DeepRob

Lecture 19
3D Perception and Point Clouds
University of Michigan | Department of Robotics

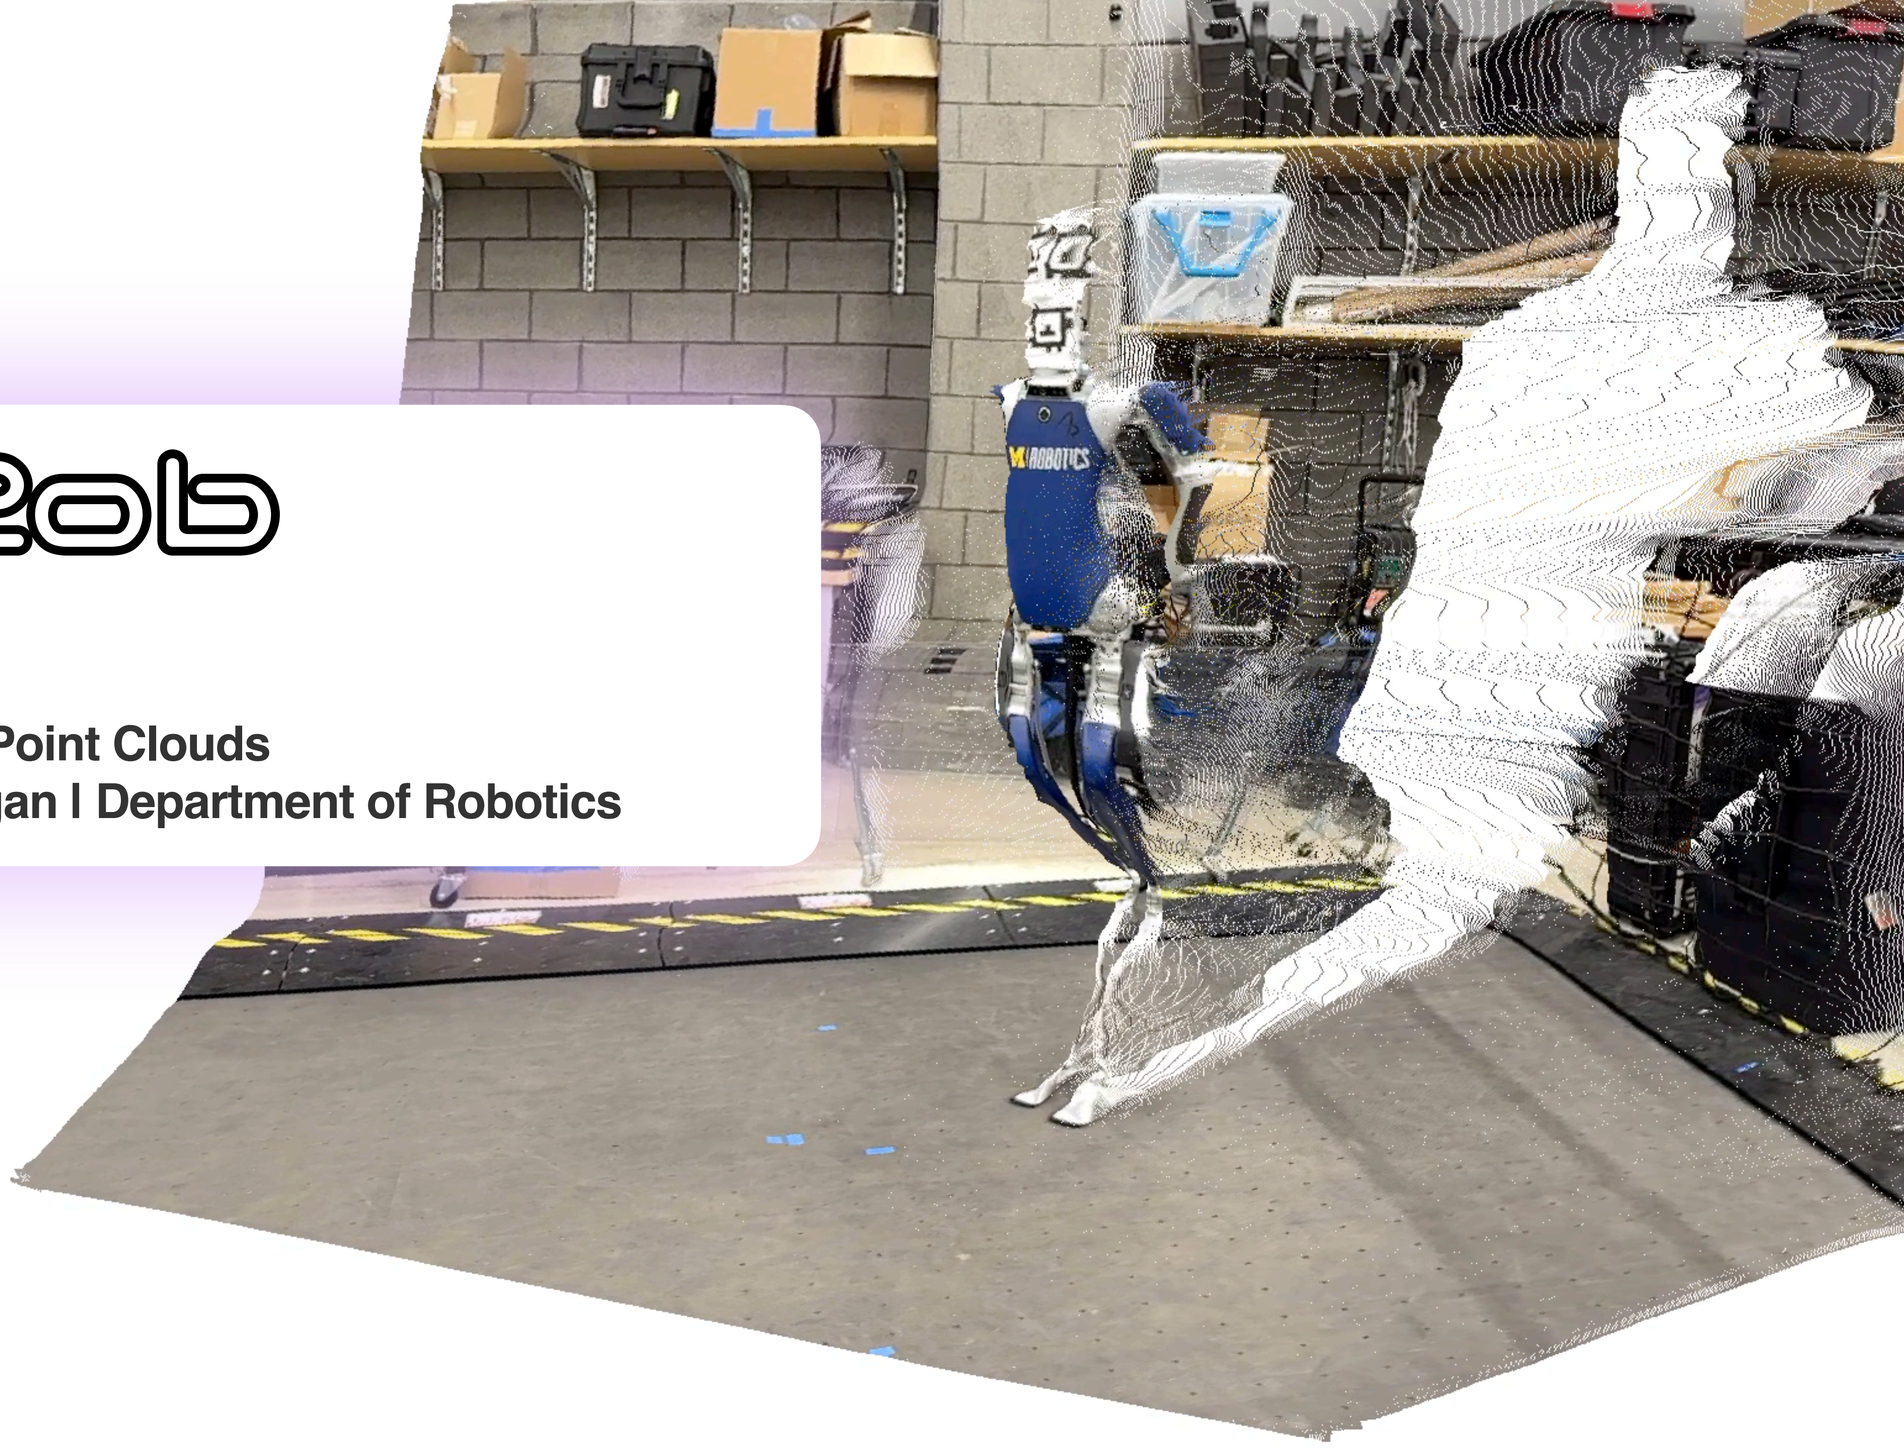

DeepRob

Recall: 2D representation

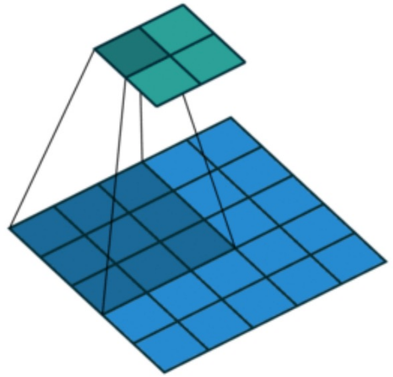

Convolution

2D pixel-grid

3D Vision

3D Point clouds

Example: MUUFL Gulfport dataset; LiDAR BEV

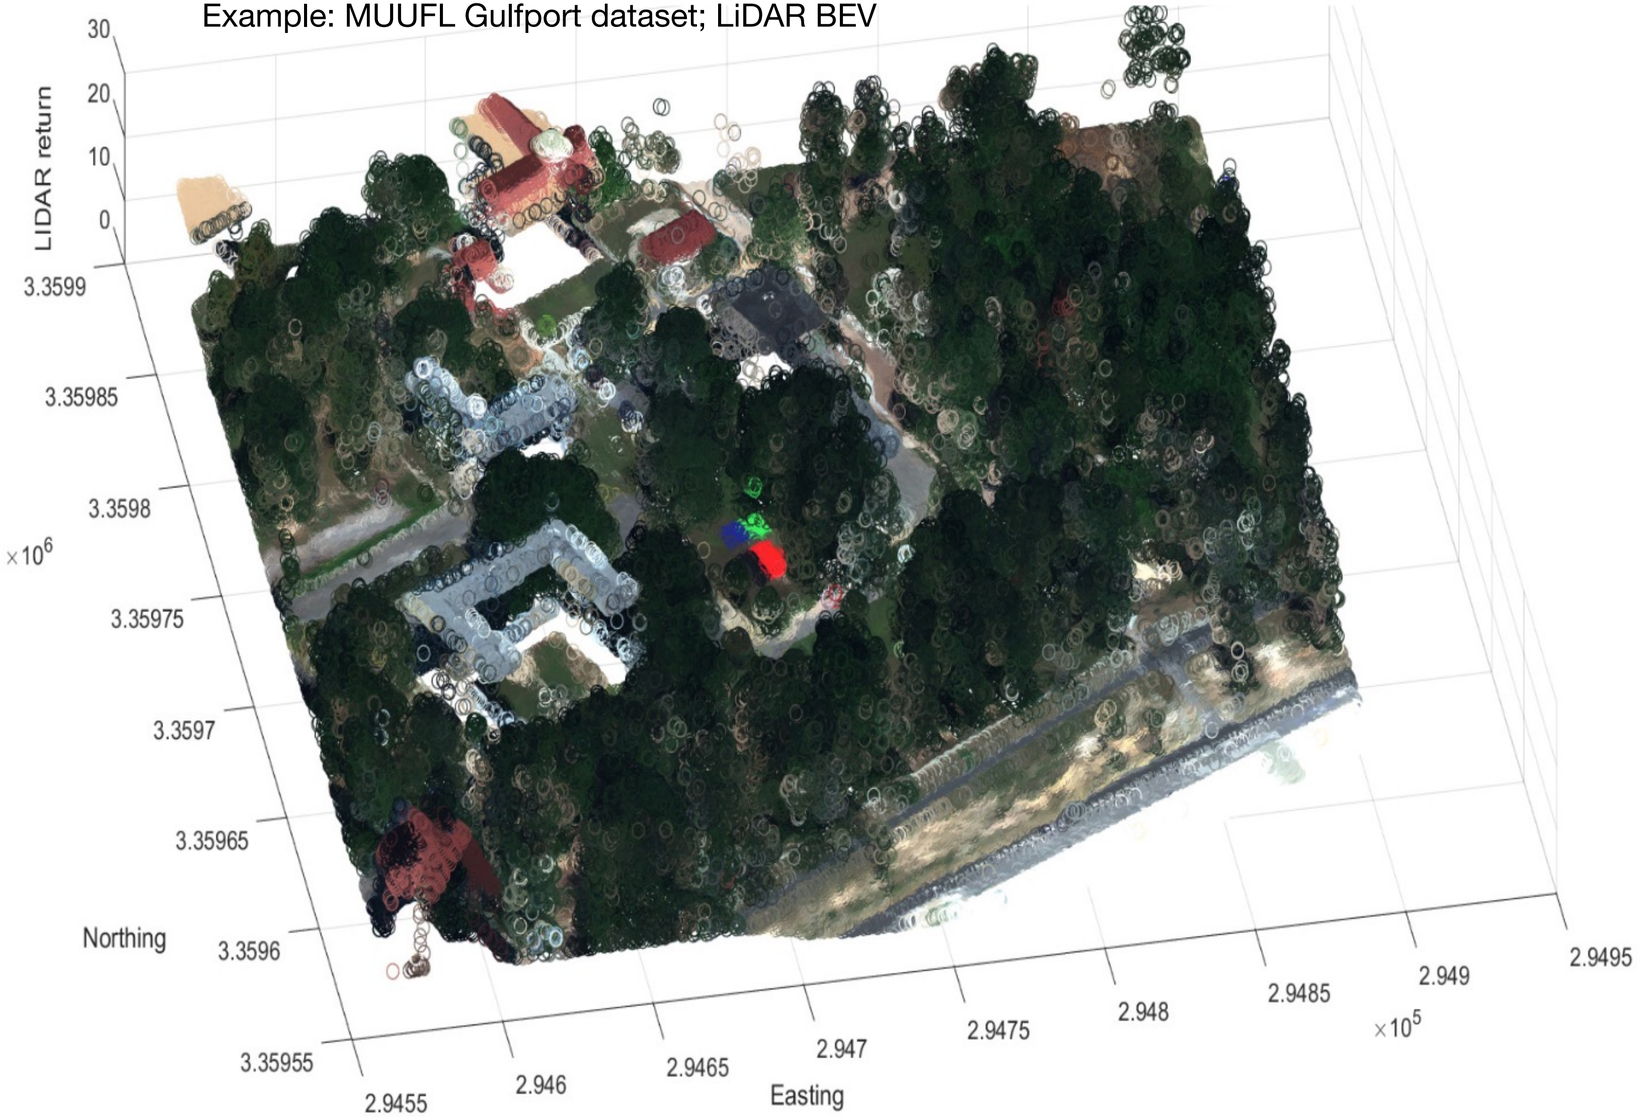

3D representations

Voxels

Represent a shape with a V x V x V grid of occupancies (in 3D!)

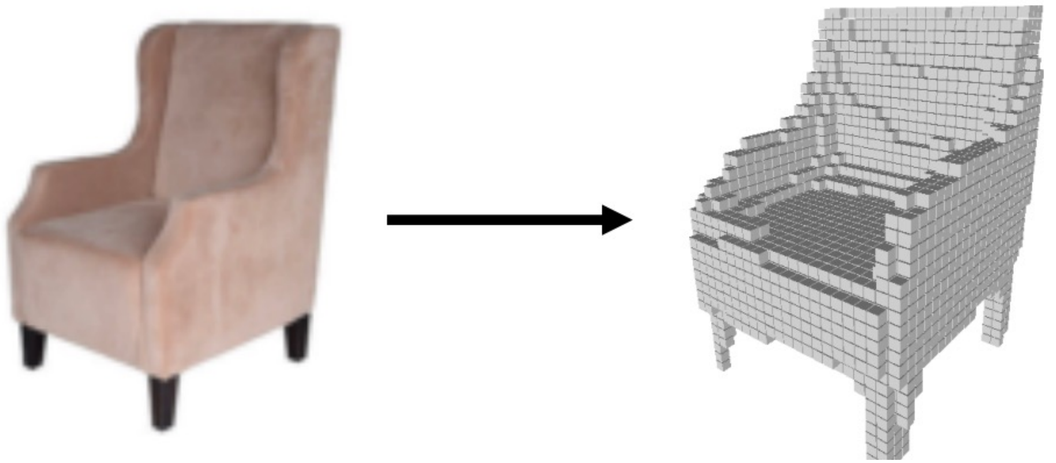

Limitations:

-Need high spatial resolution to capture fine structures

-Scaling to high resolutions is nontrivial!

Adapted from Prof. Johnson's slides
Choy et al., "3D-R2N2: A Unified Approach for Single and Multi-view 3D Object Reconstruction",

Processing Voxel Inputs

- 3D Convolution

torch.nn.functional.conv3d(*input*, *weight*, *bias=None*, *stride=1*, *padding=0*, *dilation=1*, *groups=1*) → Tensor

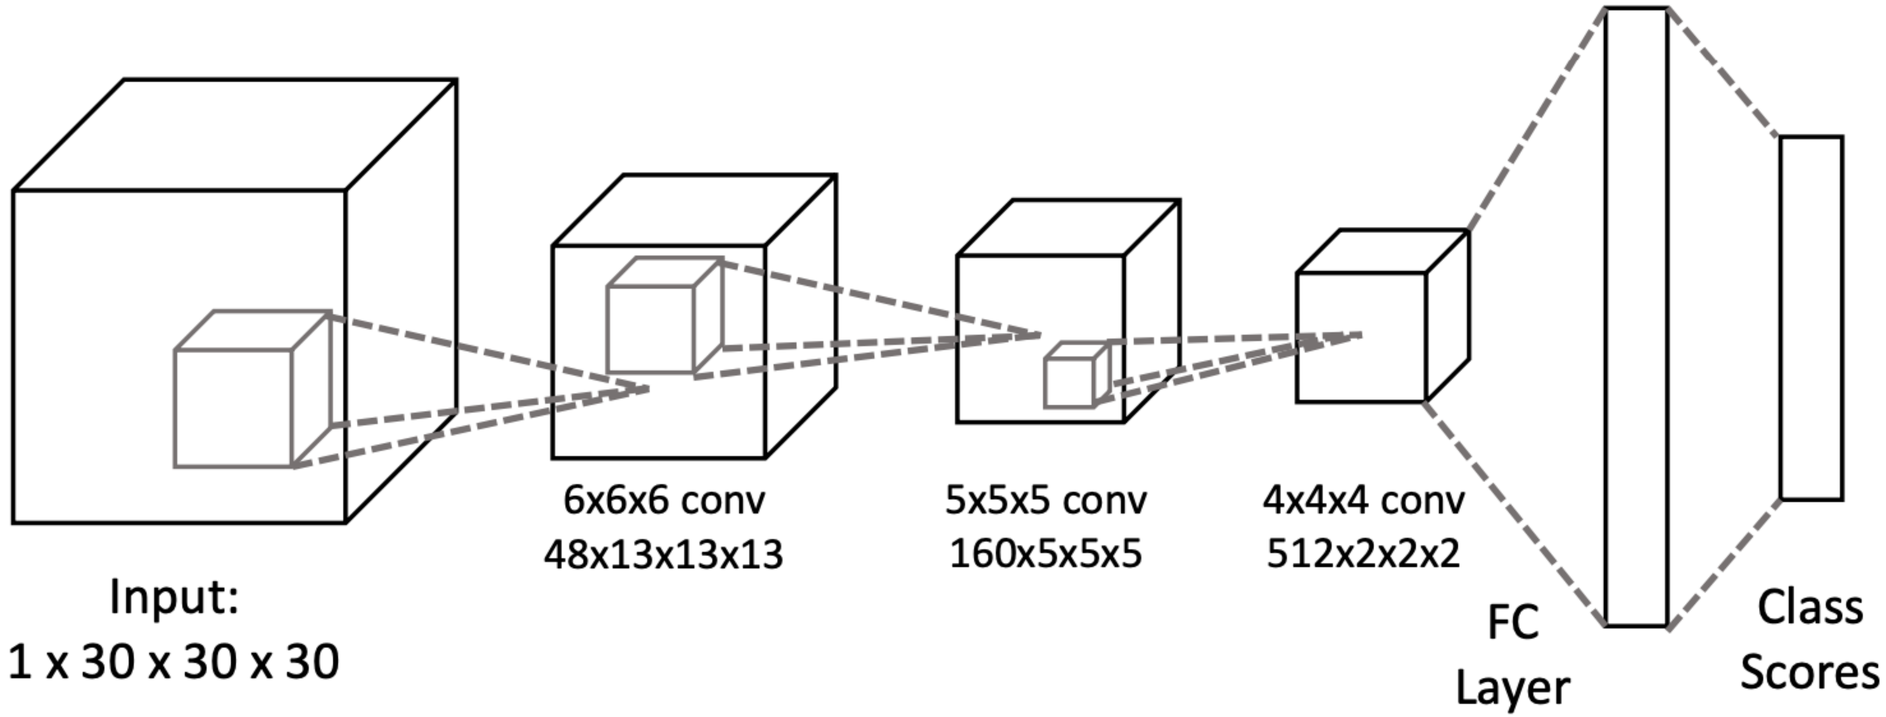

PointNet

- LiDAR-based 3D perception
- Input: Point Cloud (set of points)

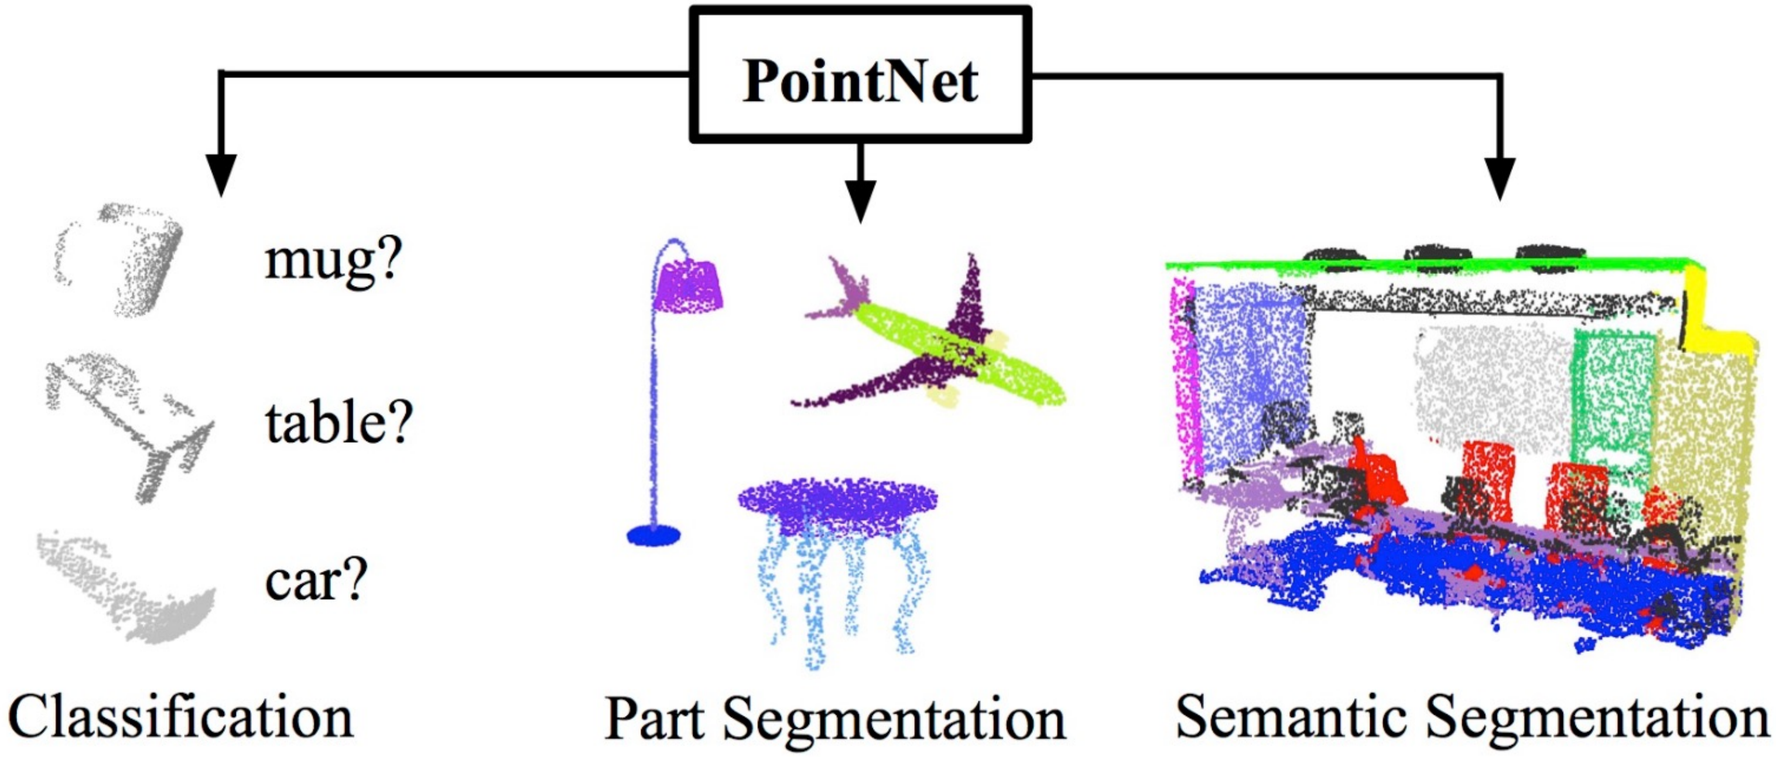

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet

*Order does not matter!
*Invariance & Interaction

- Pros:
  - Only using a set of points as input (no need for voxelization or rendering)
  - Can represent fine structures without huge numbers of points

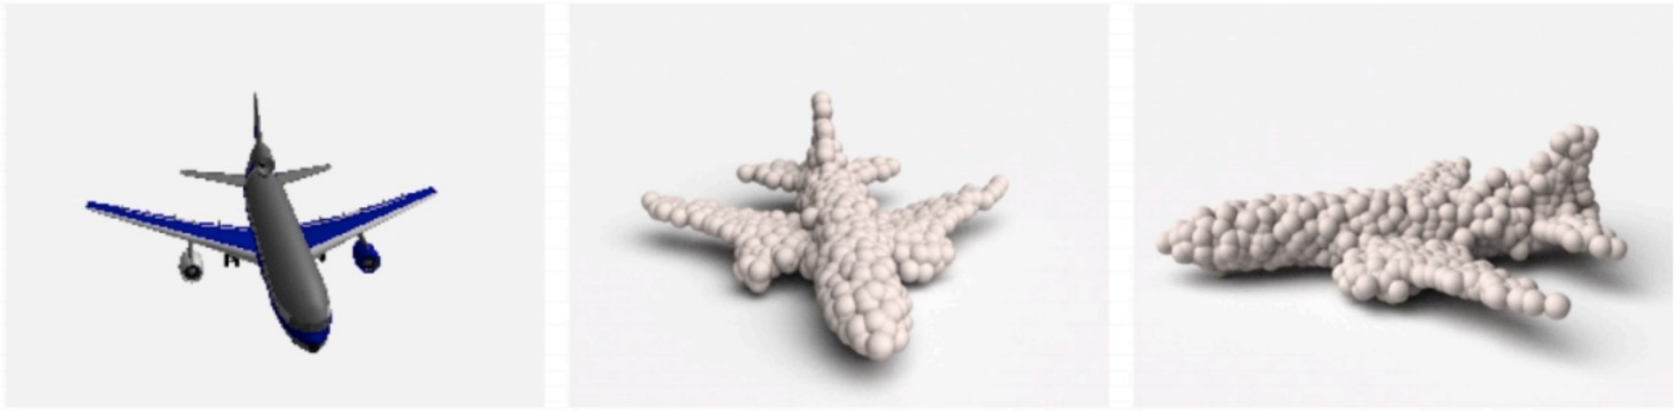

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet

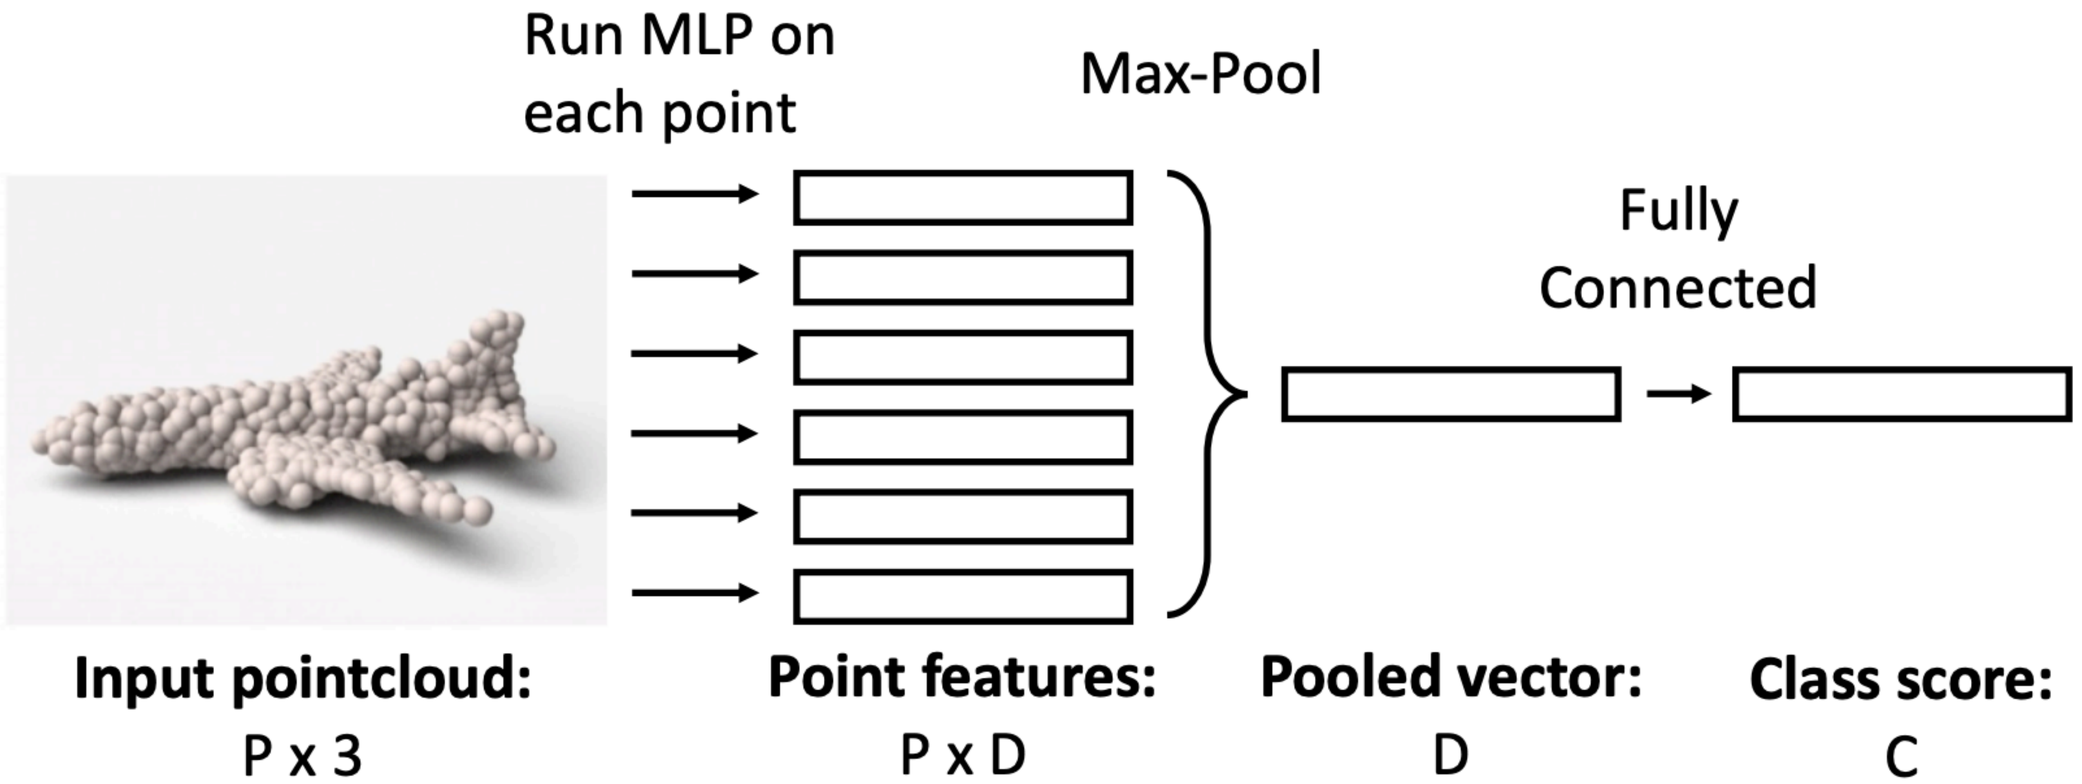

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet

- Architecture

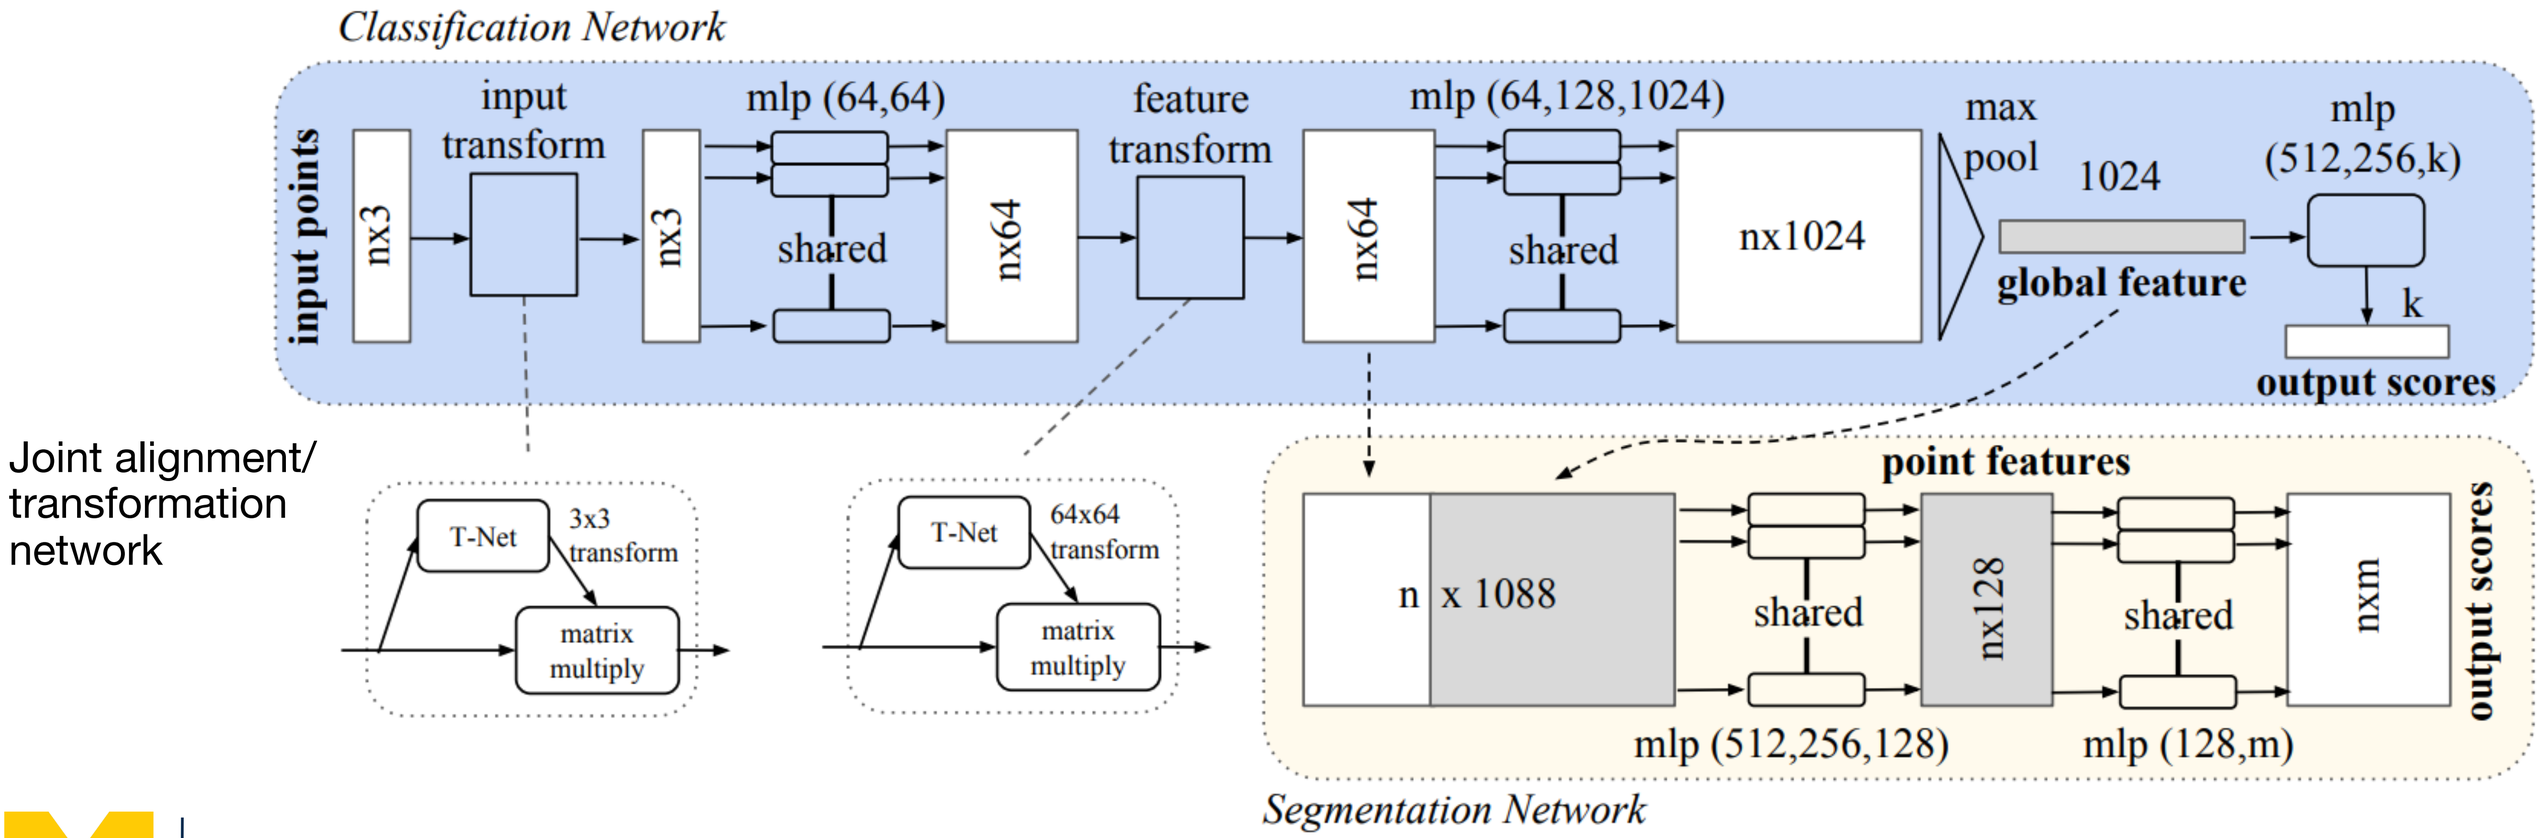

Joint alignment/ transformation network

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet

- T-Net

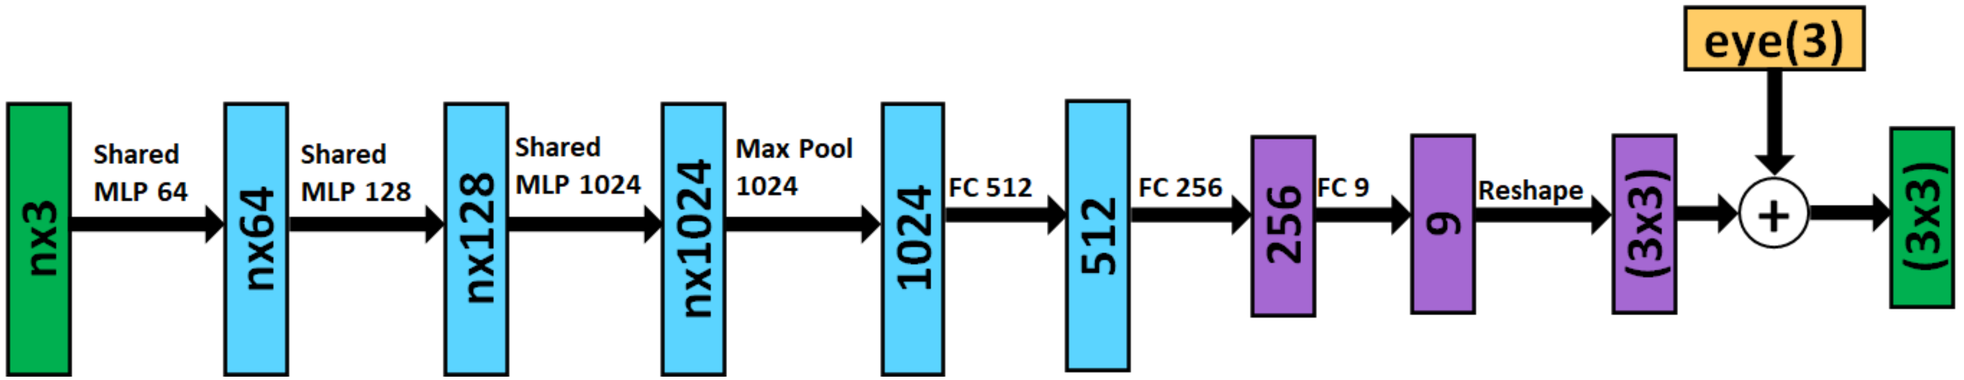

Joint alignment/ transformation network

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet

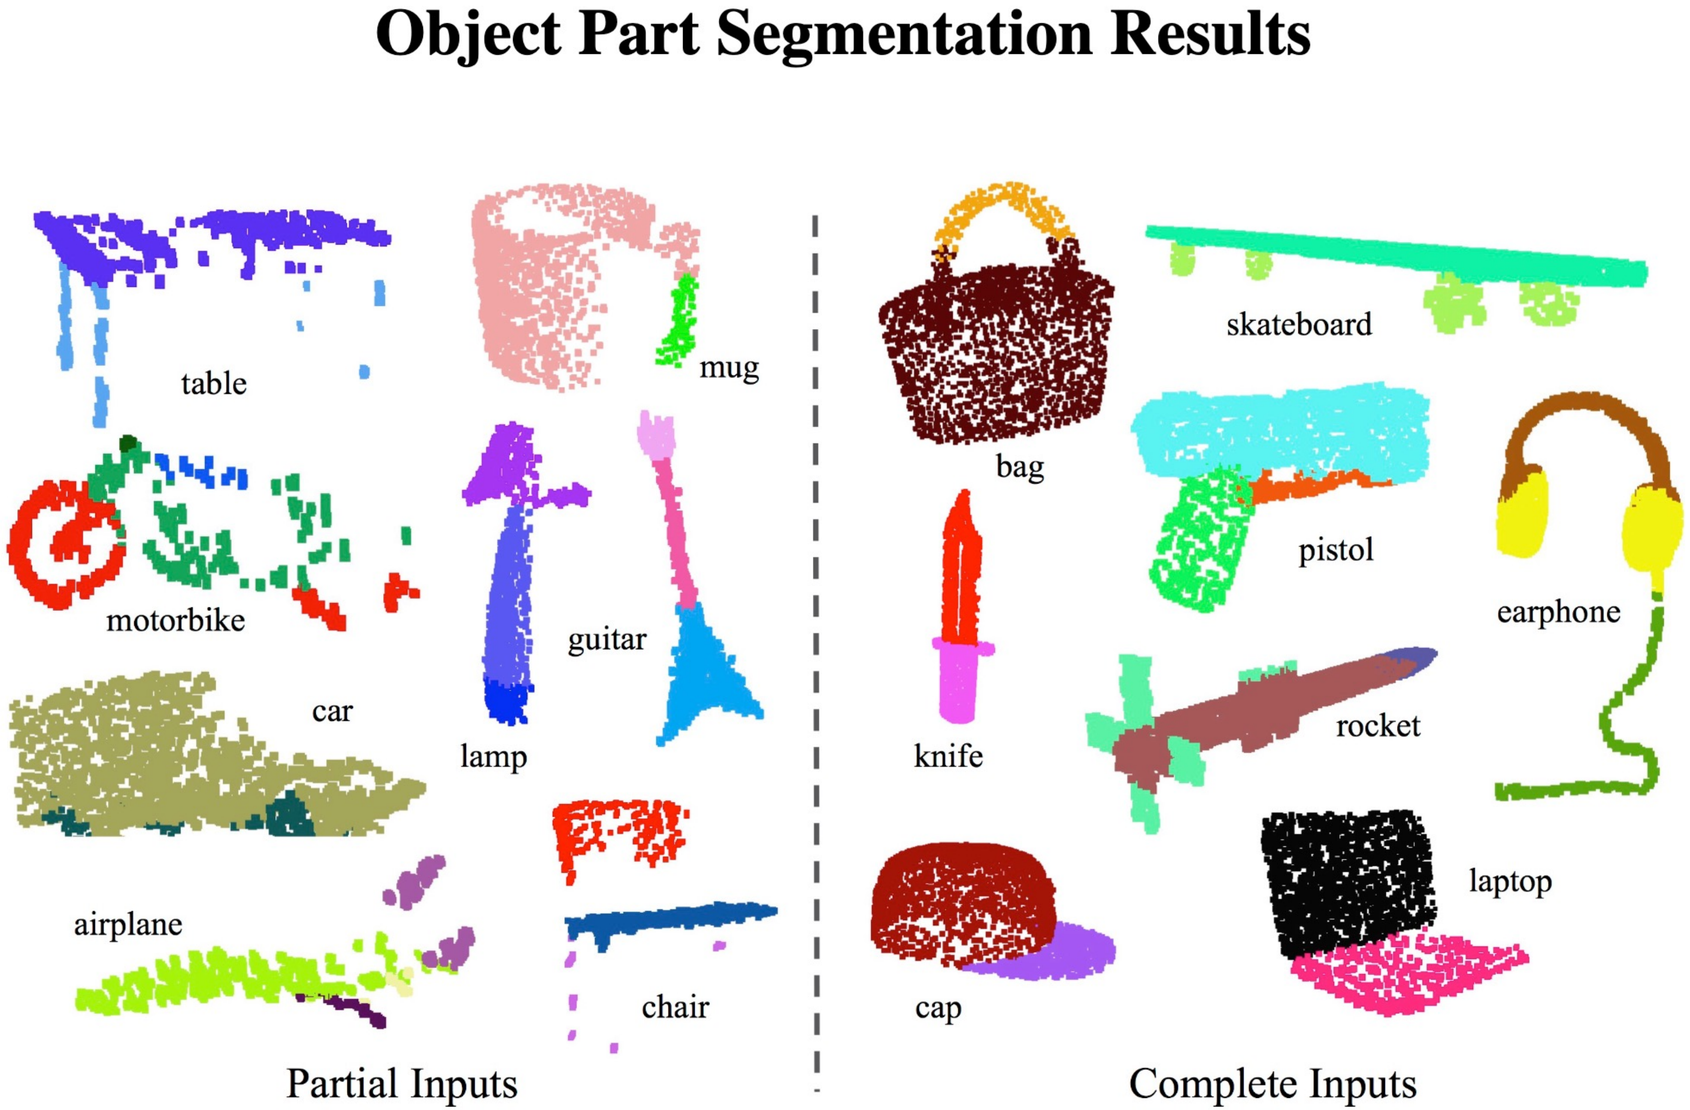

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

PointNet - semantic segmentation results

Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.

Loss function for point clouds

Chamfer distance is the sum of L2 distance to each point's nearest neighbor in the other set

$$d_{CD}(S_1, S_2) = \sum_{x \in S_1} \min_{y \in S_2} \|x - y\|_2^2 + \sum_{y \in S_2} \min_{x \in S_1} \|x - y\|_2^2$$

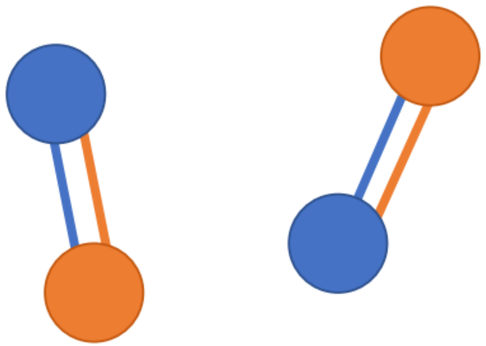

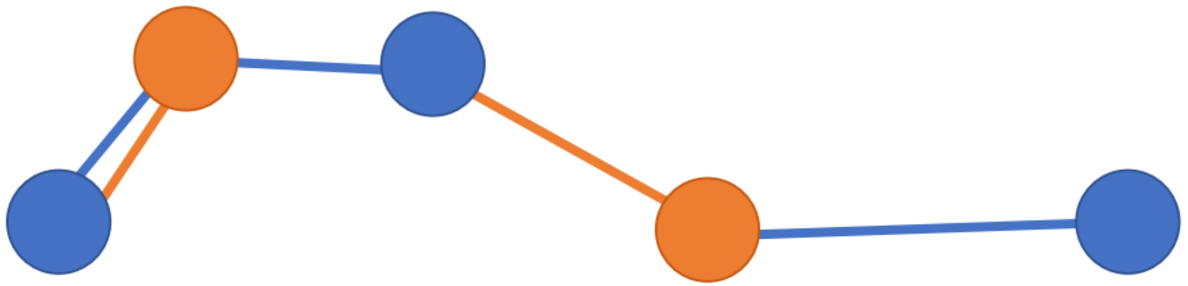

Fan et al, "A Point Set Generation Network for 3D Object Reconstruction from a Single Image", CVPR 2017

DETR 3D

- Camera-based 3D perception (BEV)
- Input: 2D Camera images

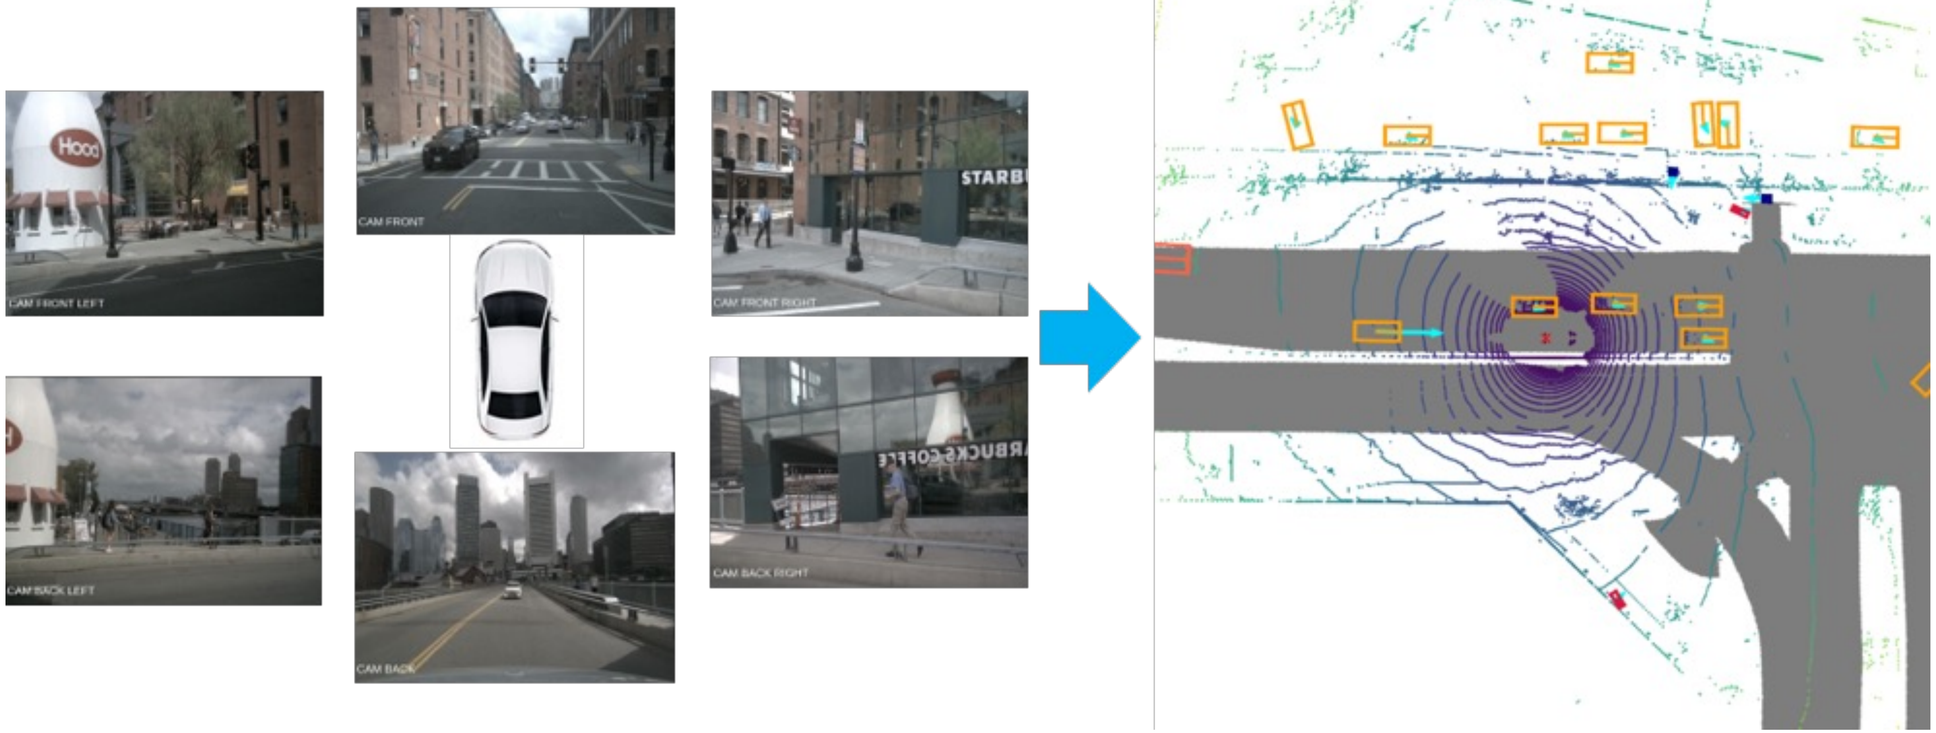

Wang., et al. "DETR3D: 3D Object Detection from Multi-view Images via 3D-to-2D Queries"., CoRL 2021.

DETR 3D

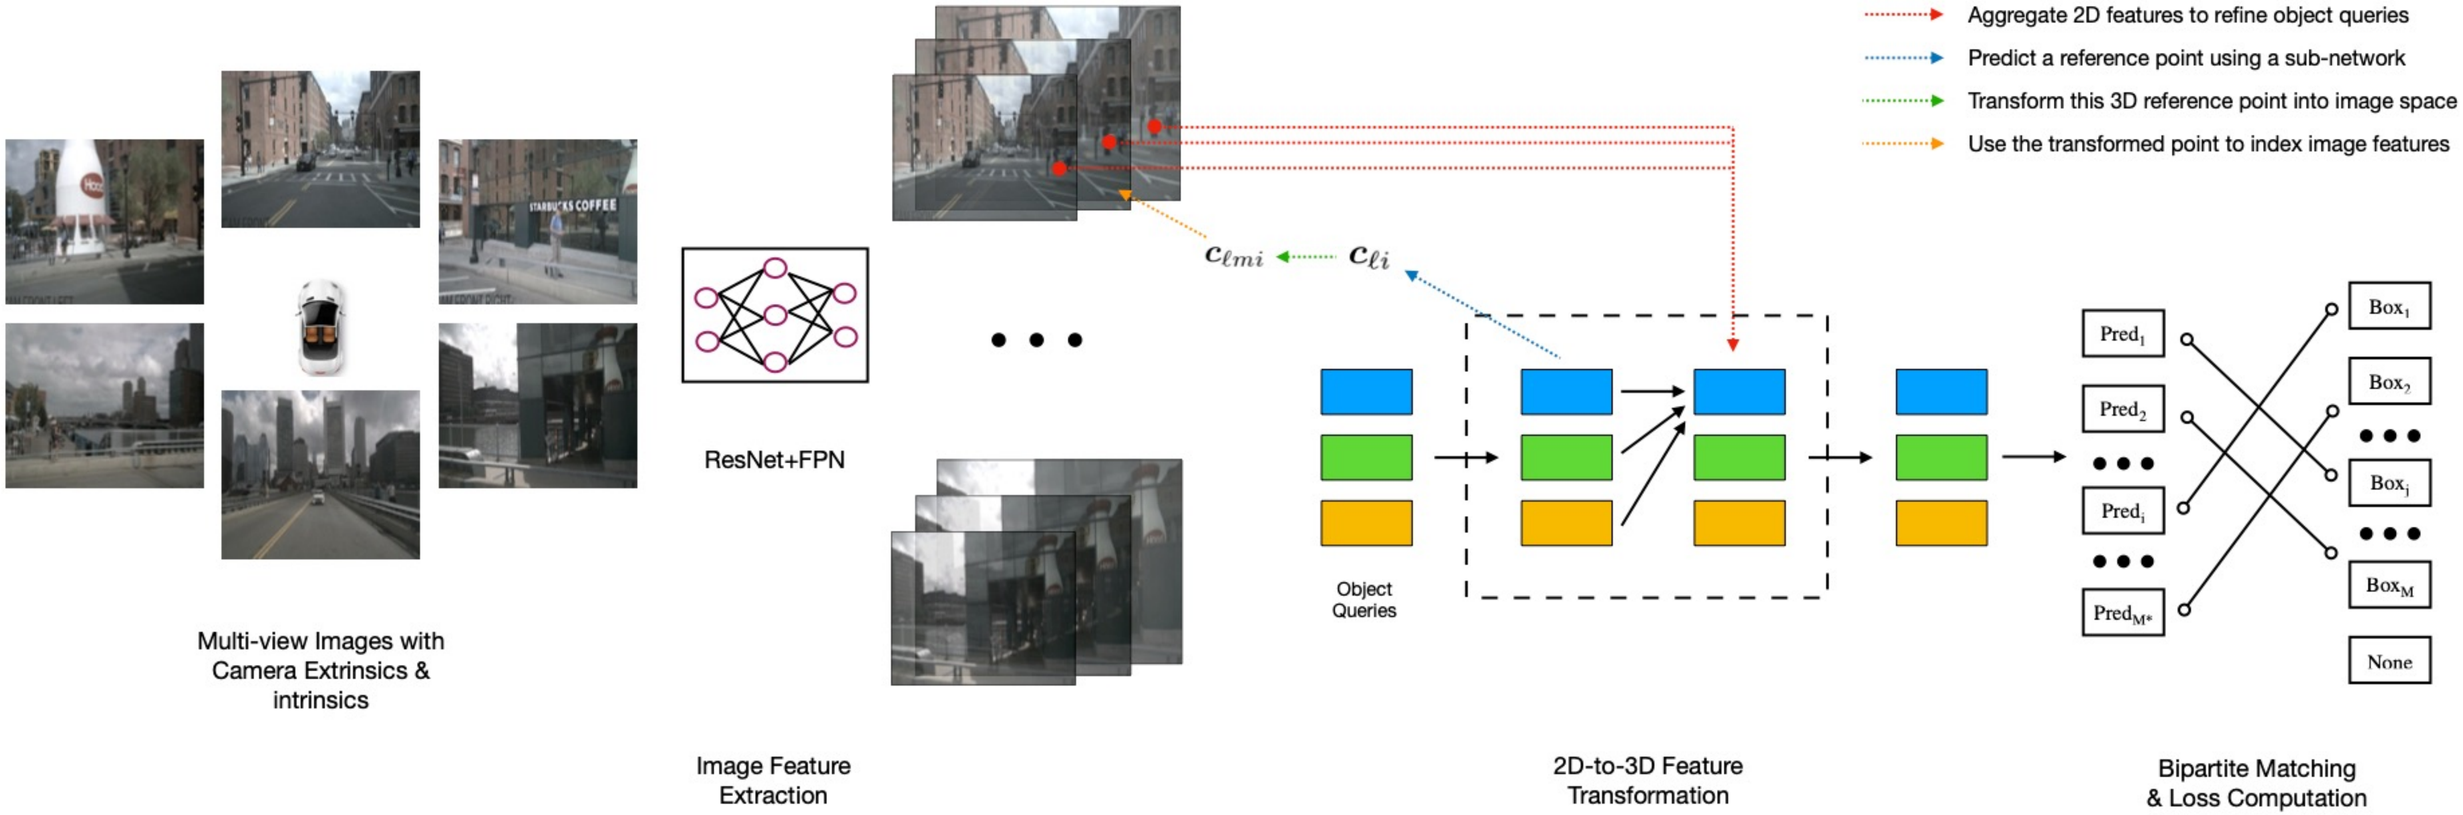

Wang., et al. "DETR3D: 3D Object Detection from Multi-view Images via 3D-to-2D Queries"., CoRL 2021.

Loss function for point clouds

Chamfer distance is the sum of L2 distance to each point's nearest neighbor in the other set

$$d_{CD}(S_1, S_2) = \sum_{x \in S_1} \min_{y \in S_2} \|x - y\|_2^2 + \sum_{y \in S_2} \min_{x \in S_1} \|x - y\|_2^2$$

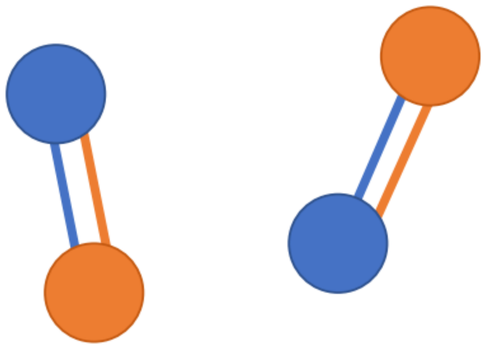

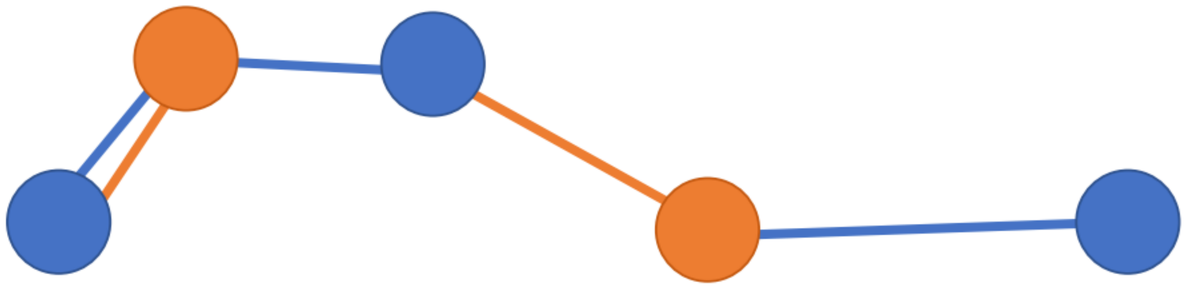

Fan et al, "A Point Set Generation Network for 3D Object Reconstruction from a Single Image", CVPR 2017

Many more topics in 3D Vision

- Multi-view stereo
- Structure from Motion
- Simultaneous Localization and Mapping (SLAM)
- View Synthesis
- Differentiable graphics
- 3D Sensors; multi-modal sensor fusion
- Non DL methods in 3DV

DeepRob

Lecture 19
3D Perception and Point Clouds
University of Michigan | Department of Robotics

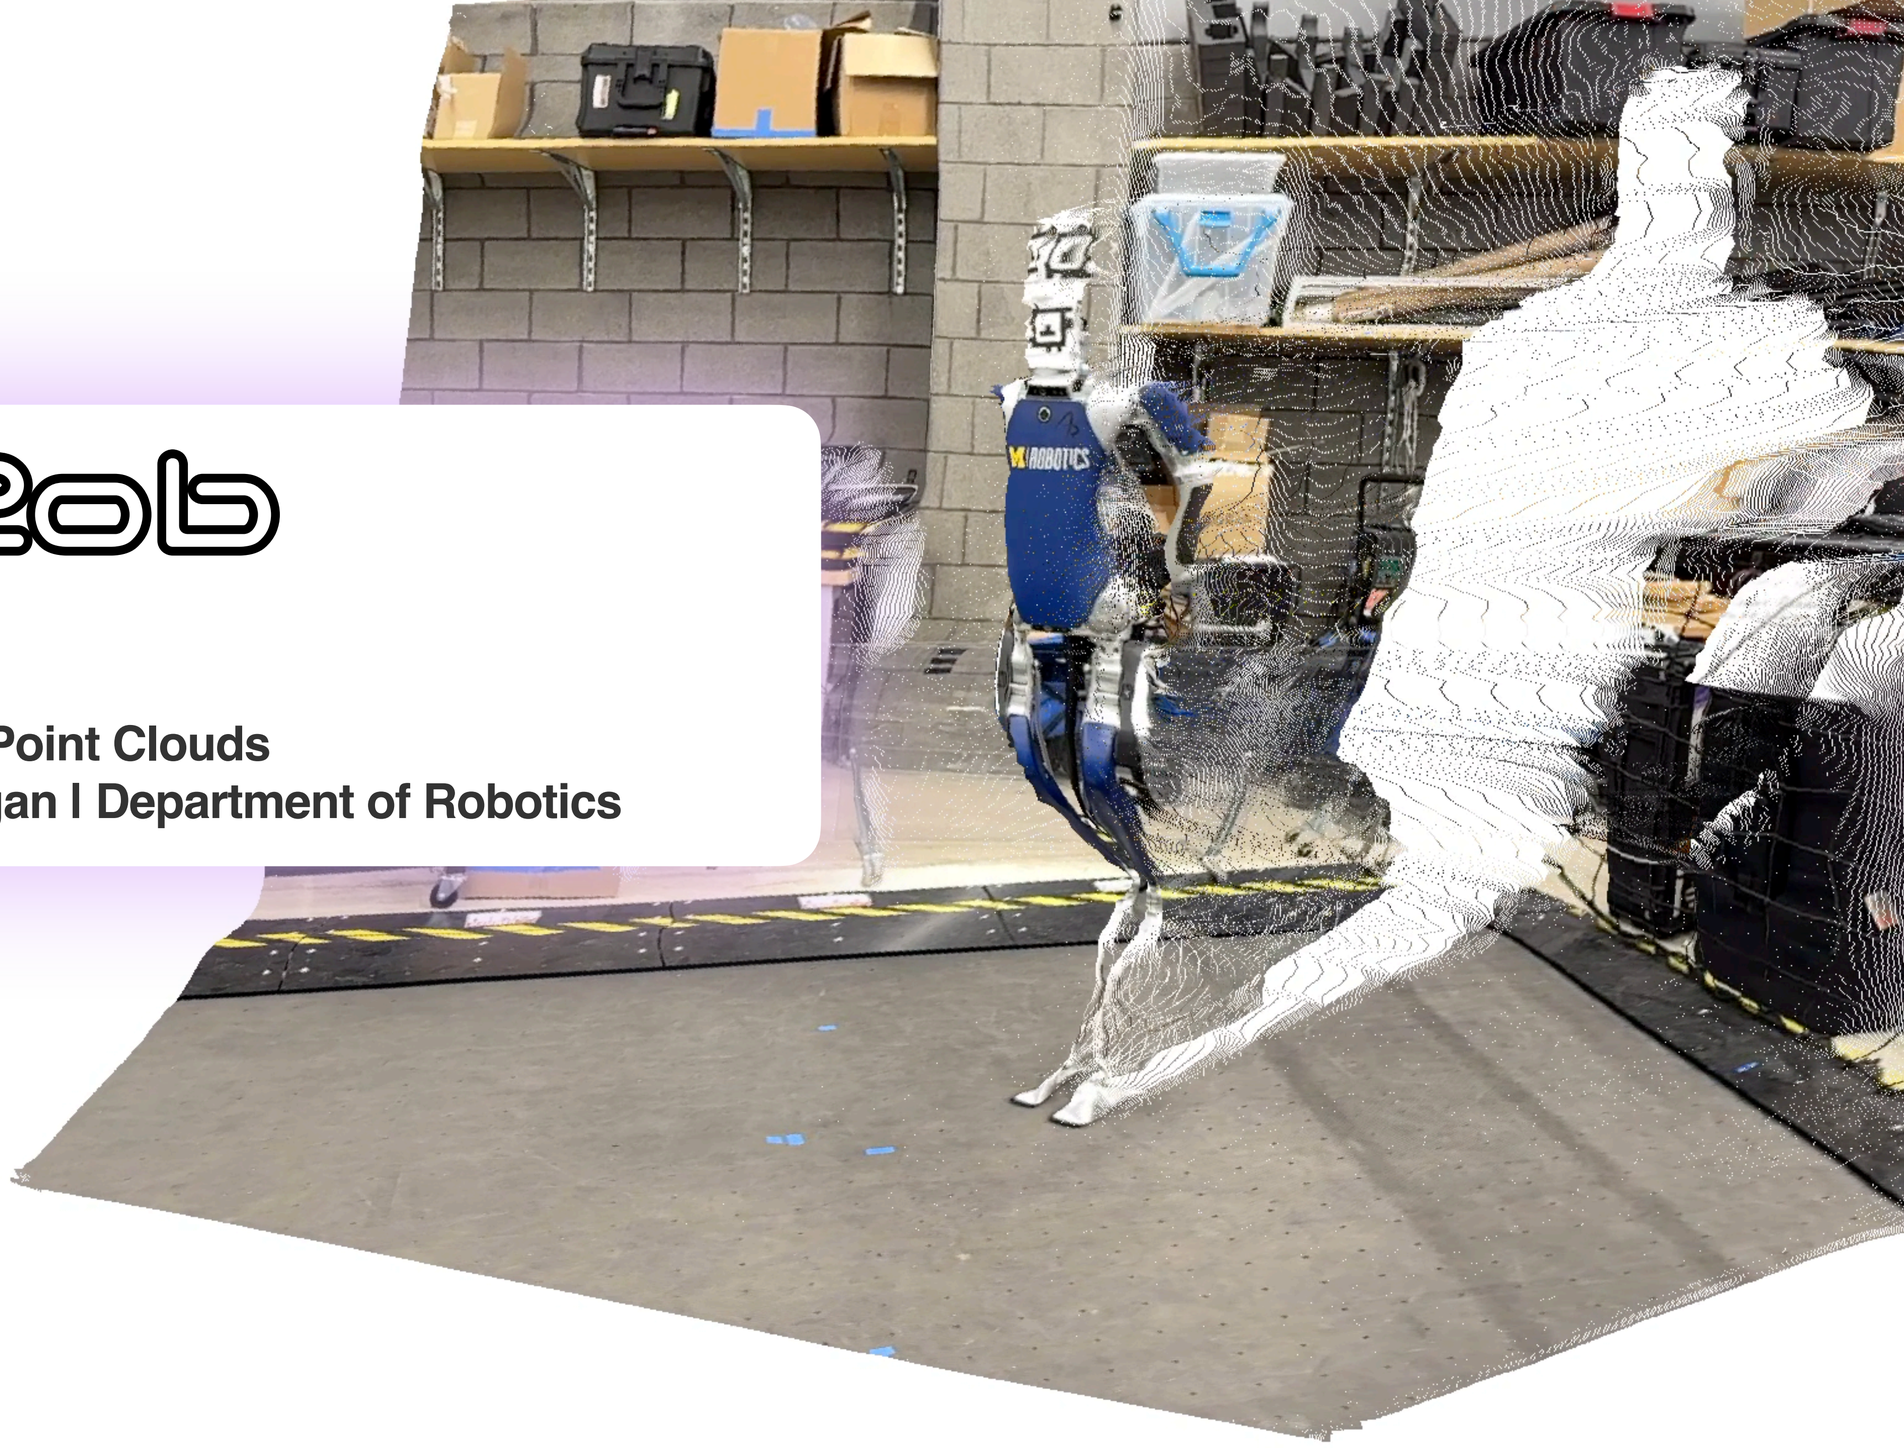

DeepRob