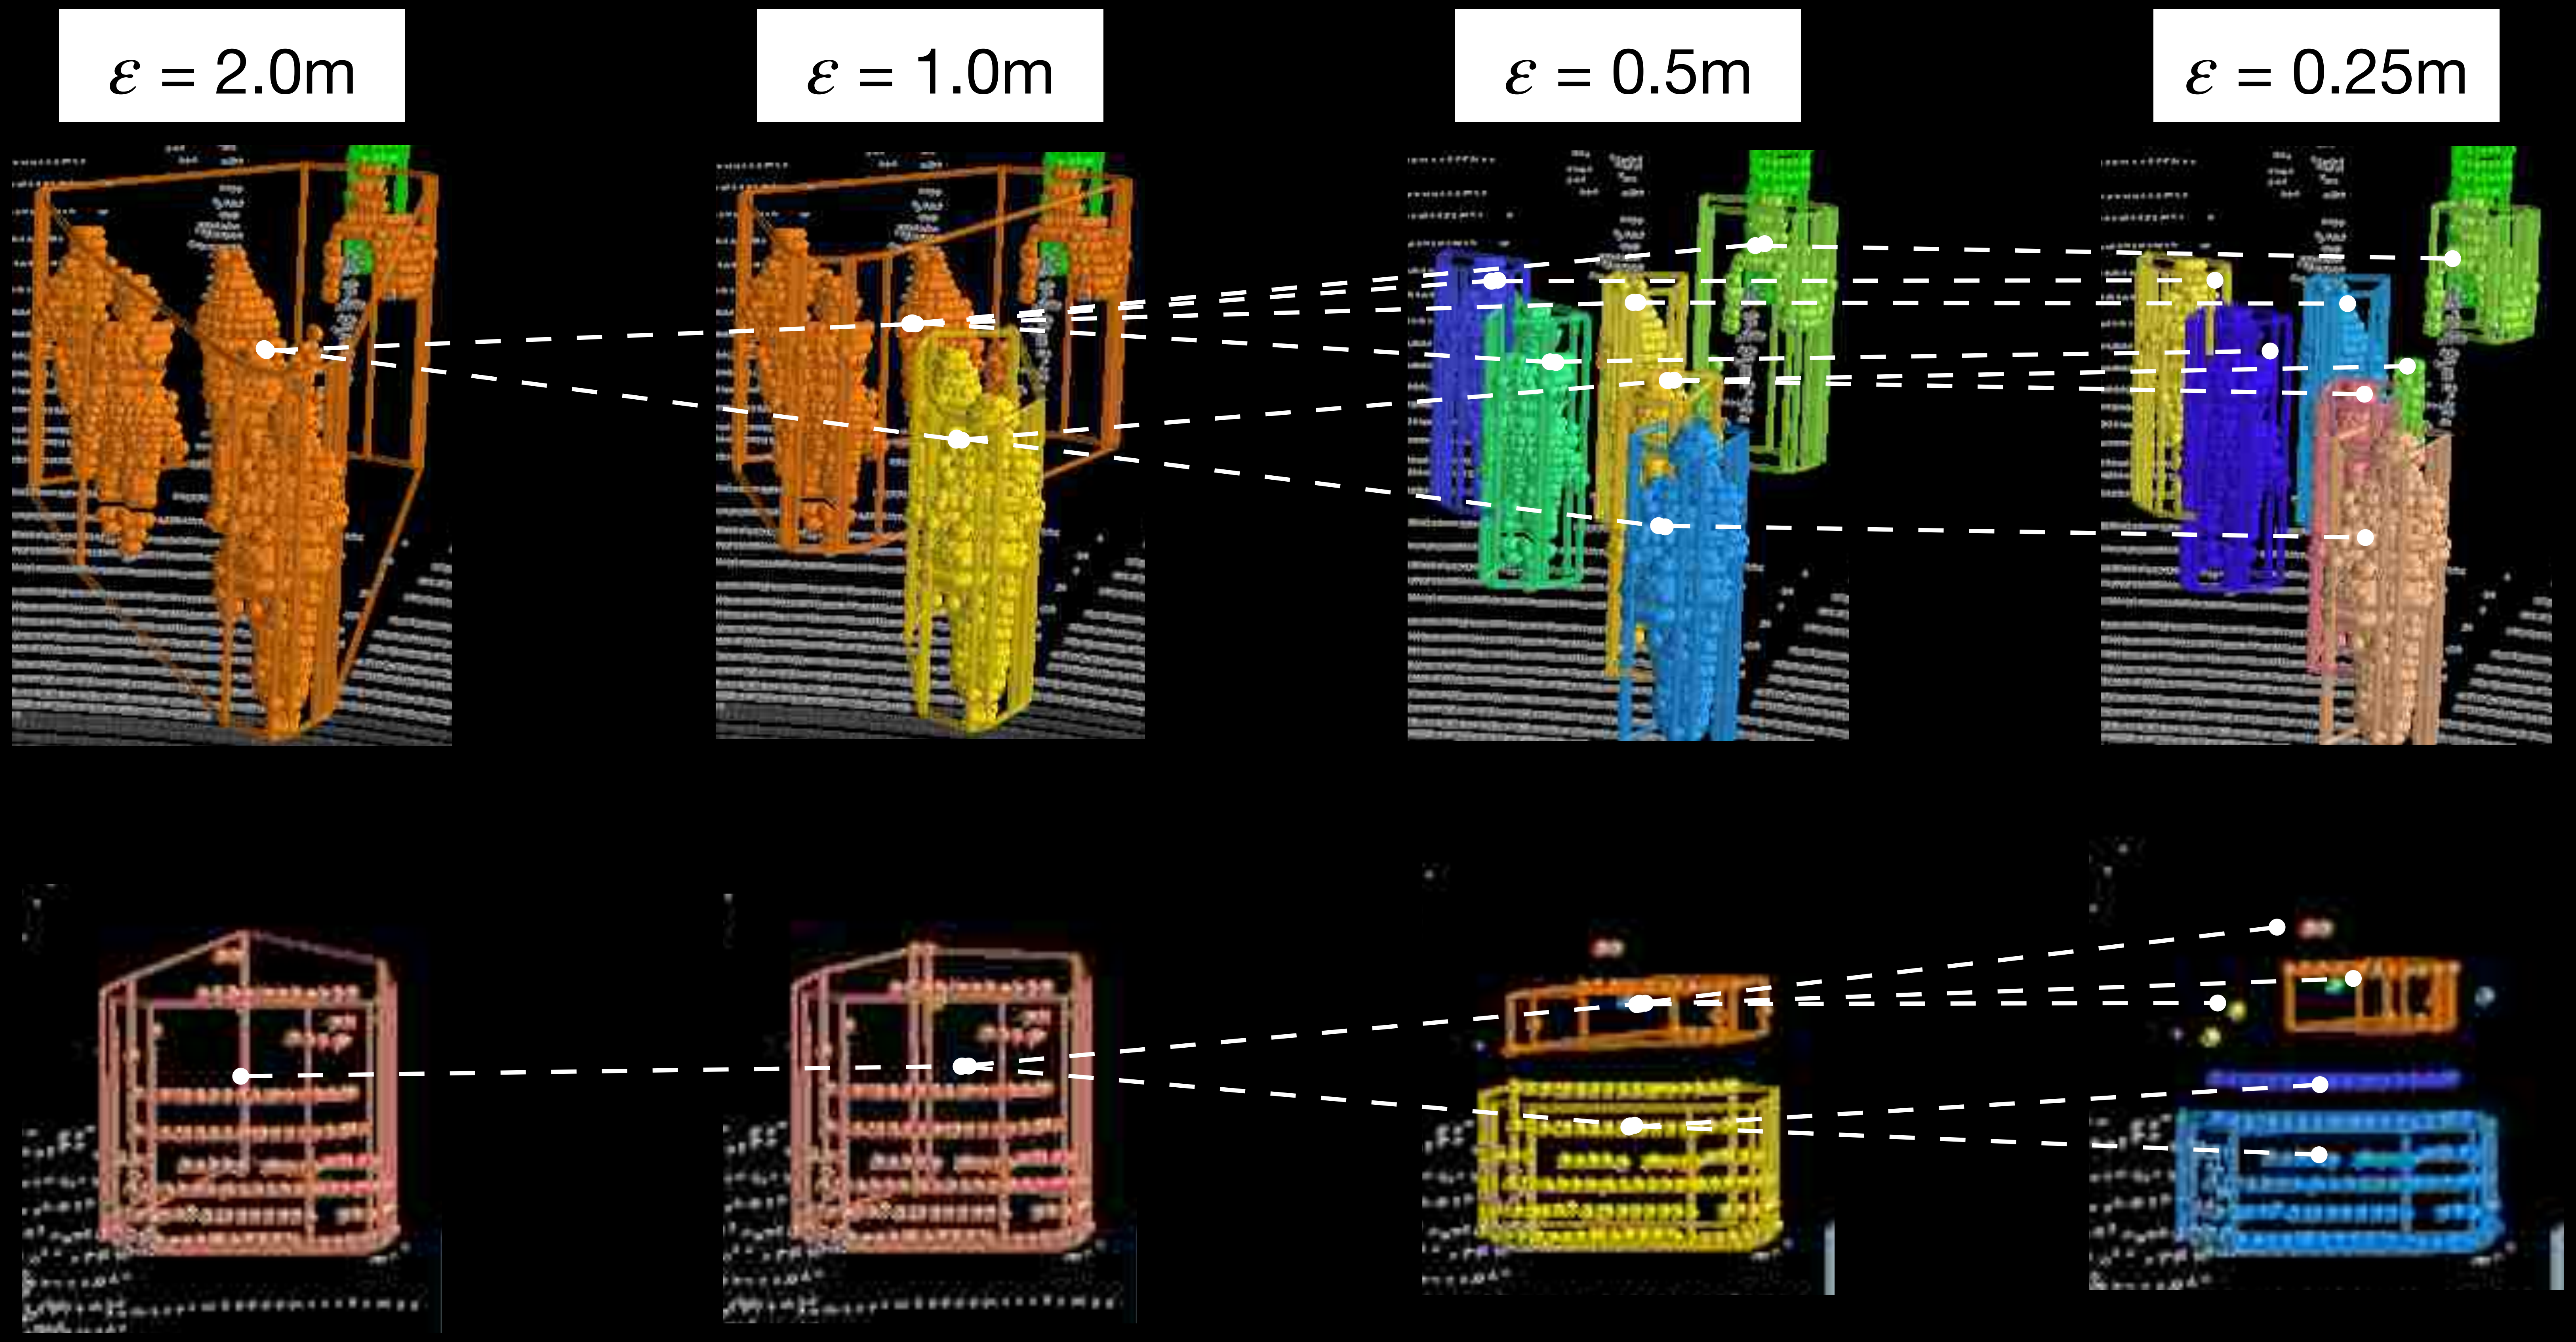

Learning Objectness Models

We learn a model to predict an objectness score for each segment in the hierarchy.

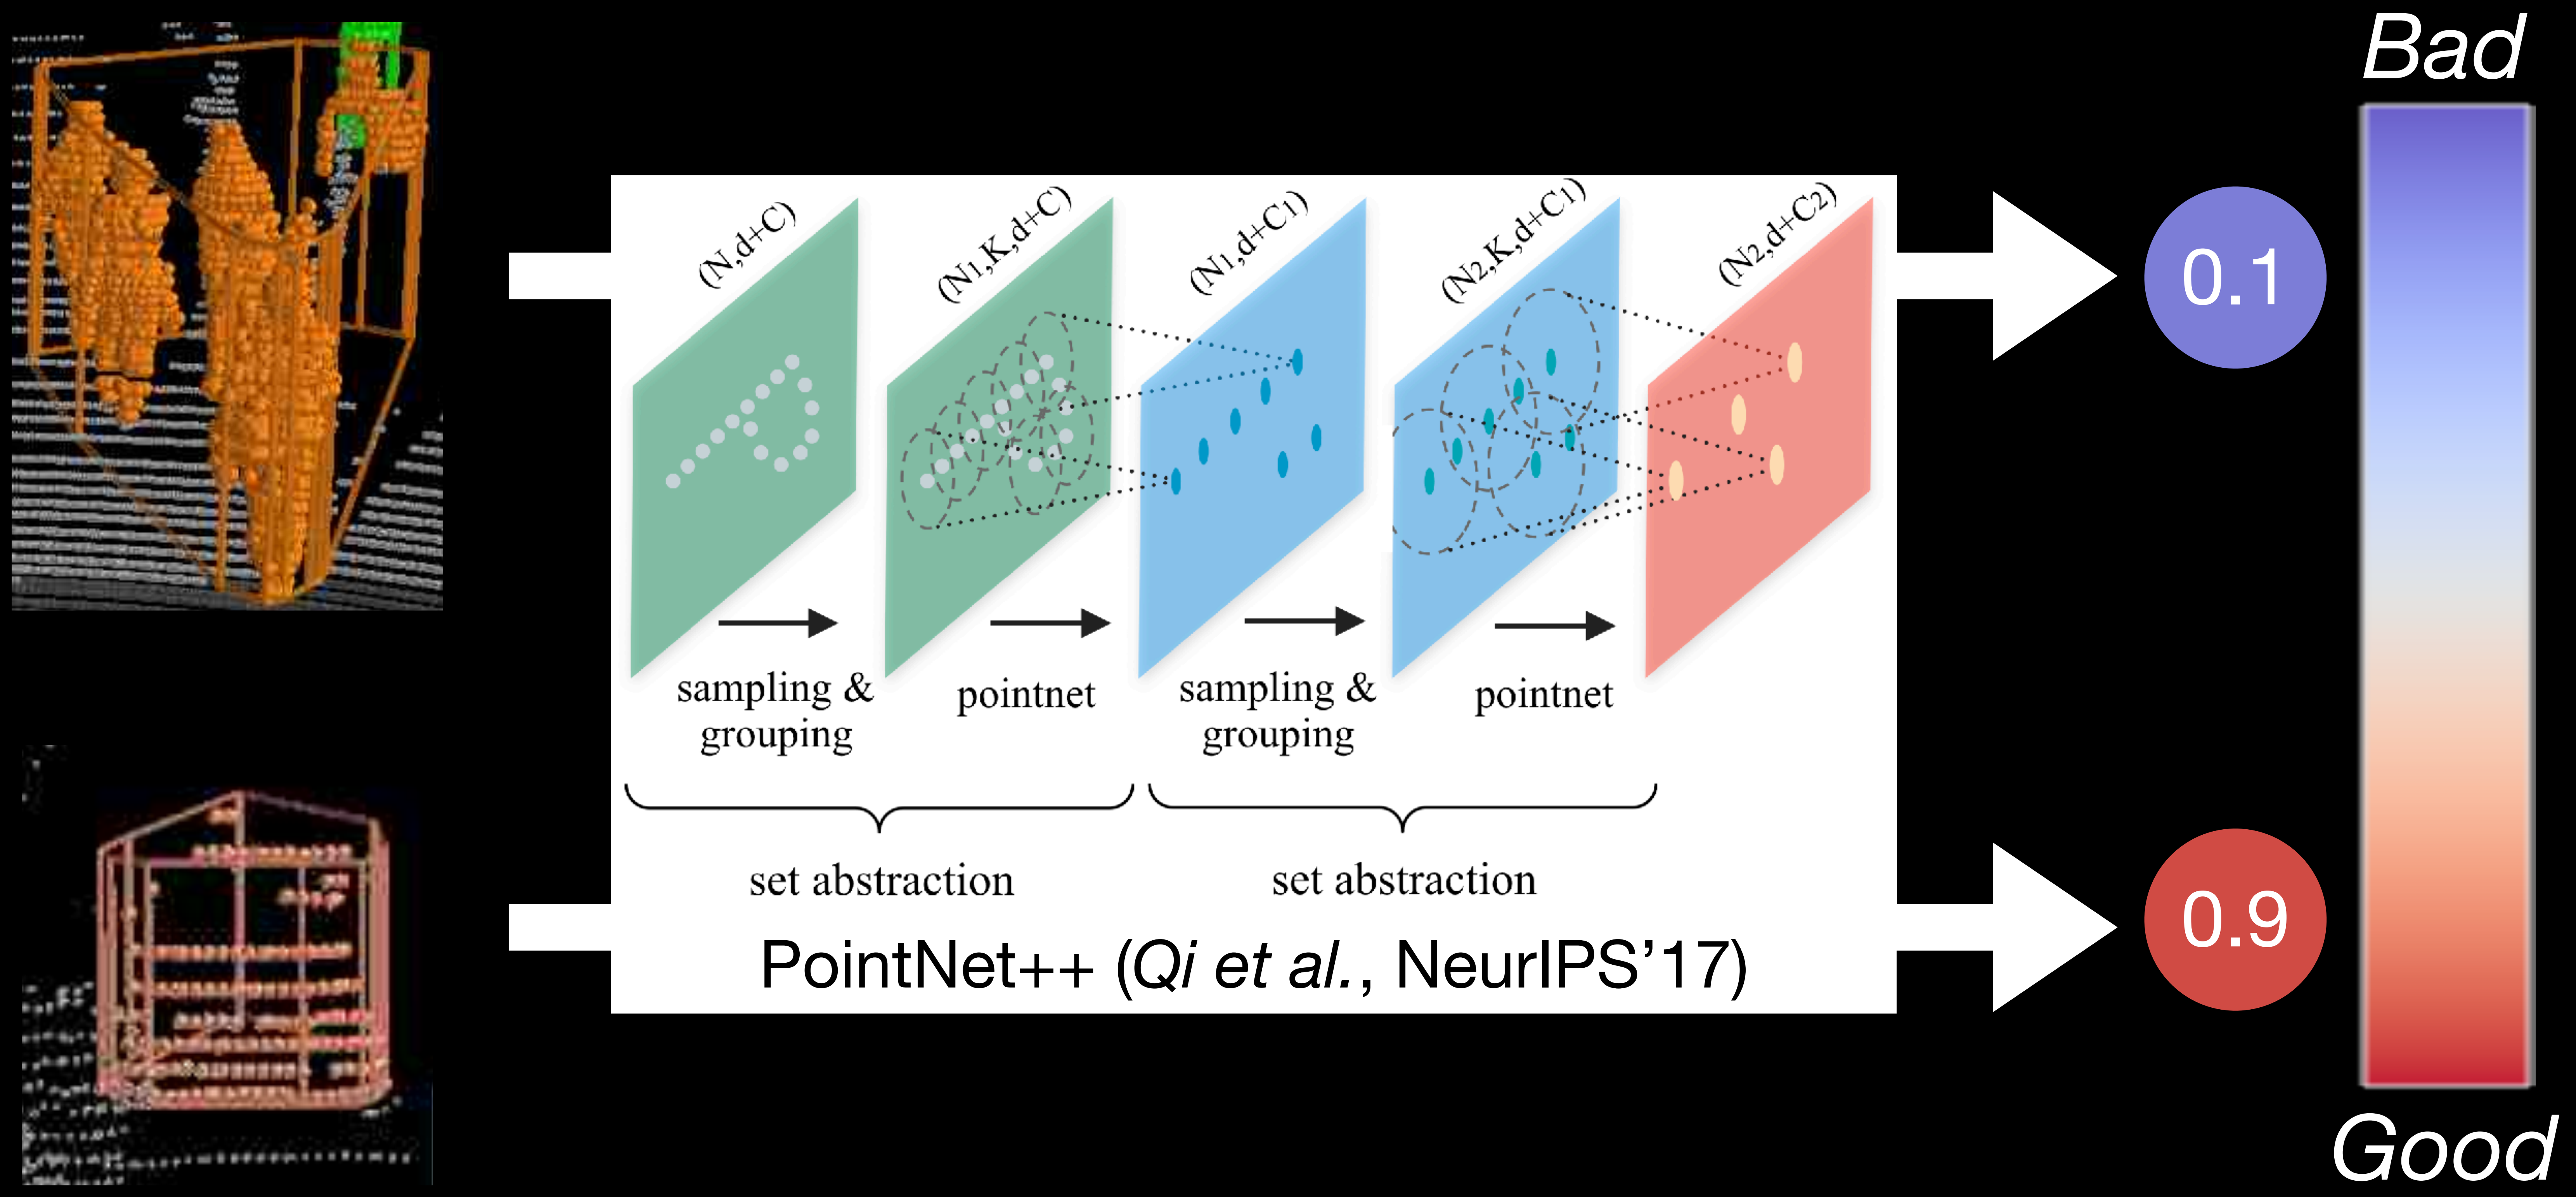

Objectness

How well a segment overlaps with ground truths.

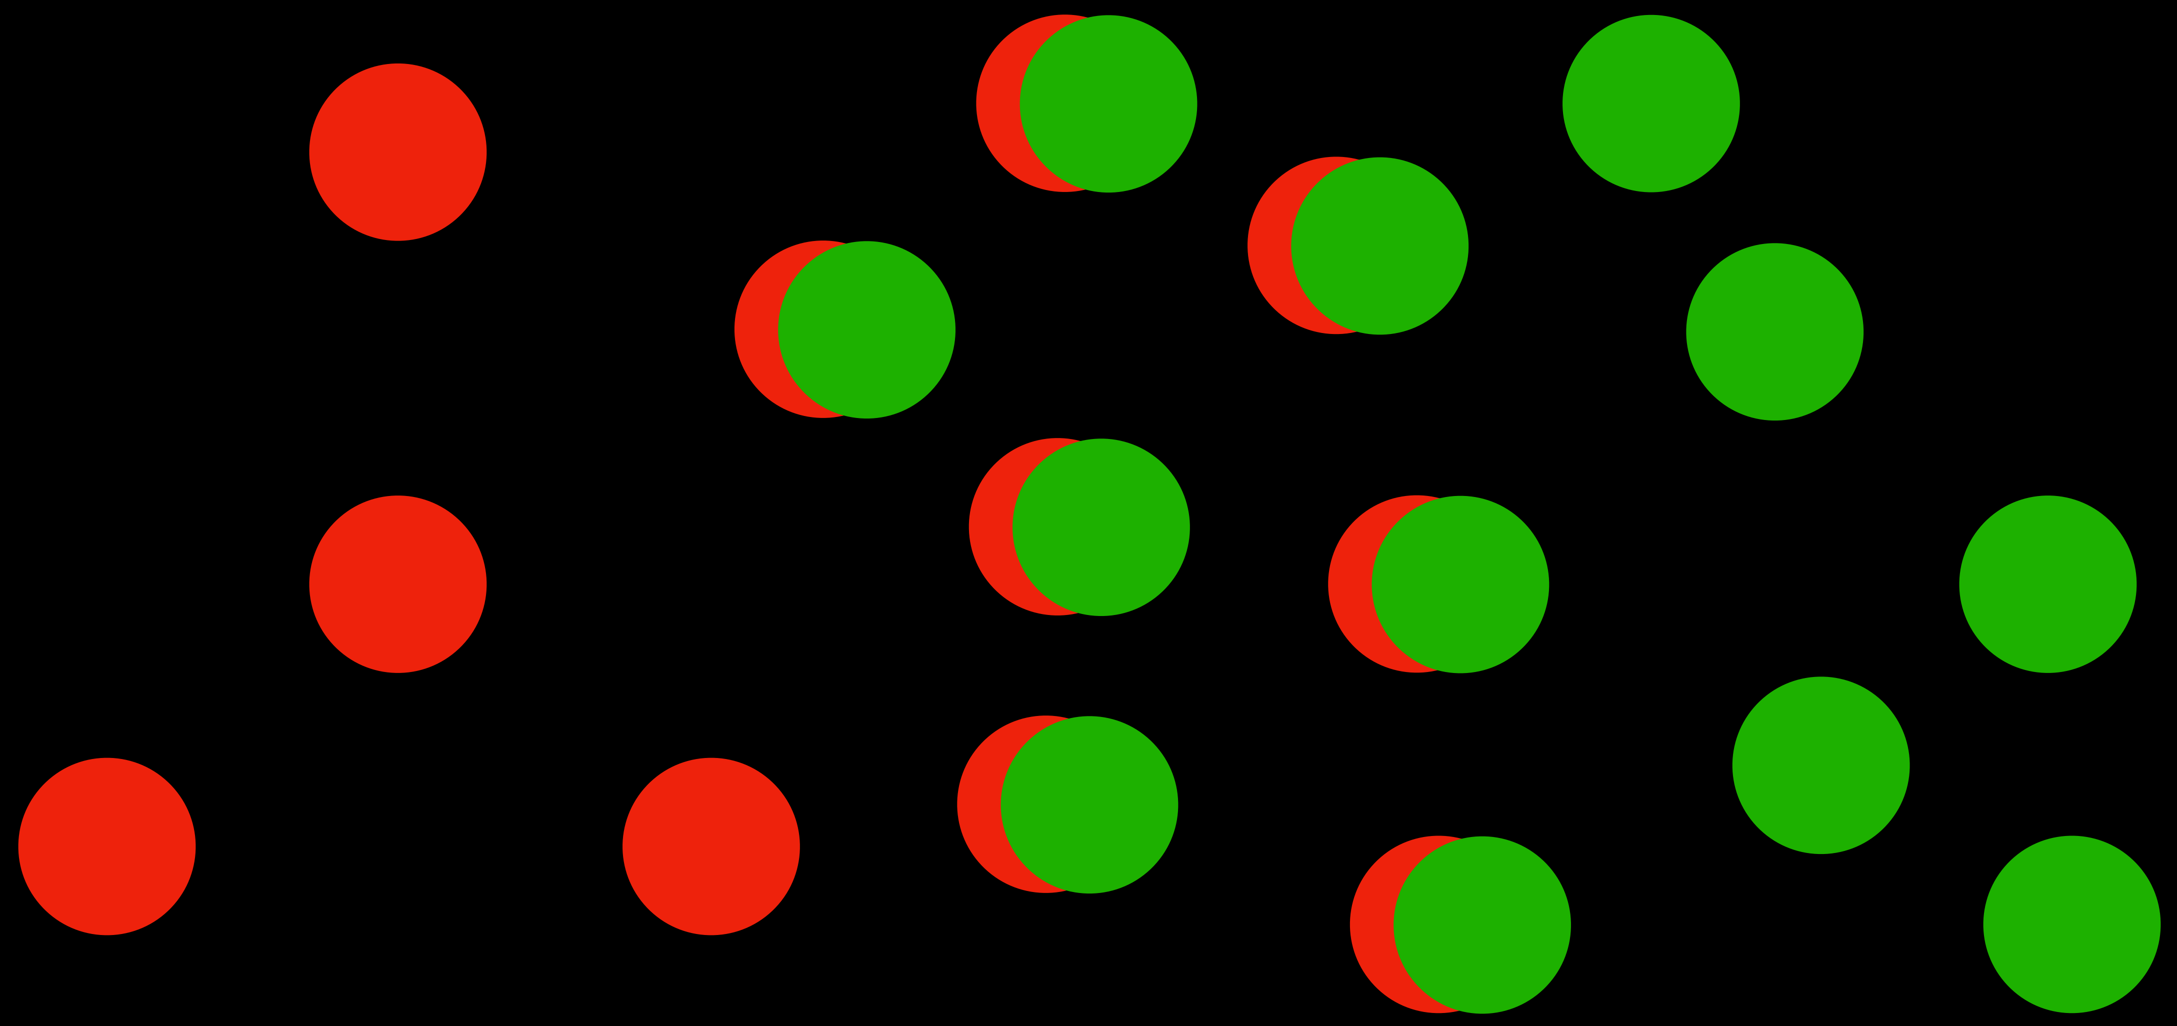

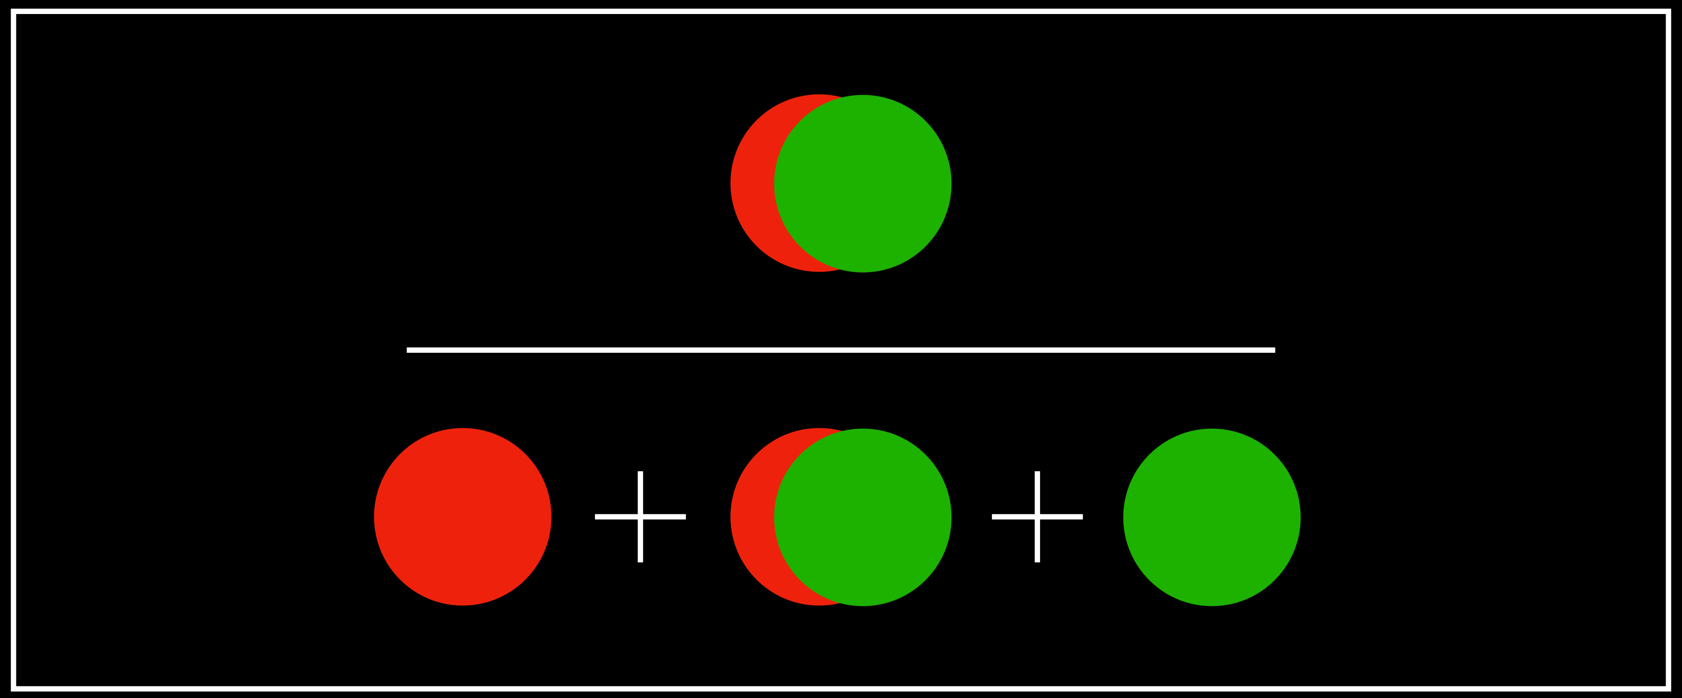

Searching for Optimality

Given a hierarchy of segments with scores, we search for the optimal segmentation.

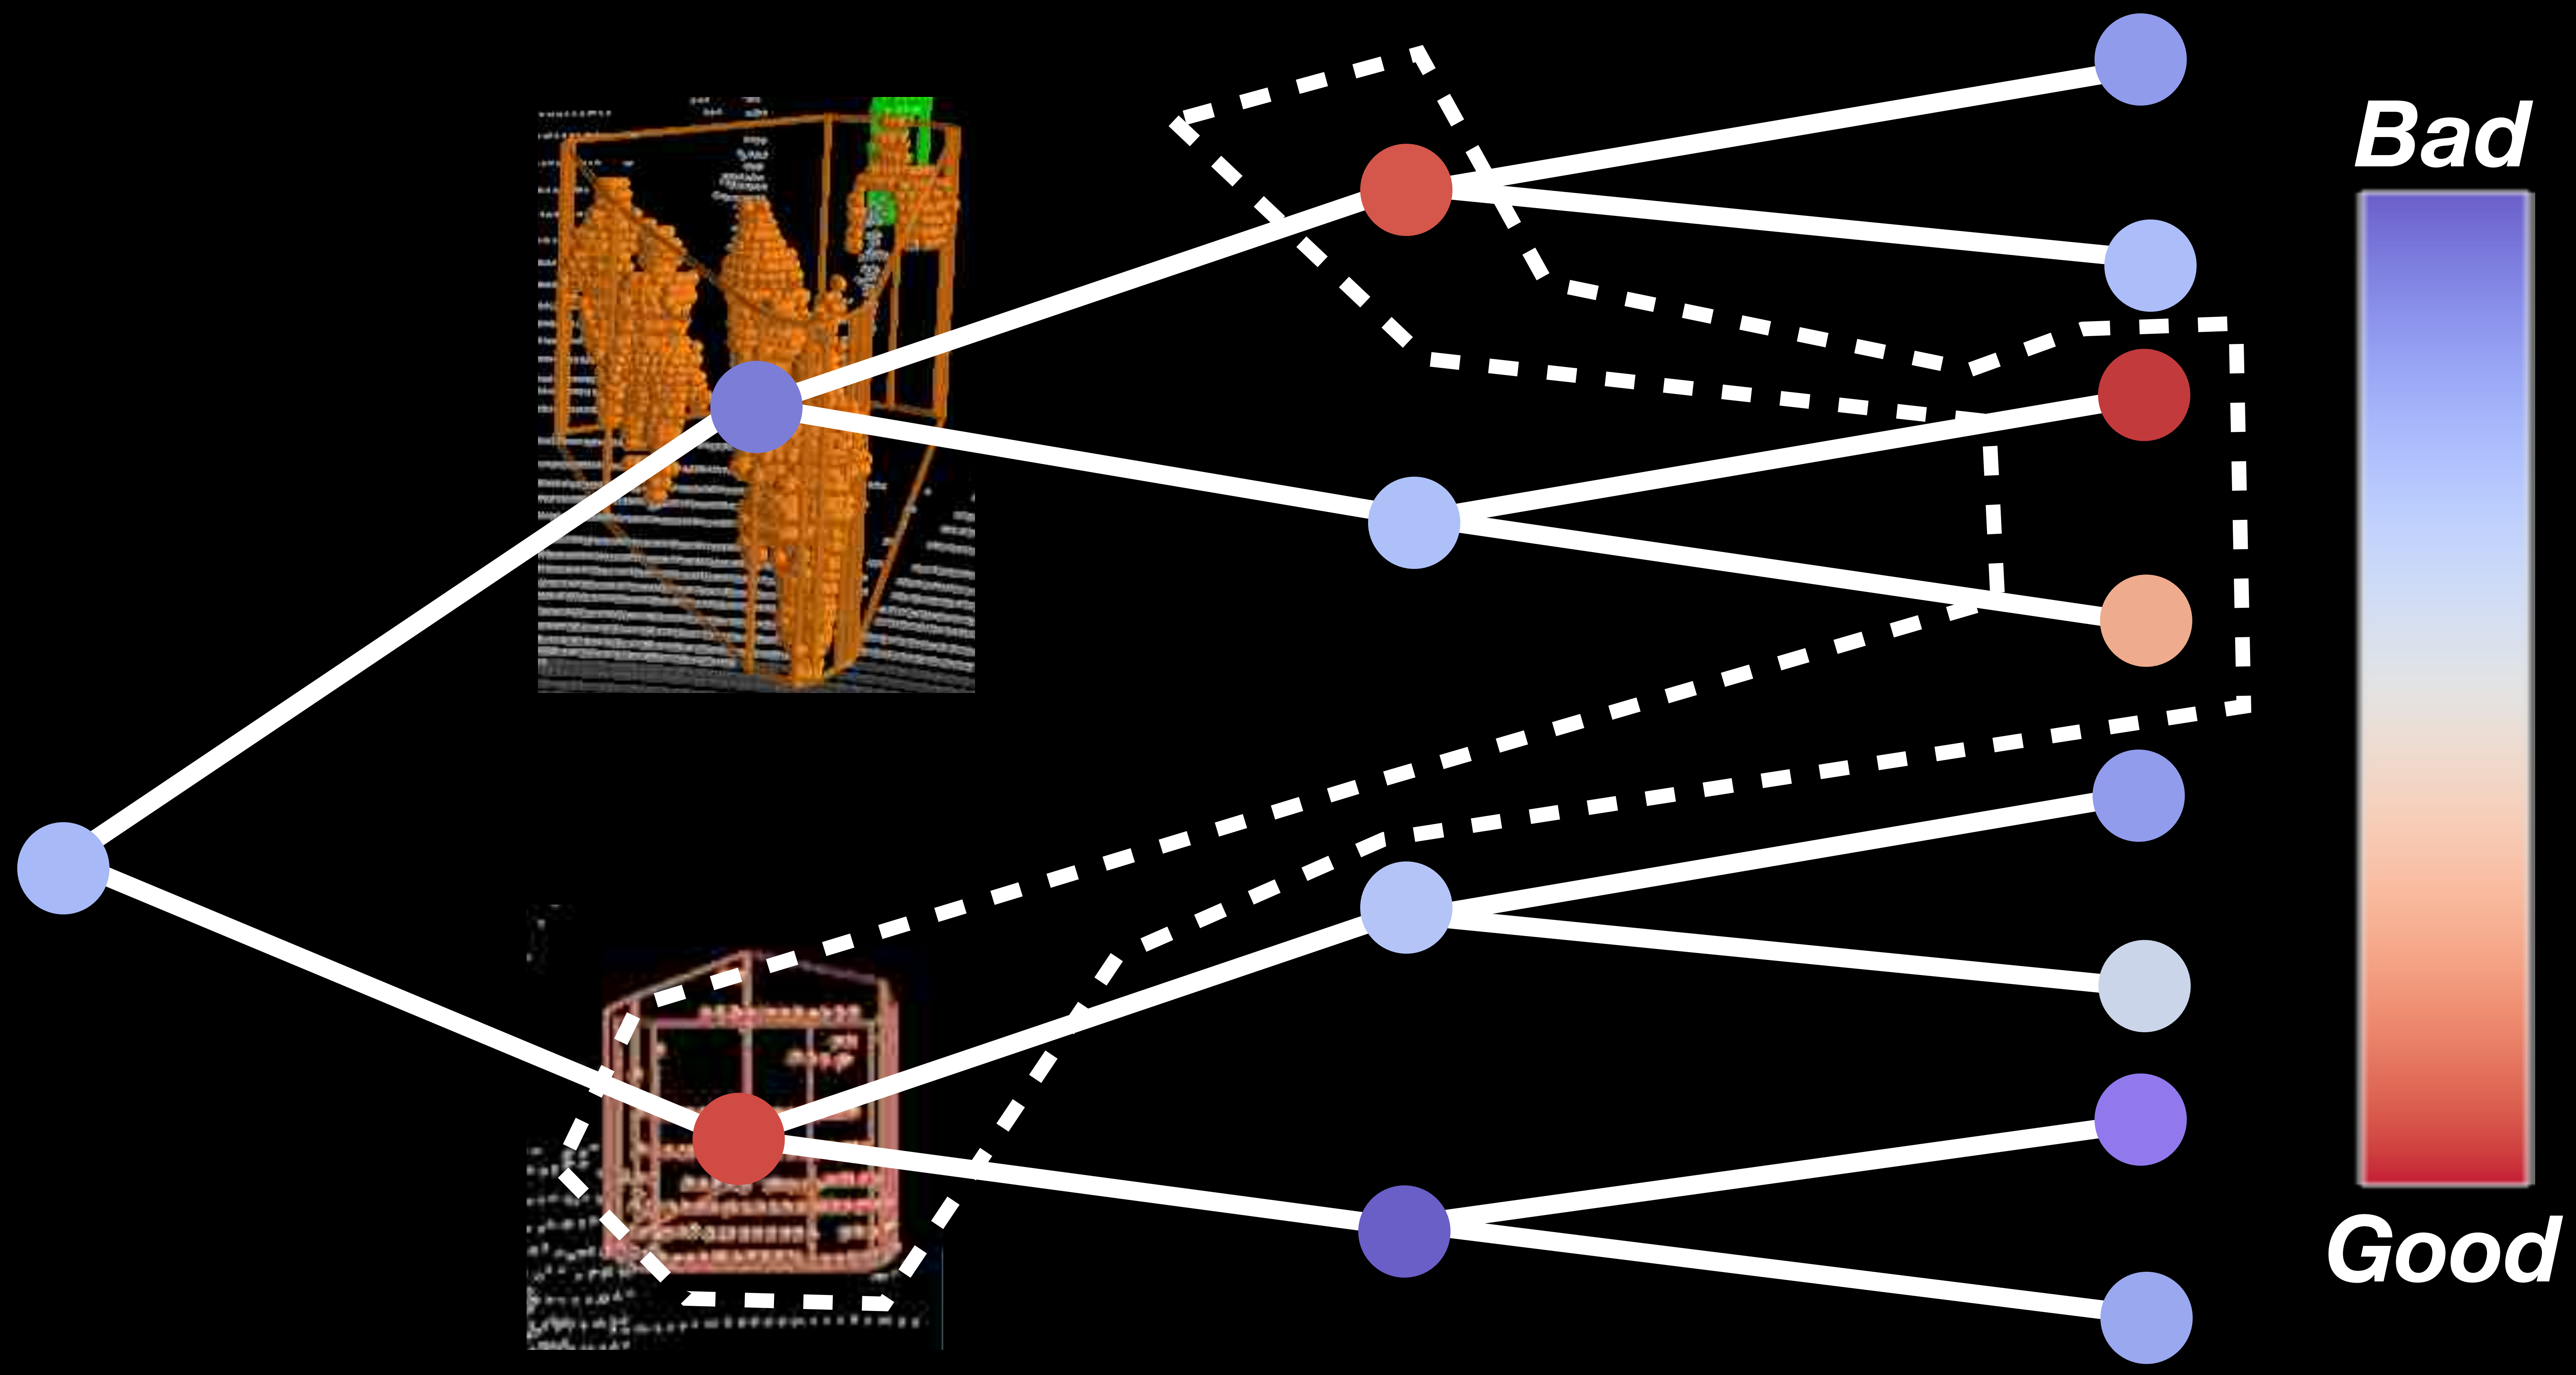

Optimal Worst-case Segmentation

We propose an efficient algorithm that produces optimal segmentation under this definition.

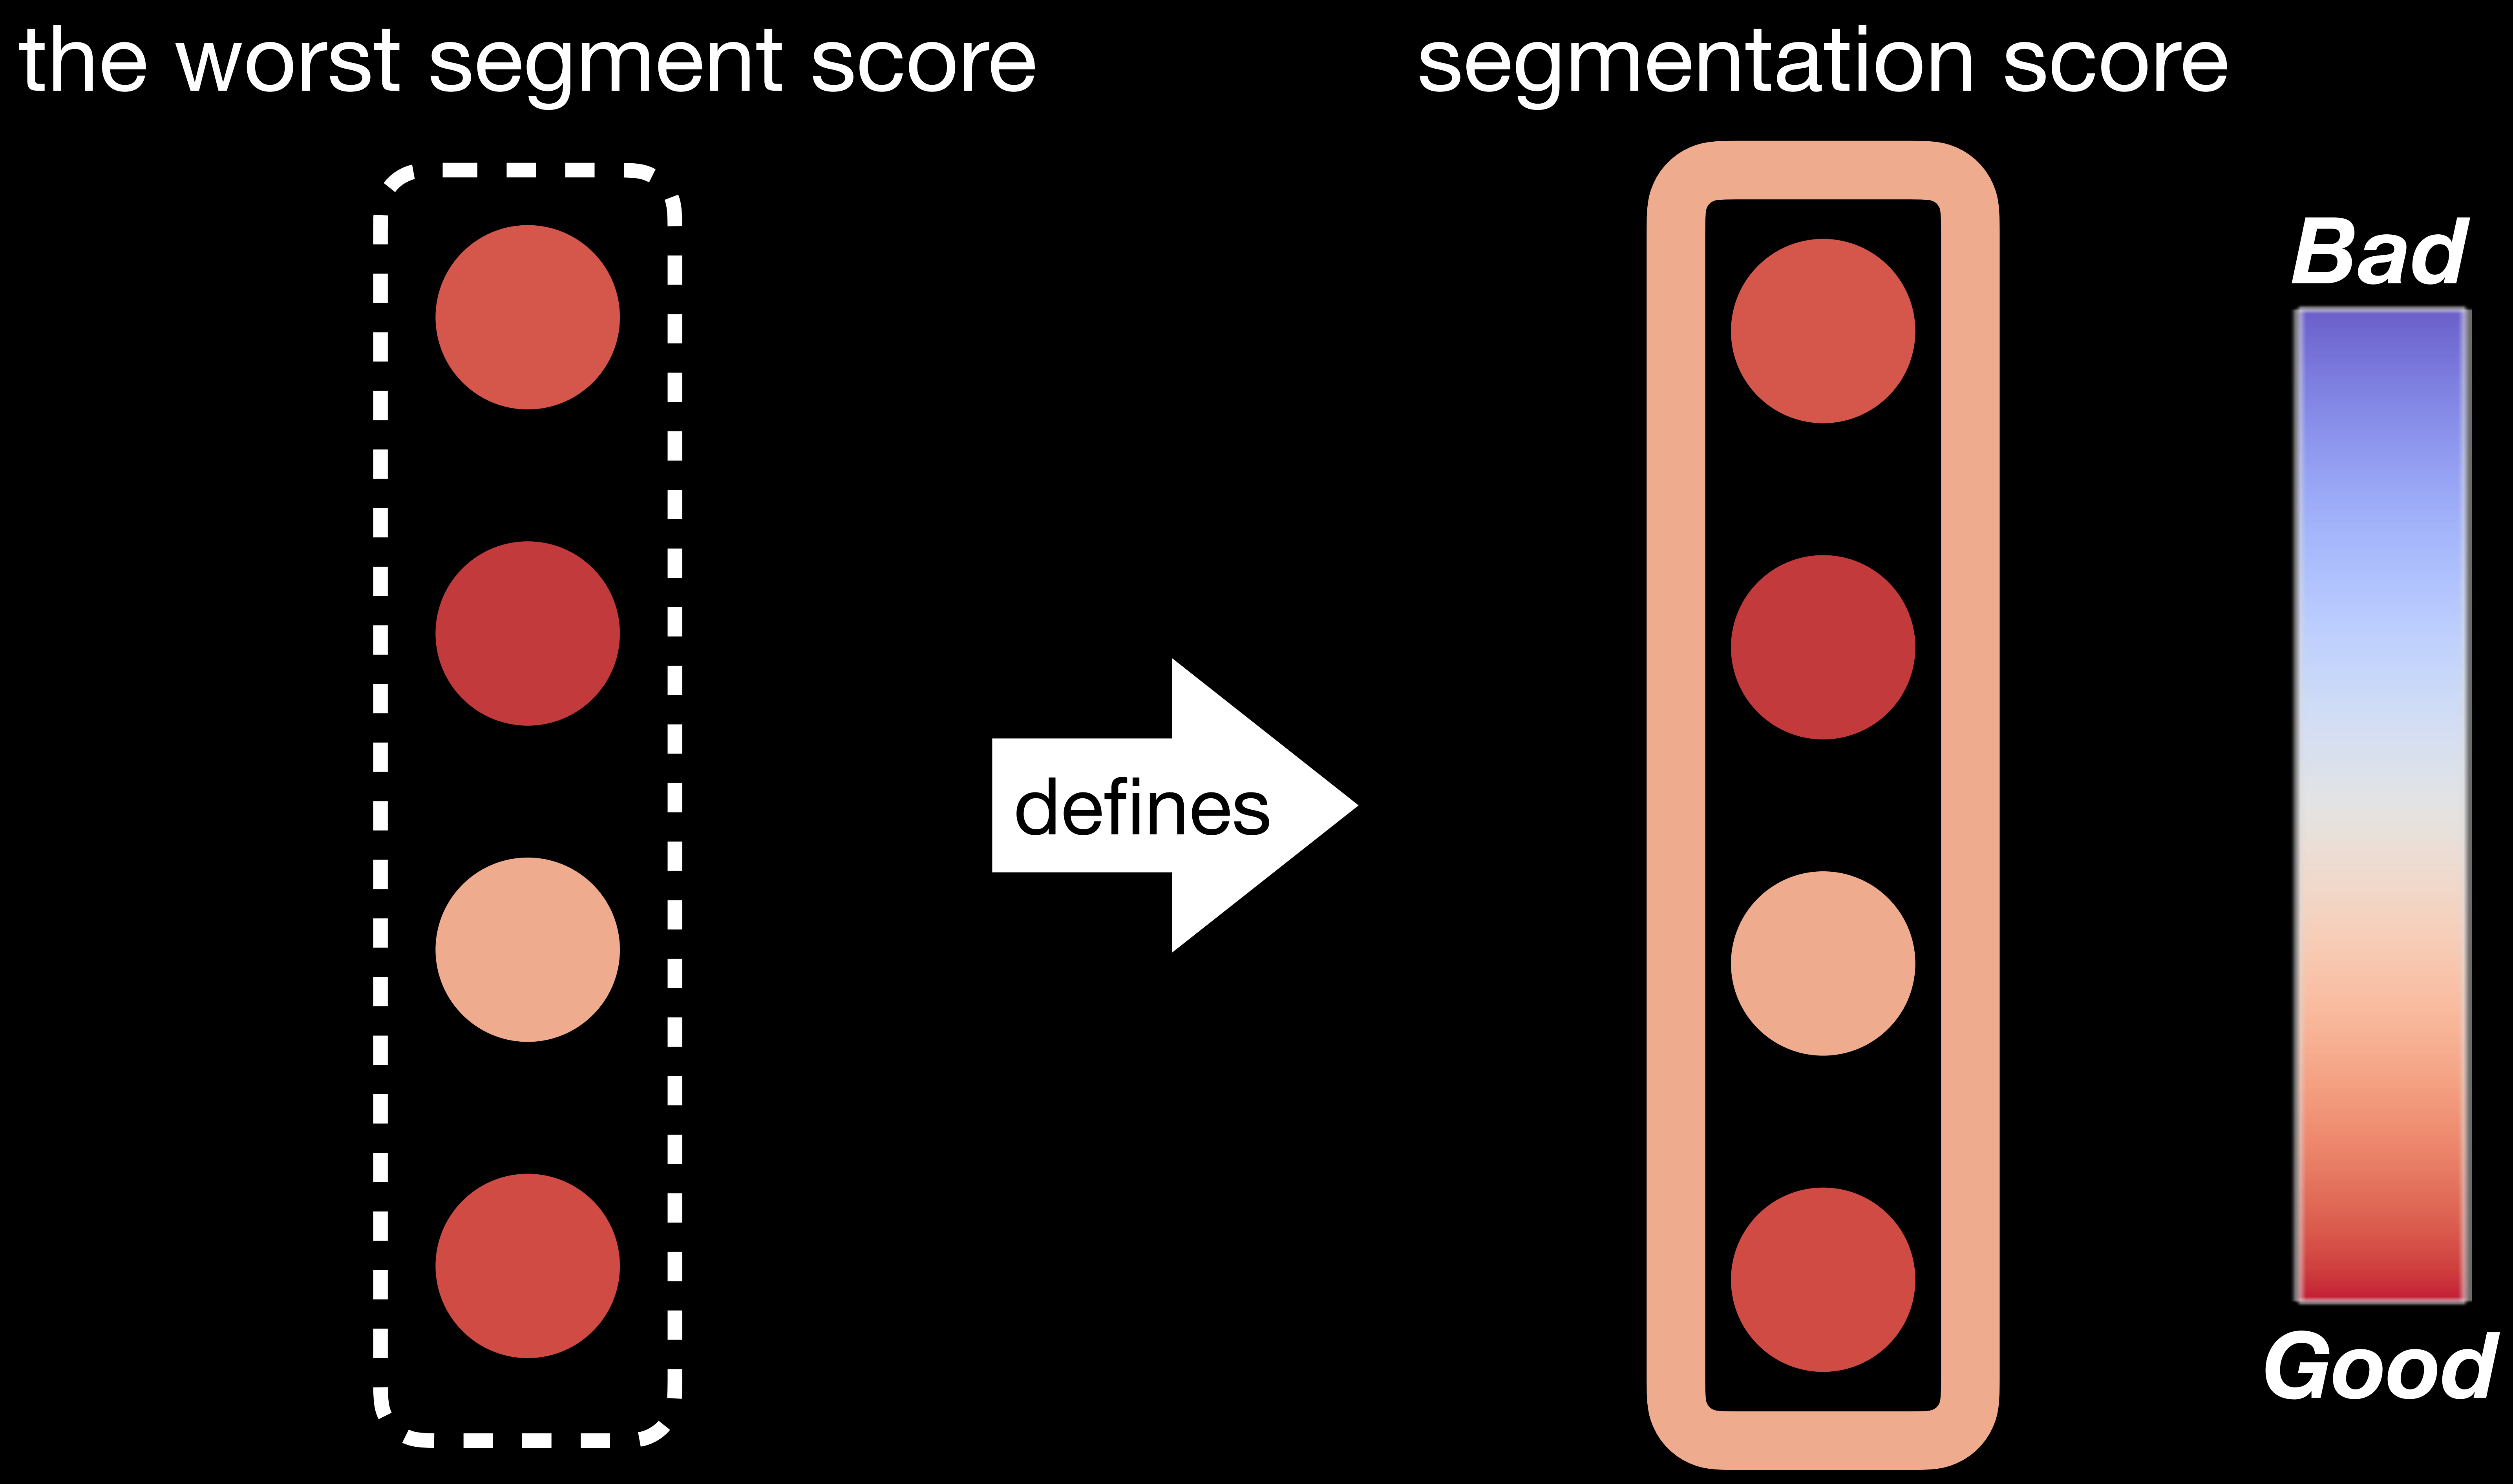

Average-case Segmentation

We also propose an efficient algorithm guided by average-case score.

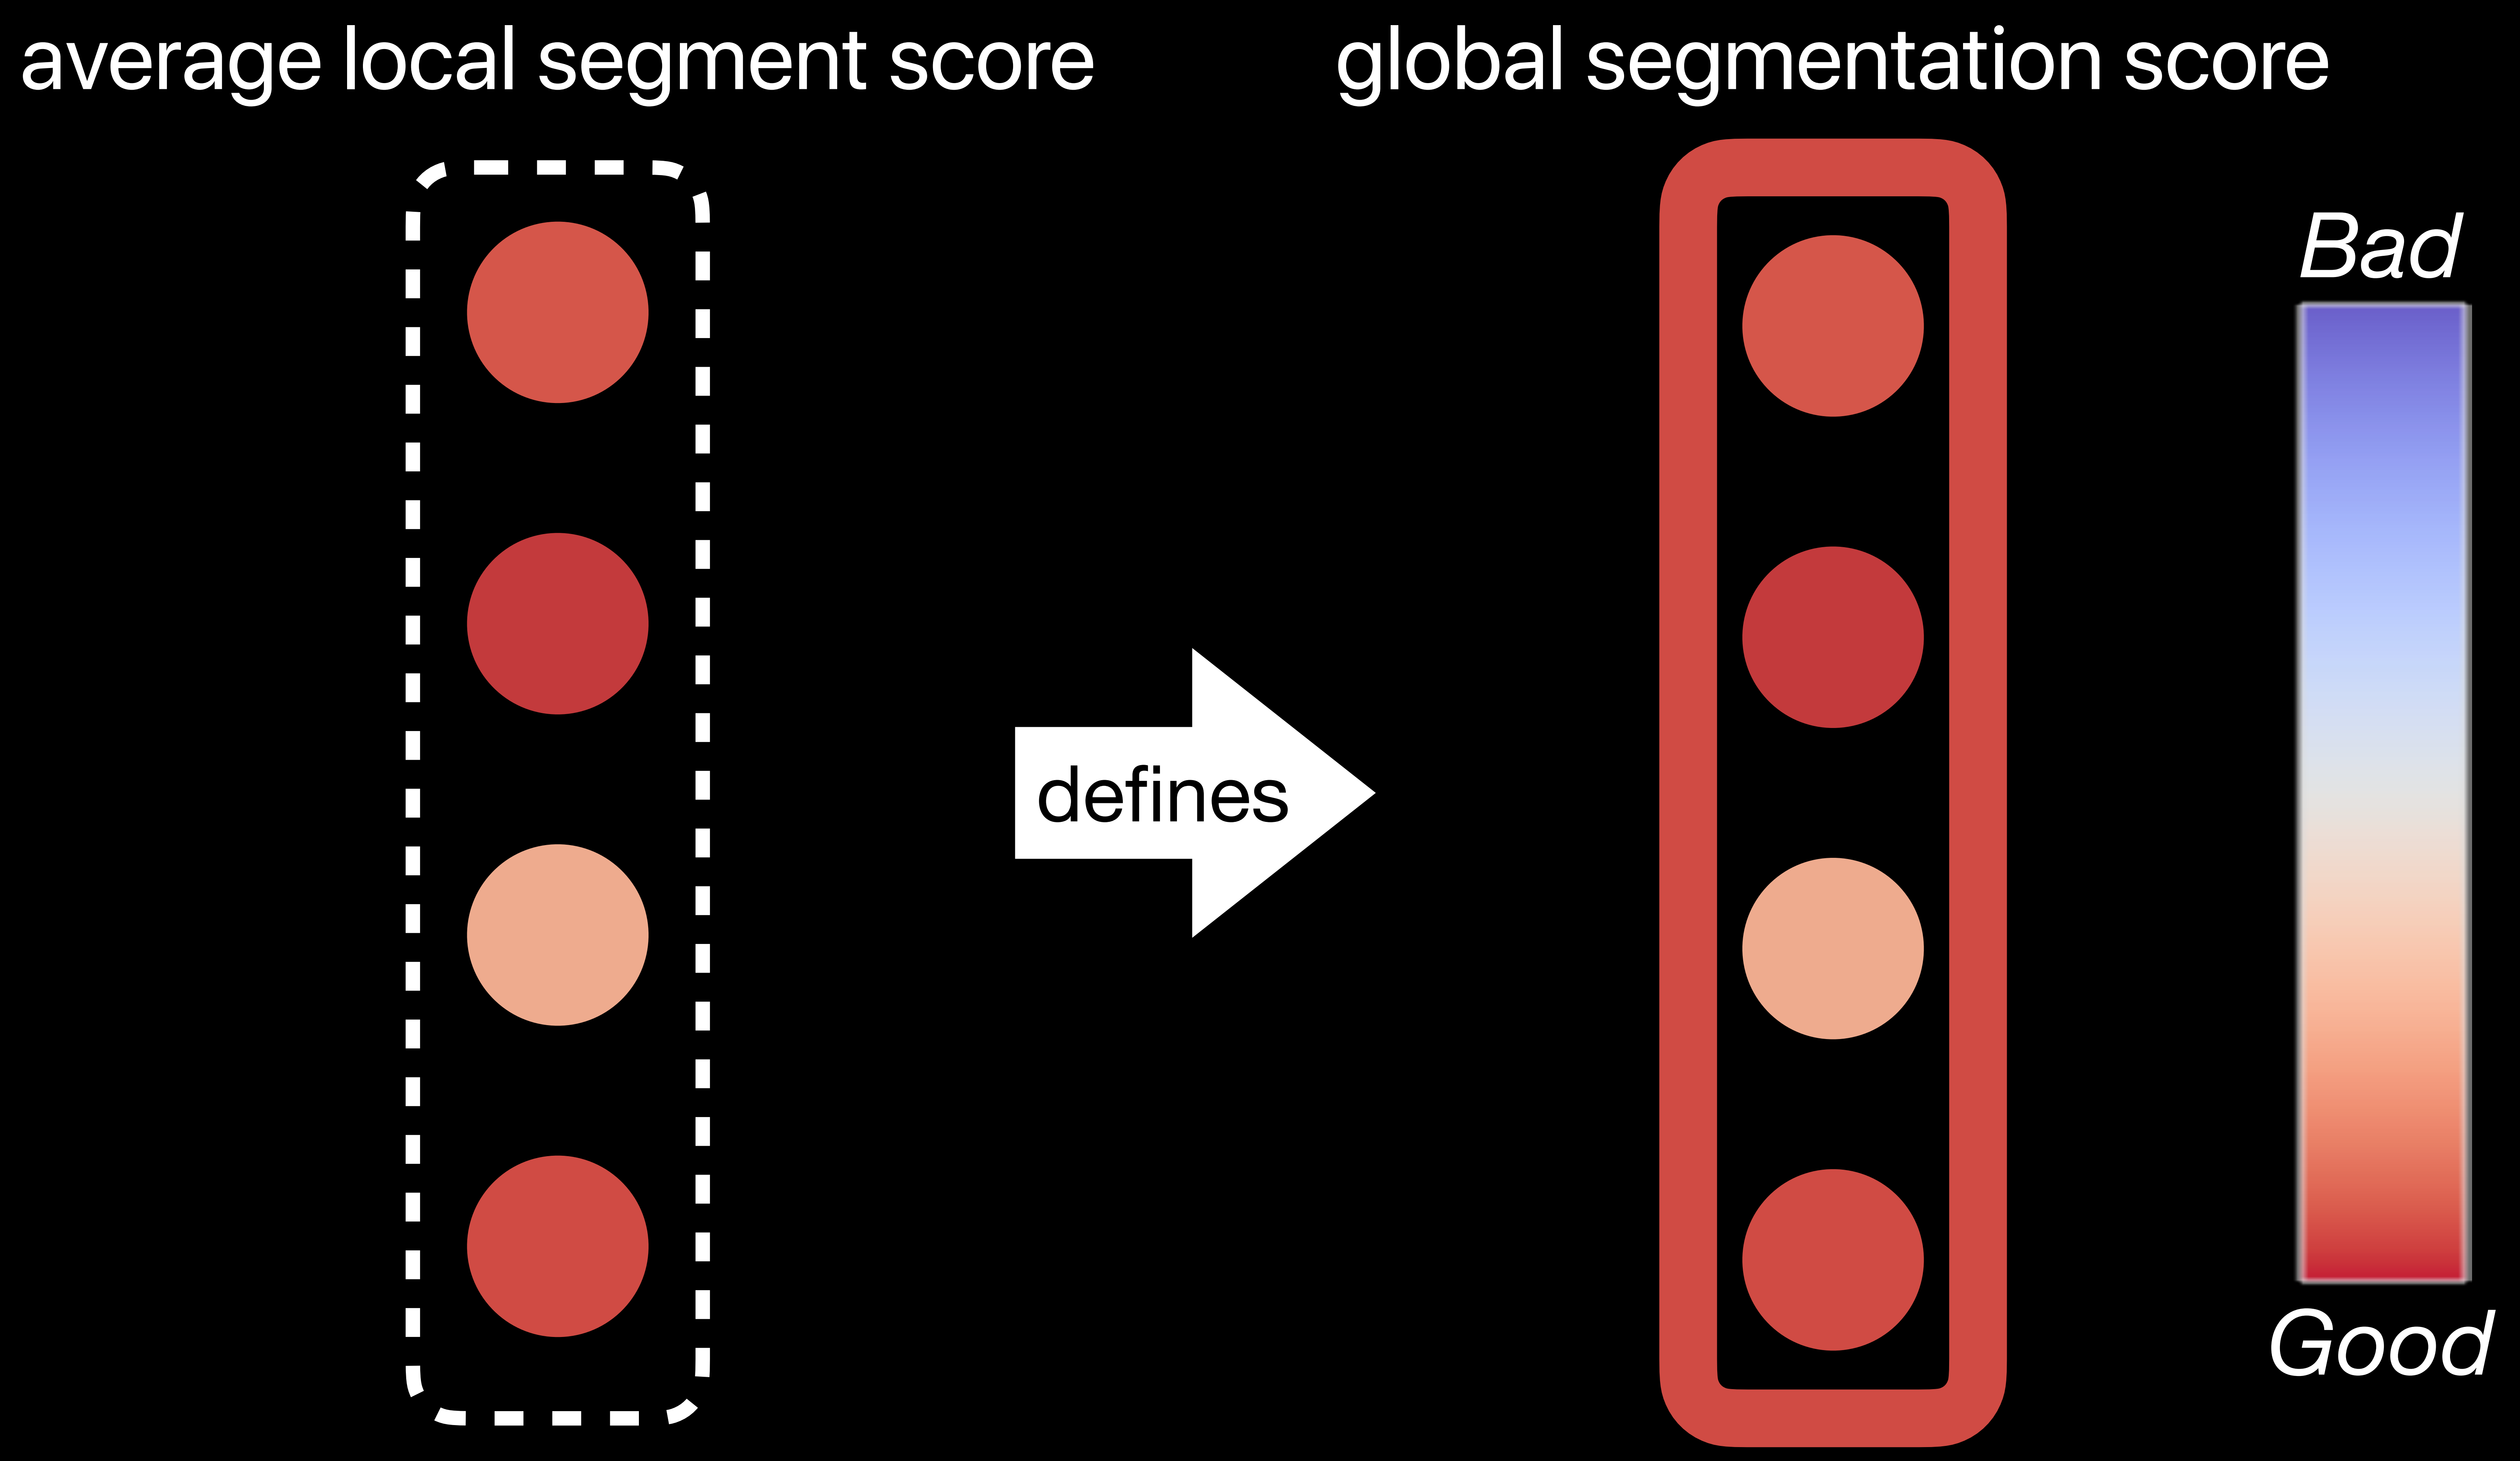

Quantitative Evaluation

Protocol: compute the percentage of objects that are under-segmented and over-segmented.

Assumption: output is a valid partition.

$$U = \frac{1}{L}\sum_{l=1}^{L} \mathbf{1}[\frac{|C_{i*} \cap C_l^{gt}|}{|C_{i*}|} < \tau_U] \qquad O = \frac{1}{L}\sum_{l=1}^{L} \mathbf{1}[\frac{|C_{i*} \cap C_l^{gt}|}{|C_l^{gt}|} < \tau_O]$$

2 under-segmented pedestrians

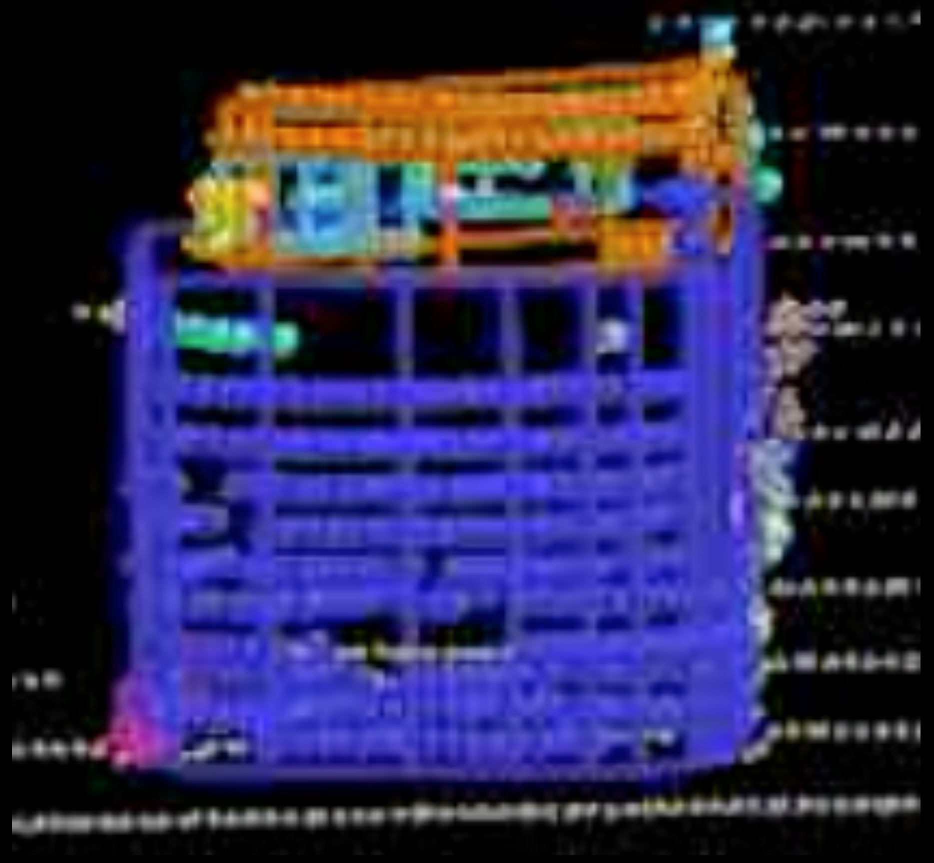

1 over-segmented car

Held et al., RSS'15

Segmentation Errors

(1) As distance threshold increases, more under-segmentation and less over-segmentation.
(2) Our adaptive algorithm significantly outperforms each single-parameter baseline.
(3) We also plot the lower-bound errors for the search space, showing room for improvement.

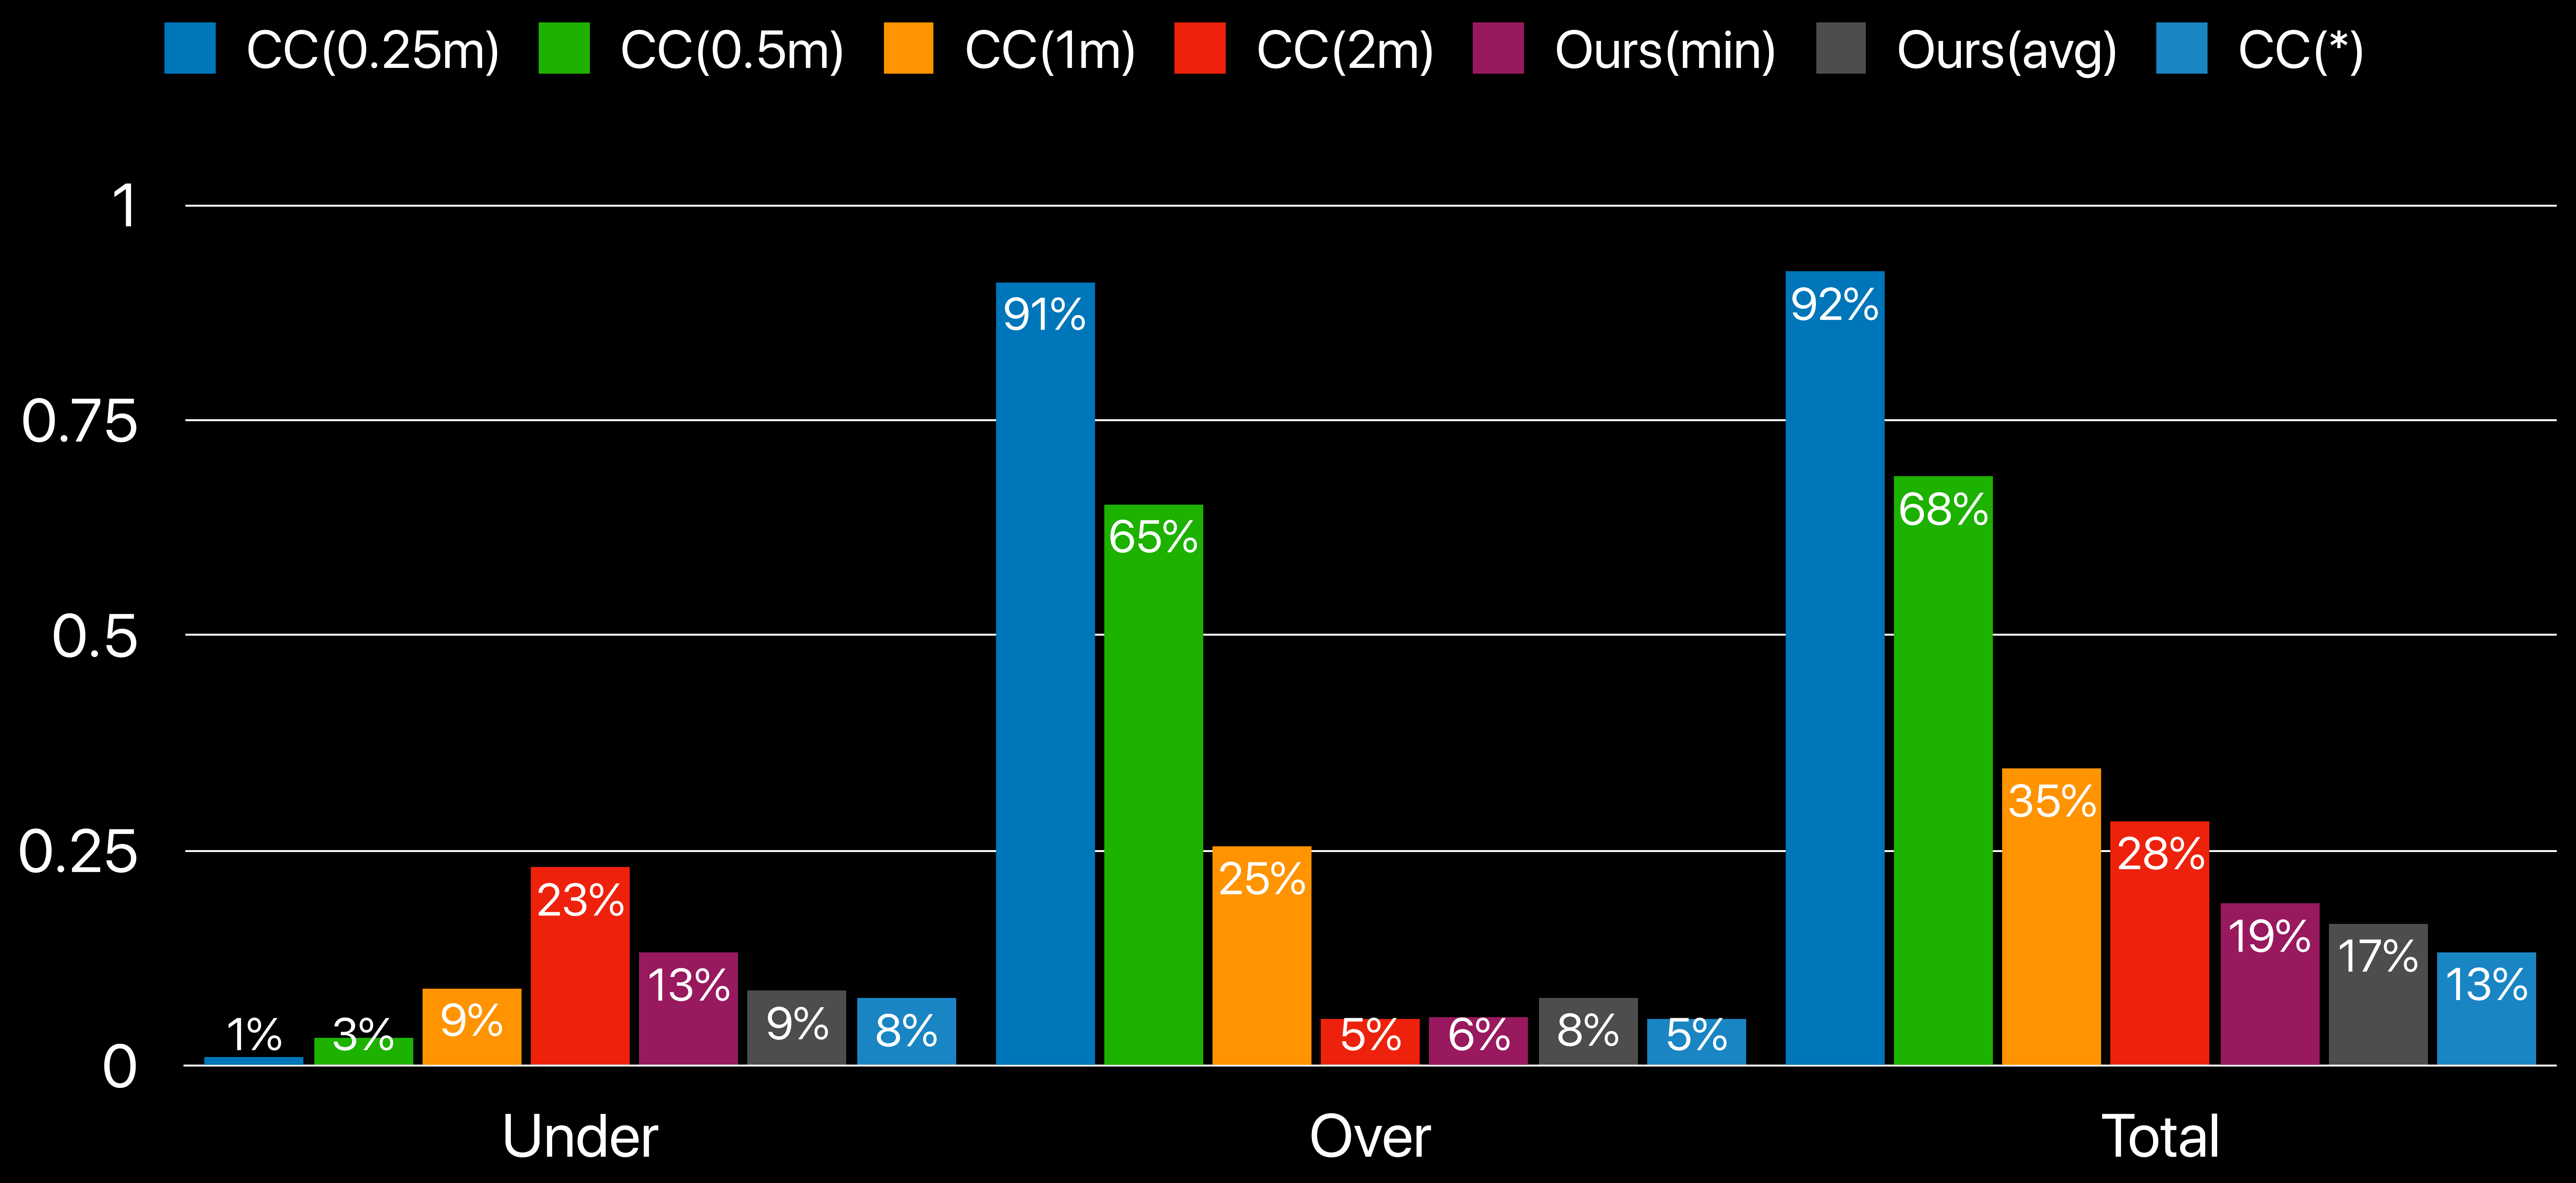

No One-Fits-All Solution

The best distance threshold often varies from scenario to scenario.

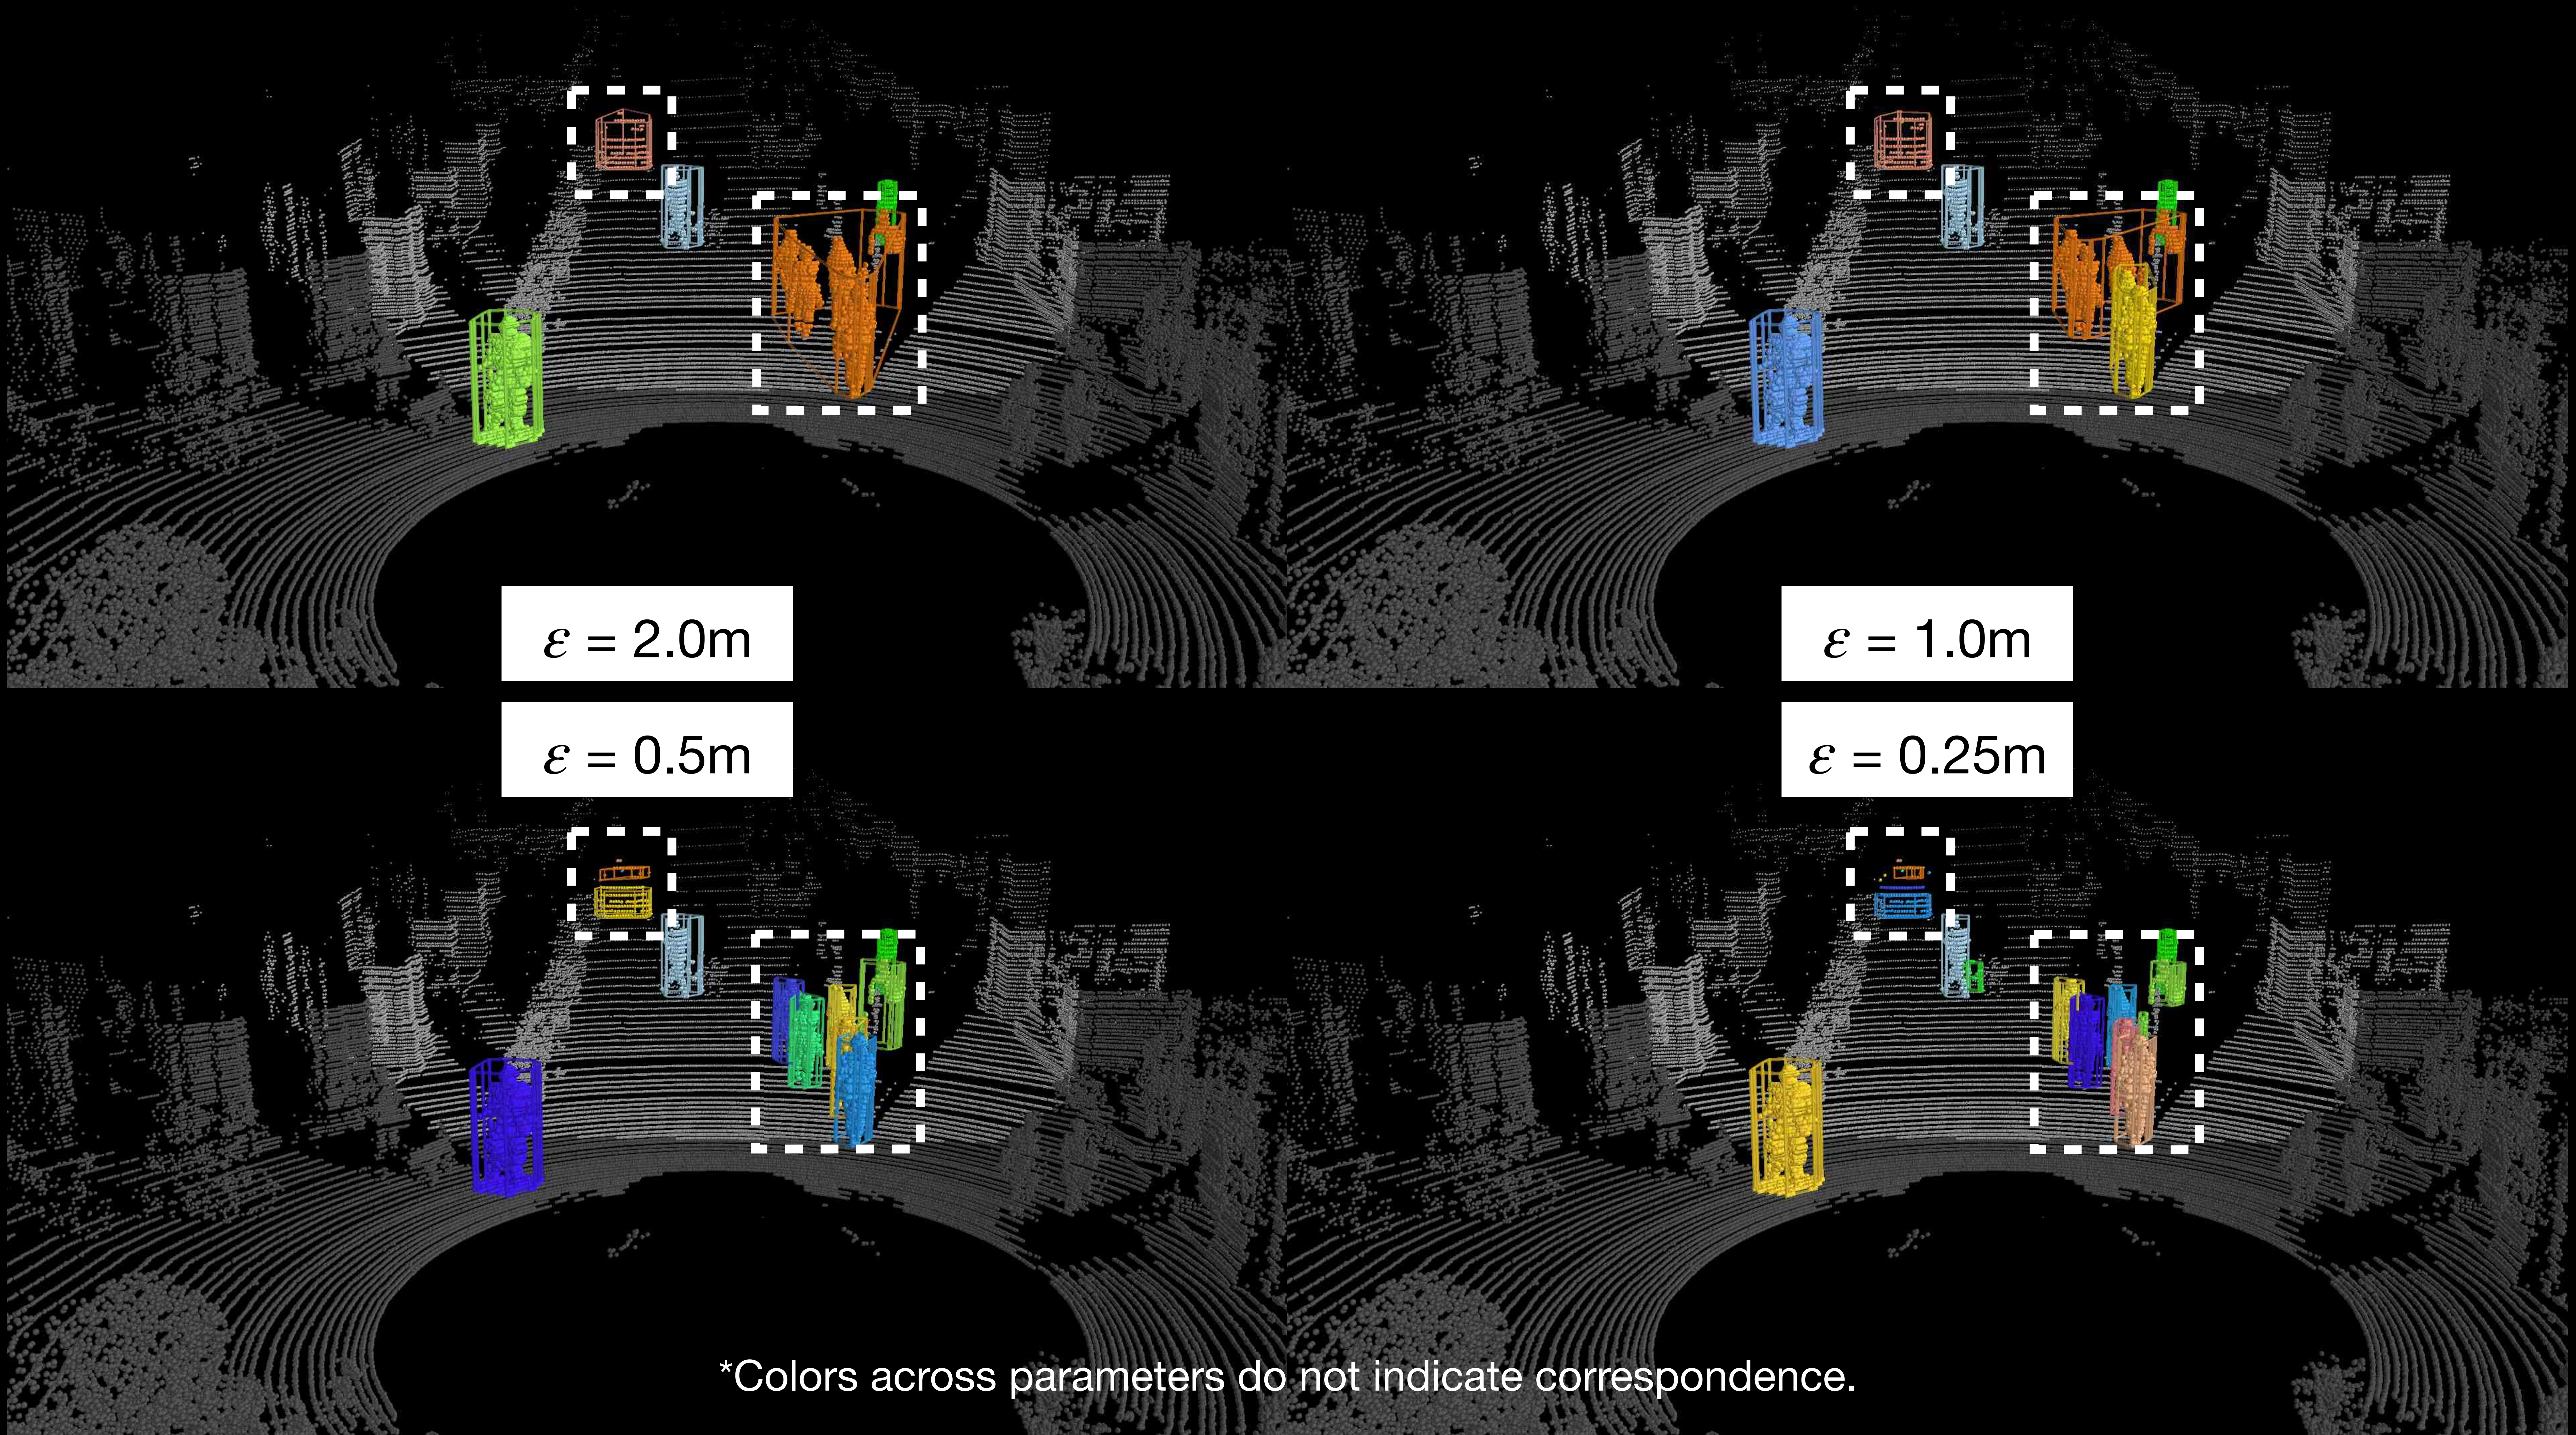

*Colors across parameters do not indicate correspondence.

Algorithmic Output

Our algorithm can adaptively choose the best distance threshold for each scenario.

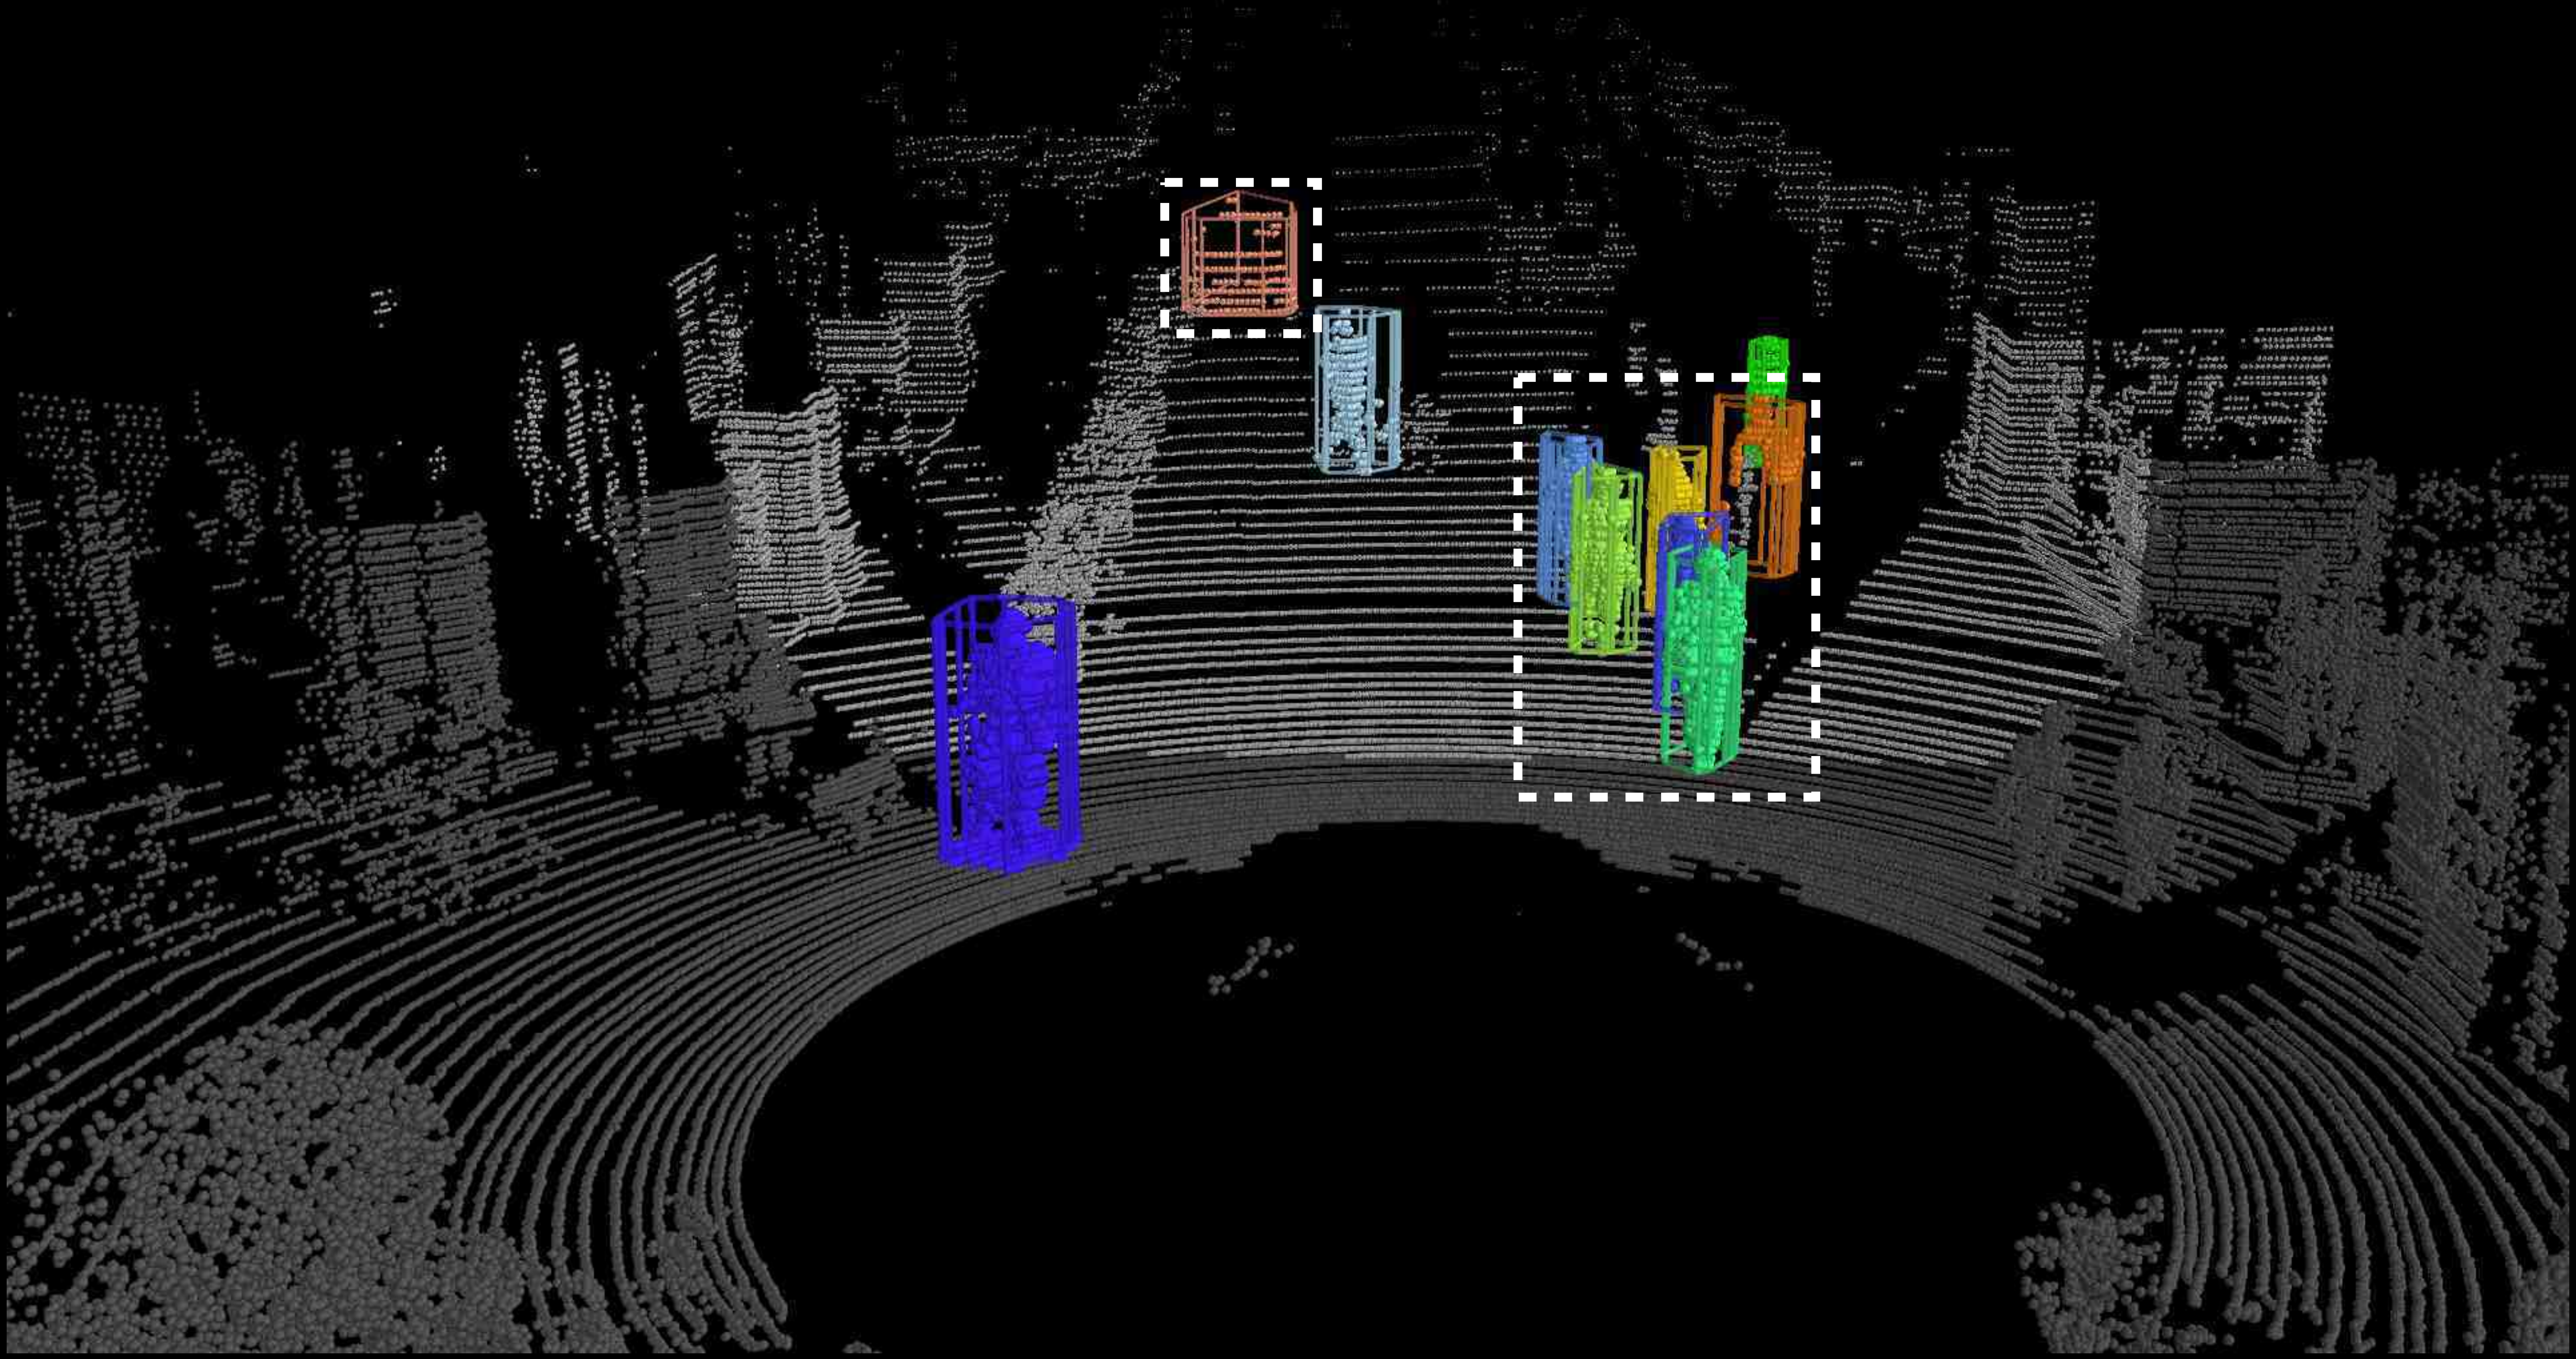

Learning to Optimally Segment Point Clouds

Peiyun Hu[1] David Held[1*] Deva Ramanan[1,2*]
[1]Robotics Institute, Carnegie Mellon University
[2]Argo AI

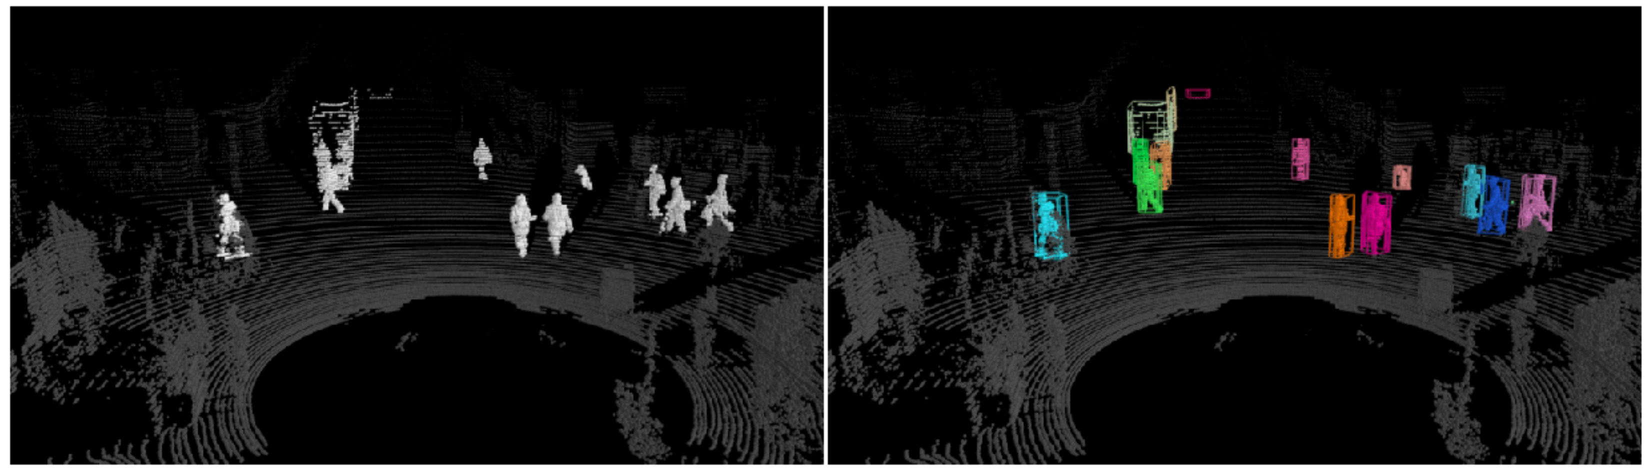

Our proposed algorithm takes a pre-processed LiDAR point cloud with background removed (left) and produces a class-agnostic instance-level segmentation over all foreground points (right). For visualization, we use a different color for each segment and plot an extruded polygon to show the spatial extent.

Qualitative results

Code

We've released our code at https://github.com/peiyunh/opcseg.

https://cs.cmu.edu/~peiyunh/opcseg