


What You See Is What You Get

Exploiting Visibility for 3D Object Detection



Peiyun Hu, Jason Ziglar, David Held, Deva Ramanan



Carnegie Mellon University
Argo AI

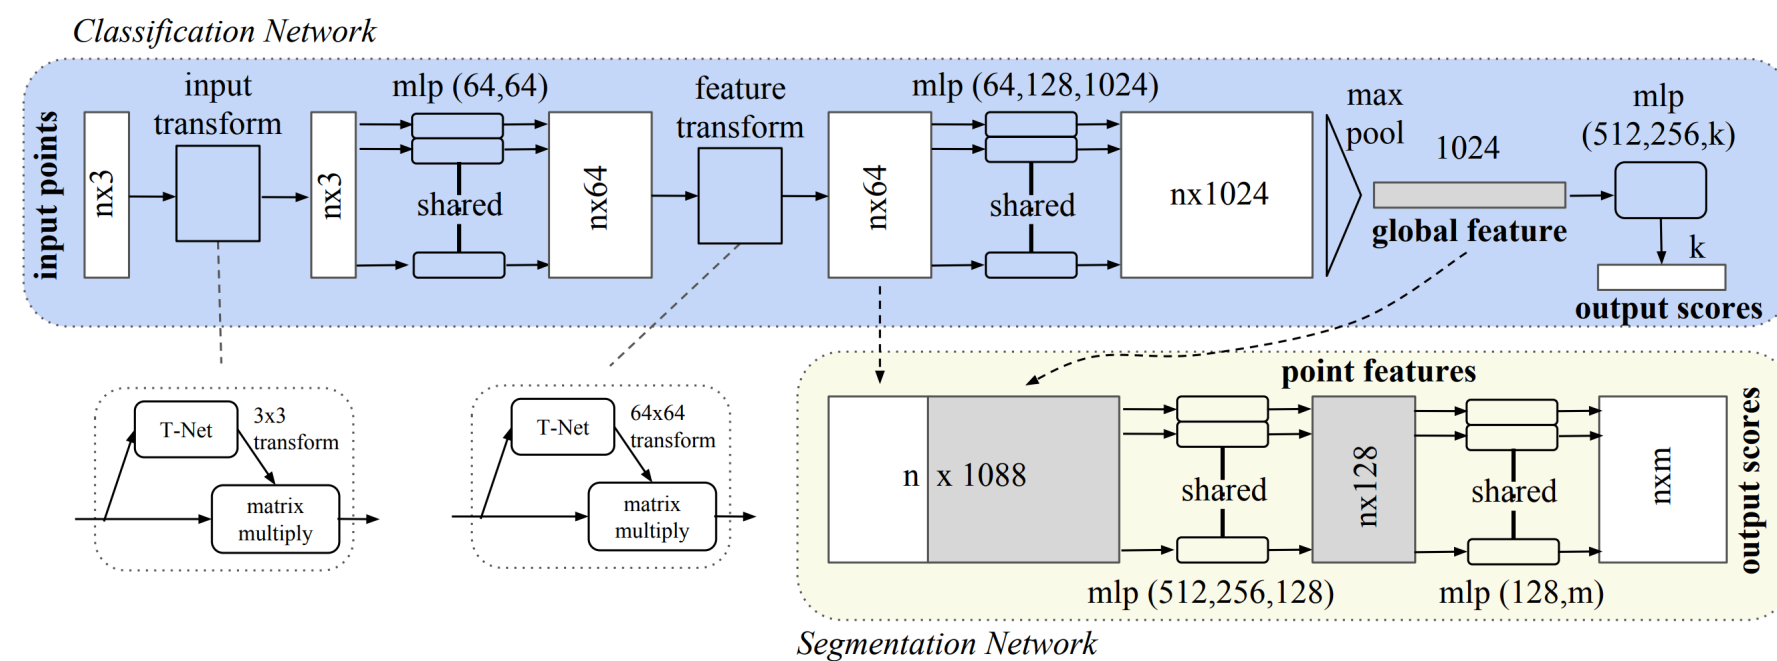


What is a good representation for LiDAR data?




- LiDAR data provides more than just point measurements
- Rays emanating from the sensor to each 3D point **must** pass through free space
- Representing LiDAR data as (x, y, z) s fundamentally destroys such freespace information

What representations do we have?



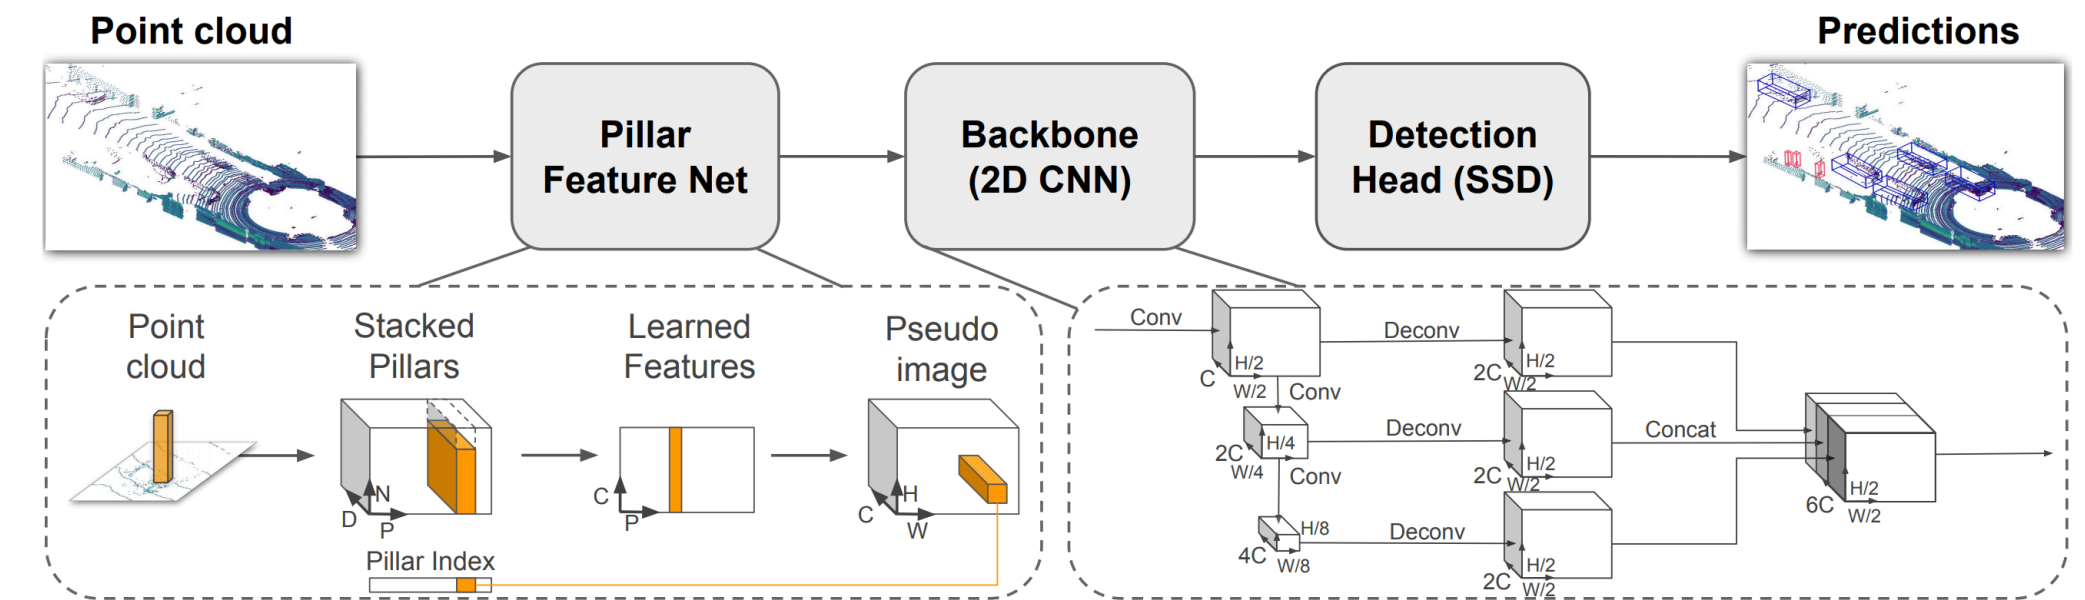
Deep Point Representation

PointNet, Qi et al., CVPR'17




Occupancy Voxels

OctoMap, Hornung et al., Autonomous Robots'13



Deep Voxel Representation

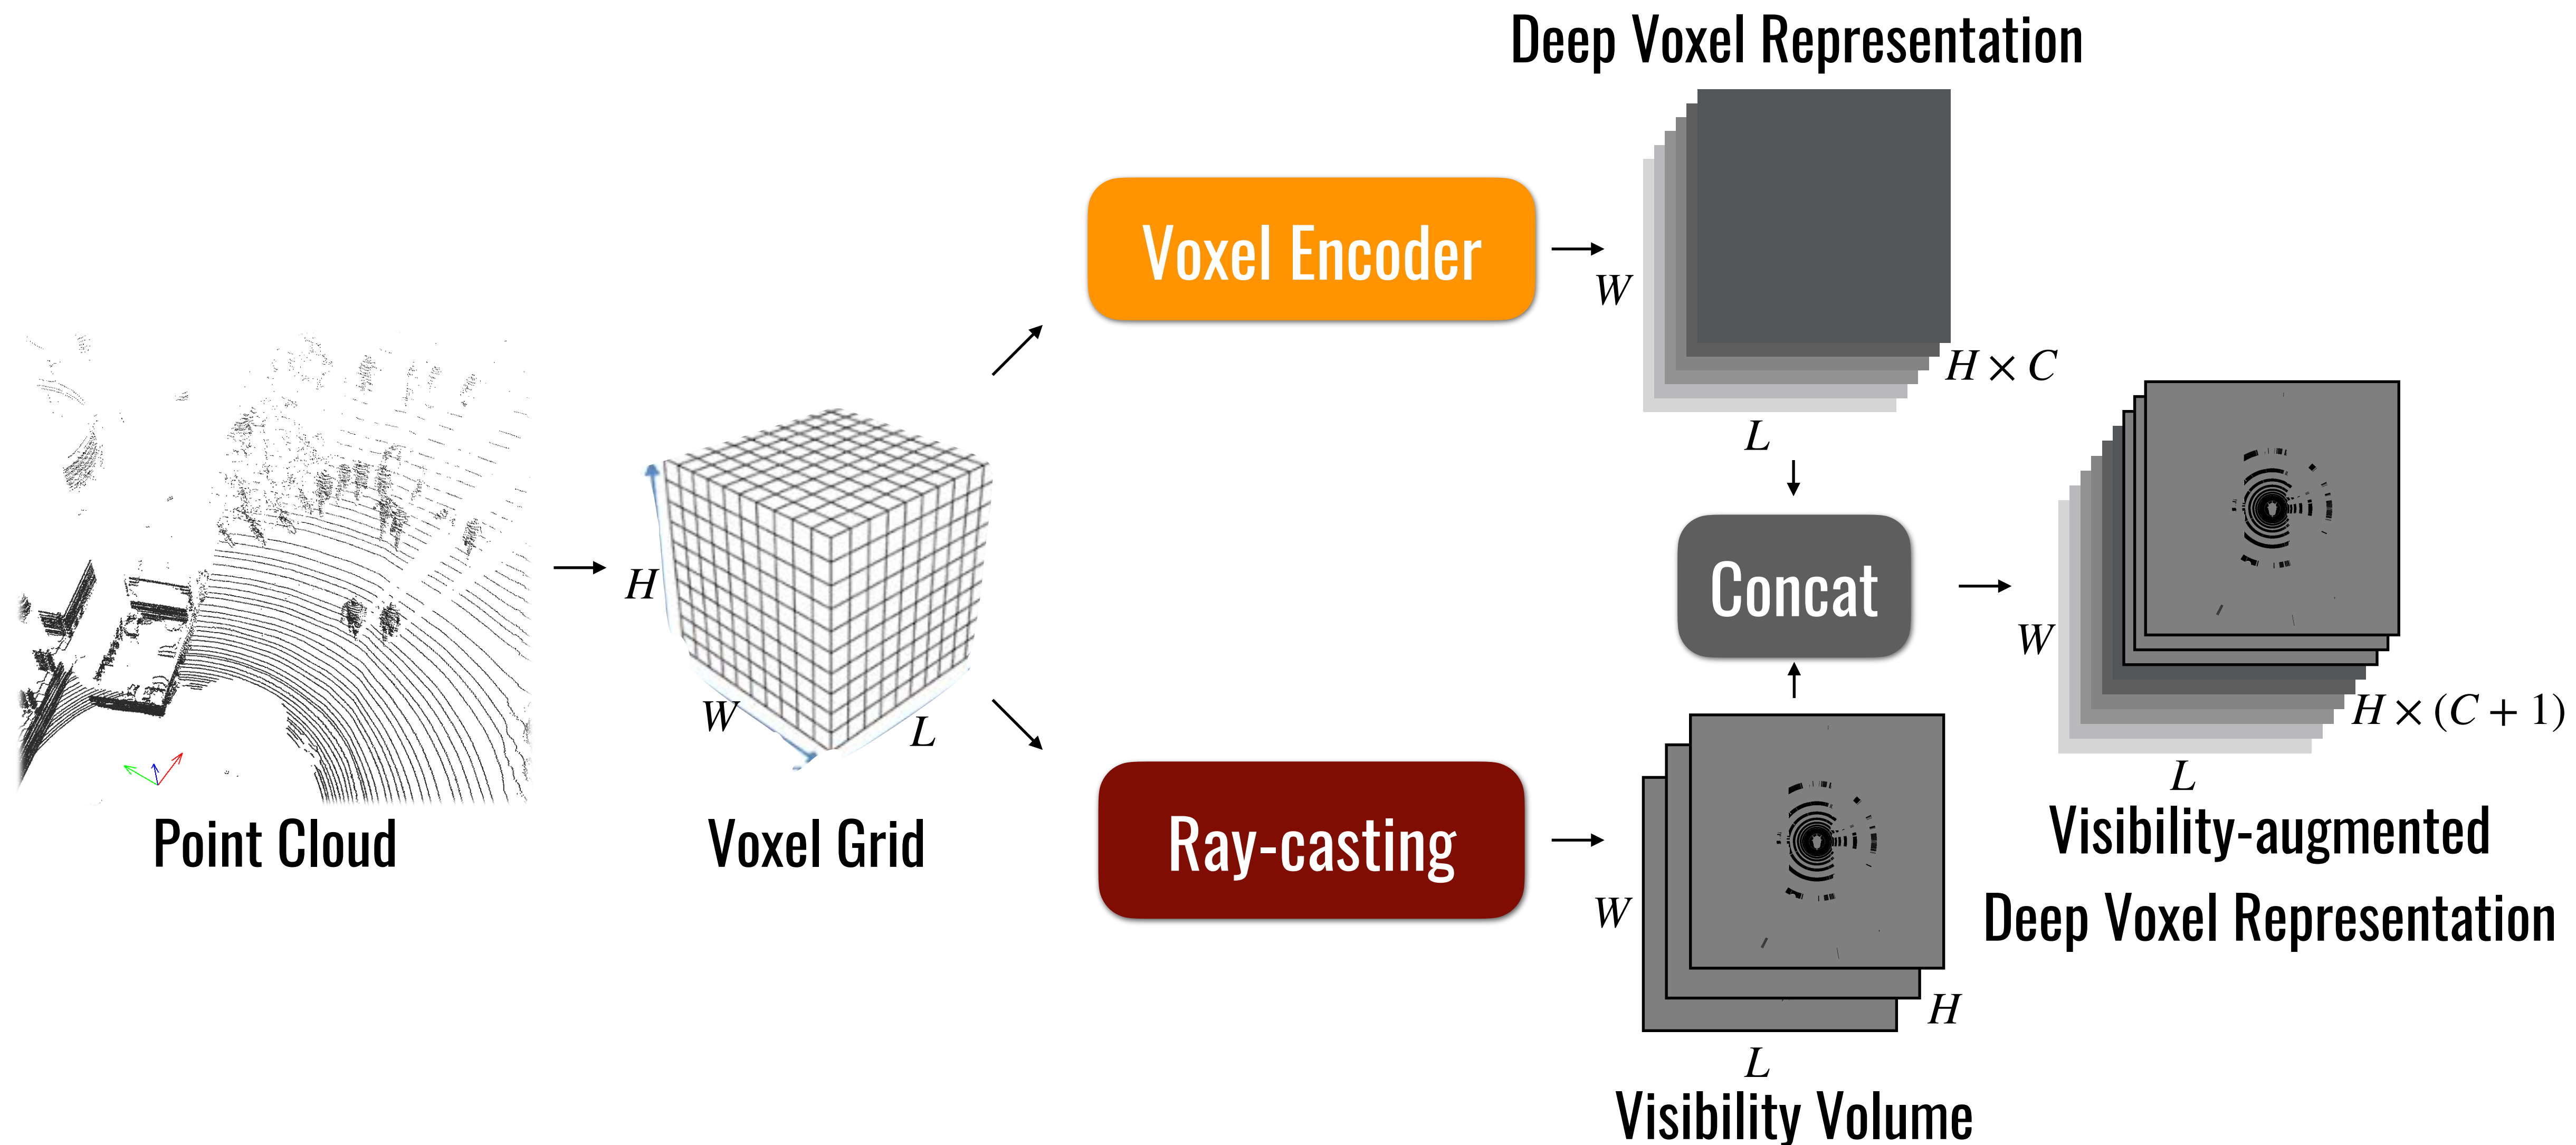
PointPillars, Lang et al., CVPR'19



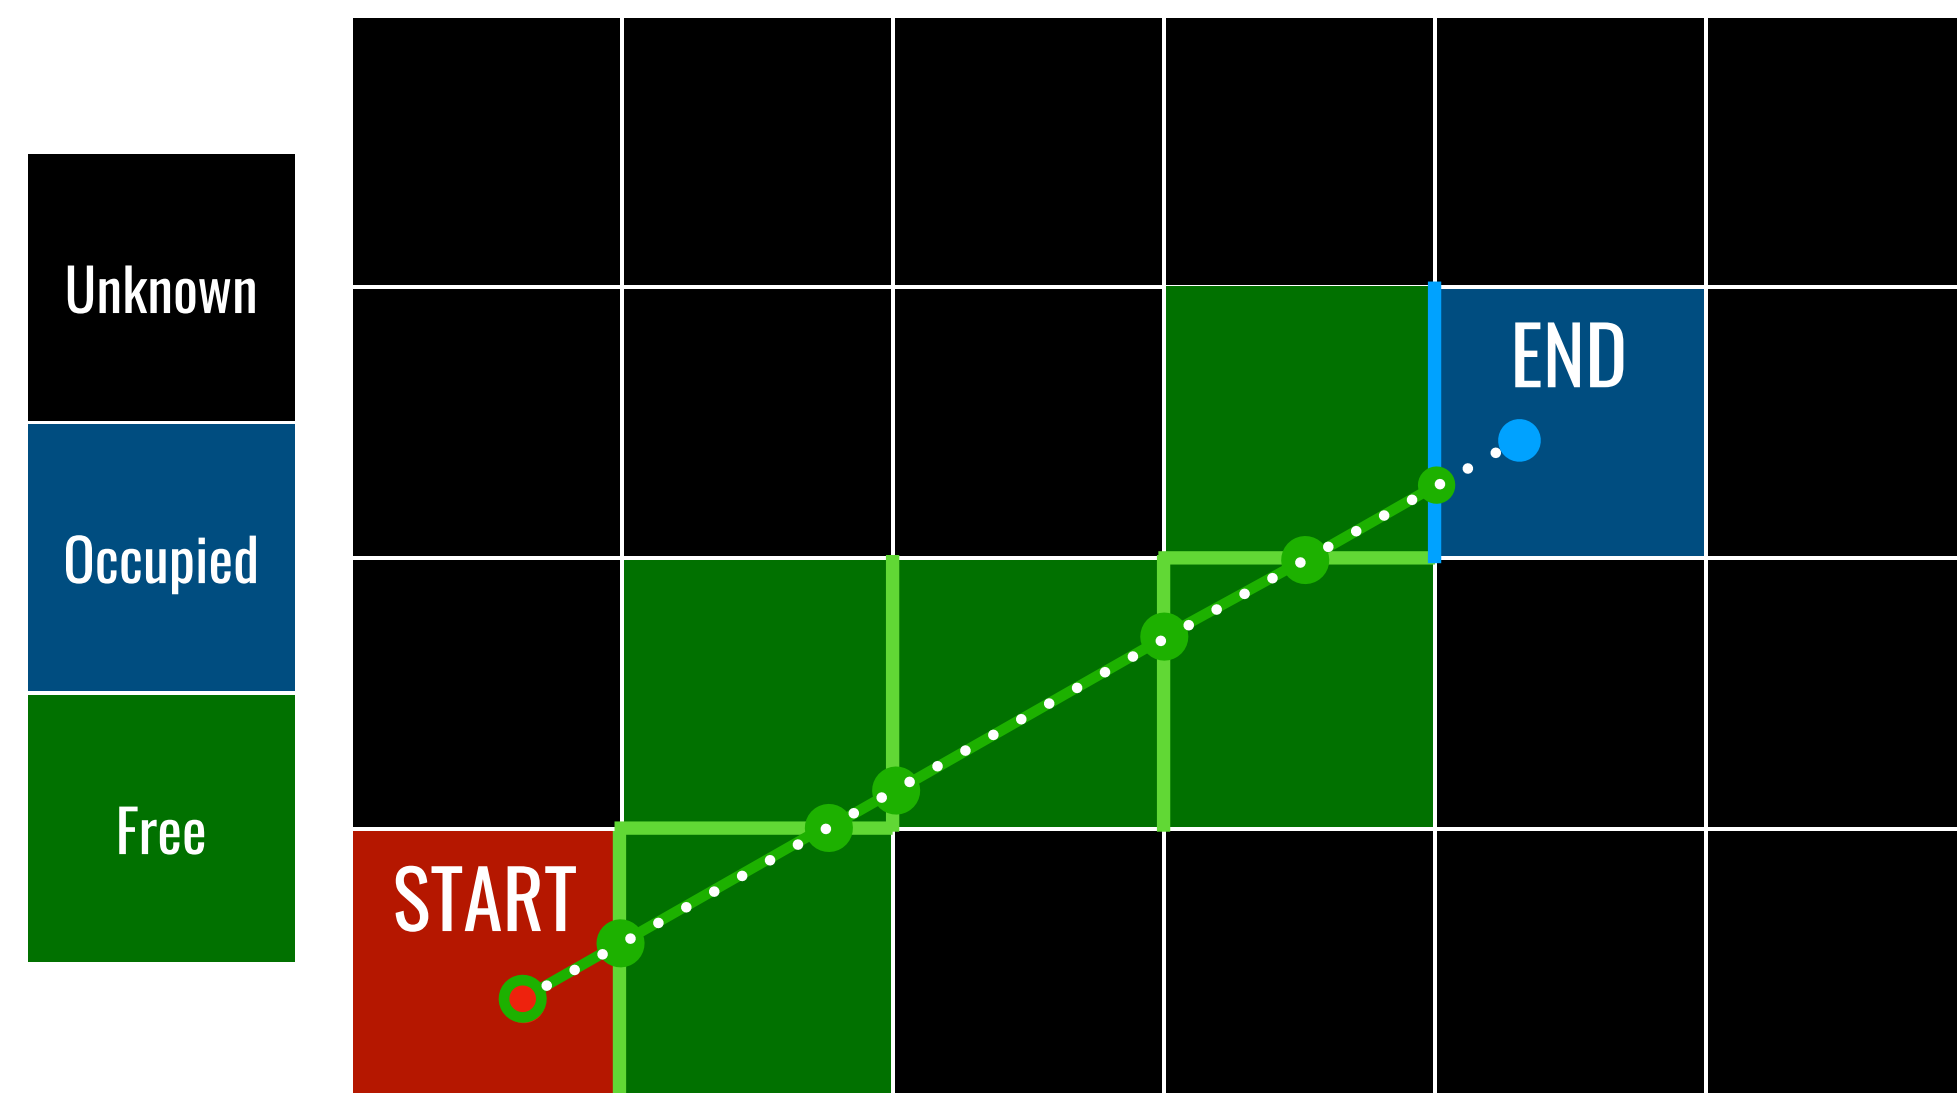
Visibility Augmented Deep Voxels

WYSIWYG, Hu et al., CVPR'20

A Simple Approach to Augment Visibility



Efficient Ray-casting via Voxel Traversal




Though animated in 2D, the idea generalizes in 3D.

A Fast Voxel Traversal Algorithm for Ray Tracing

John Amanatides, Andrew Woo

Eurographics 1987




3D Visibility Volume


Visibility over Multiple LiDAR Sweeps




Single sweep



Discrete visibility
(one slice)



Multiple sweeps




Continuous visibility
(one slice)

OctoMap, Hornung et al., Autonomous Robots'13

Visibility-aware LiDAR Synthesis

Should be occluded!

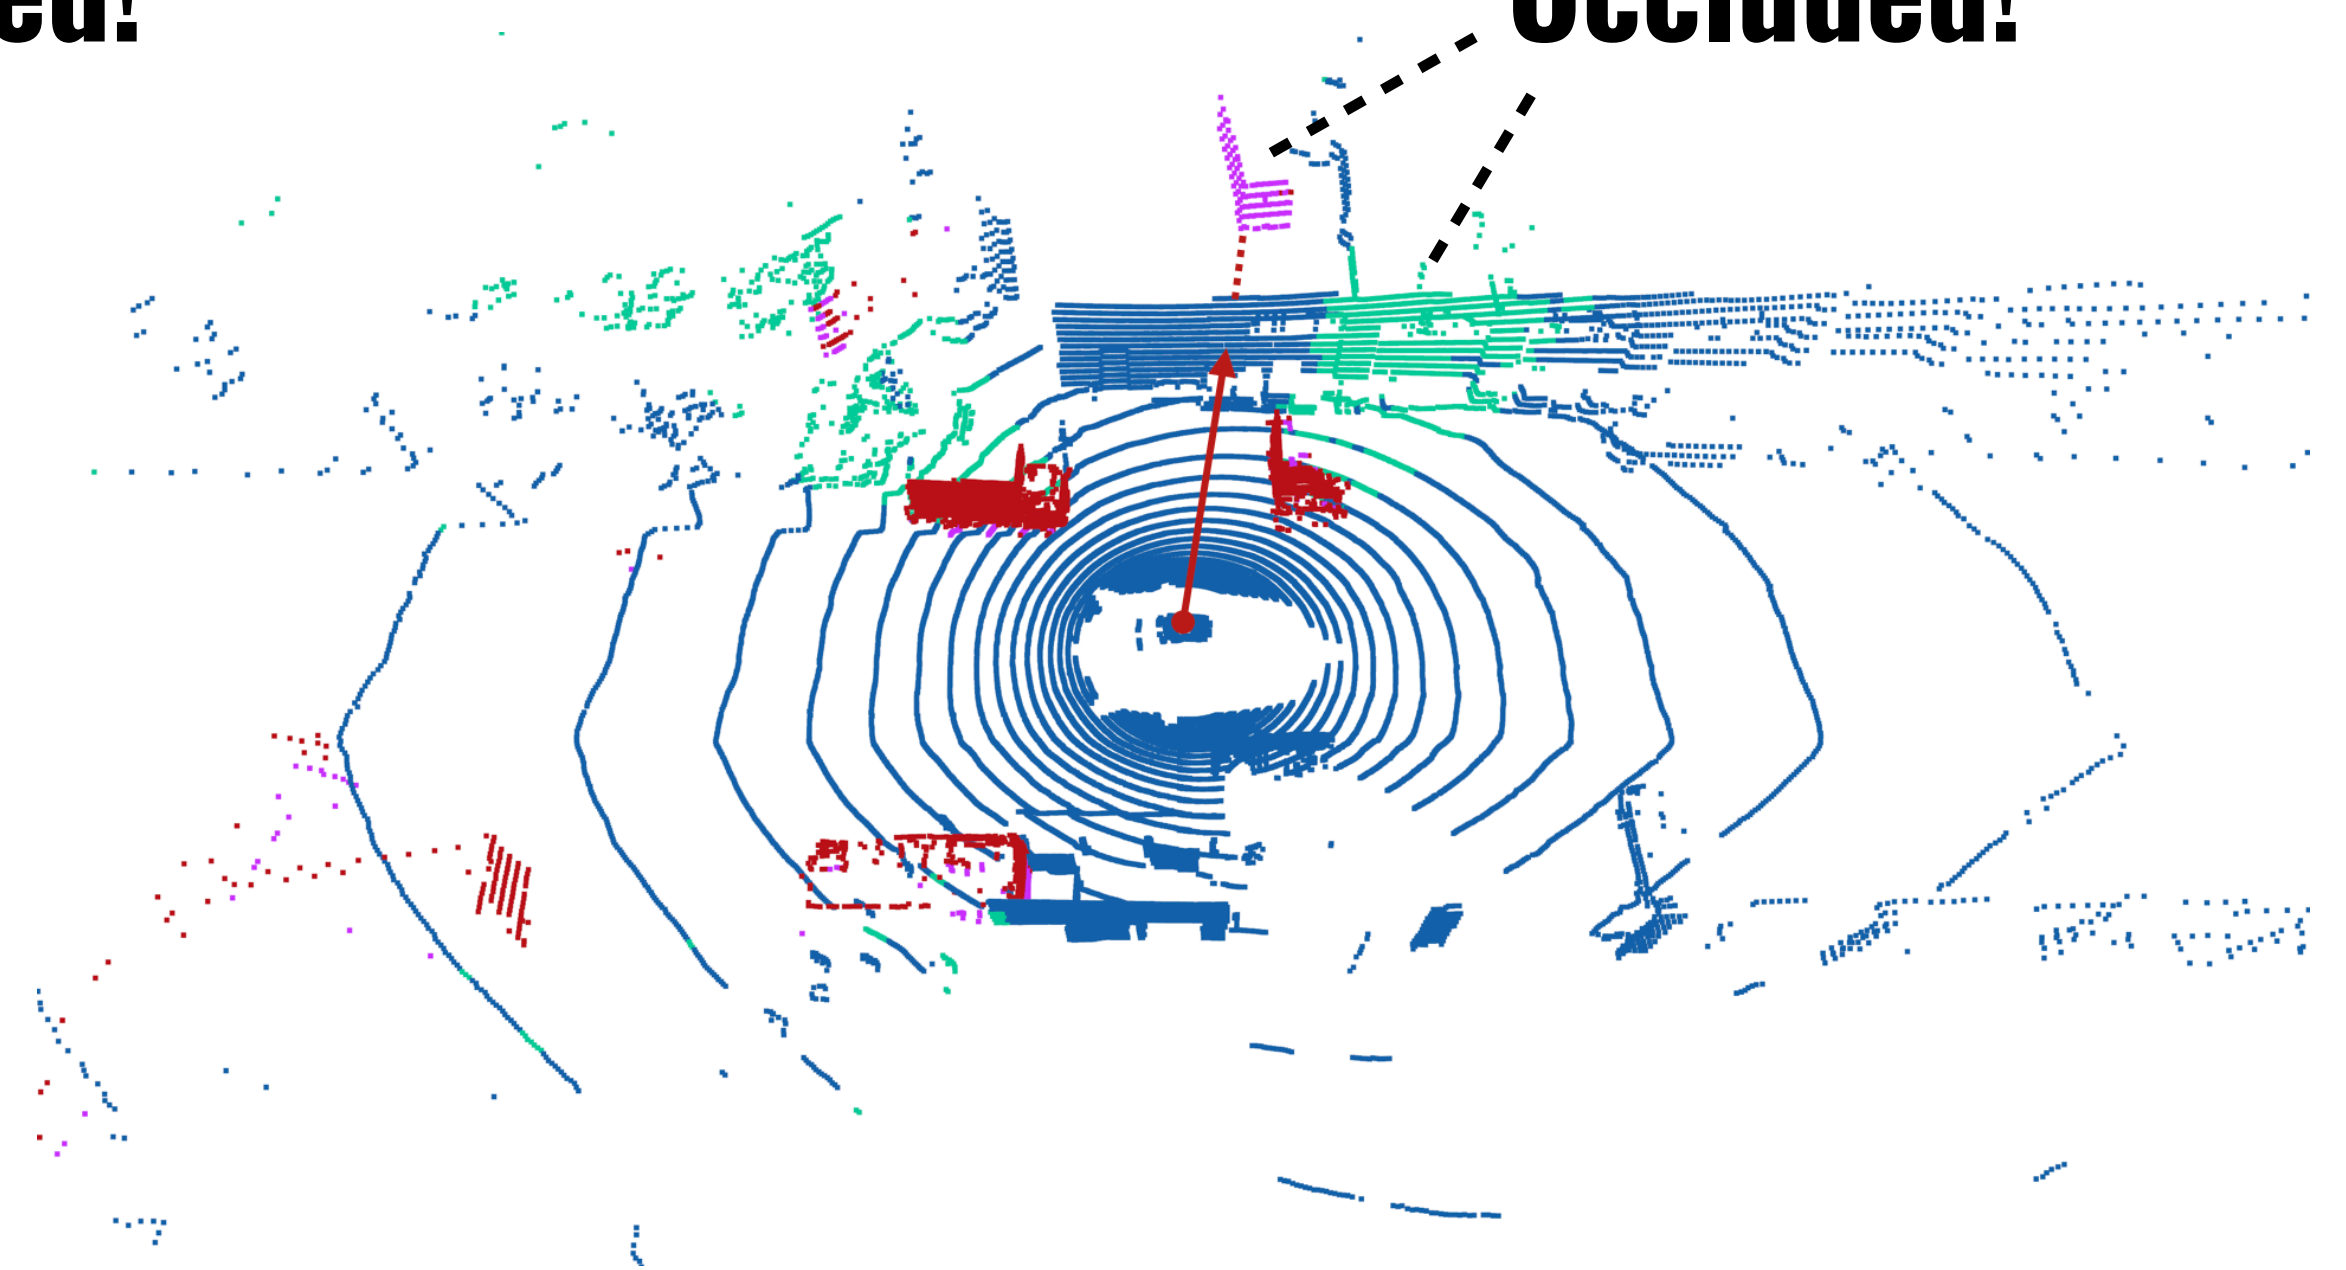


Naive Object Augmentation

PointPillars, Lang et al., CVPR'19


SECOND, Yan et al., Sensors'18

Occluded!



Visibility-aware Object Augmentation

Improve PointPillars by 4.5% in overall mAP



car

pedestrian

barrier

traffic cone

truck

bus

trailer

construction
vehicle

motorcycle

bicycle



<https://cs.cmu.edu/~peiyunh/wysiwyg>