

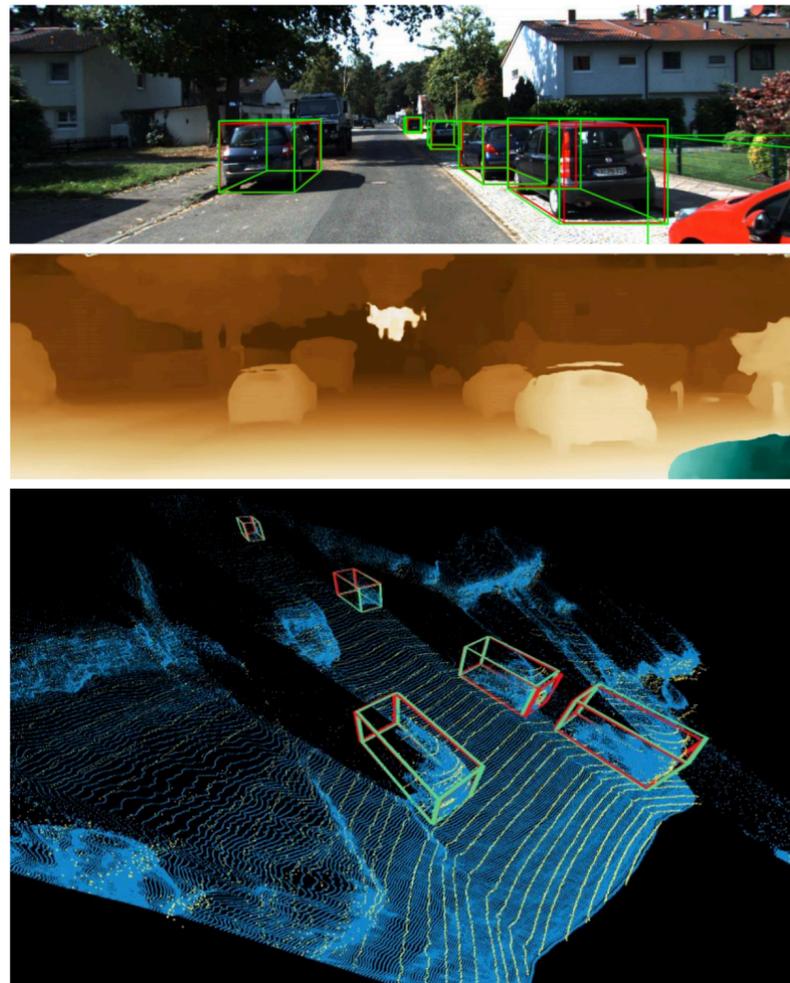
Efficient 3D Deep Learning for Particle Physics

Zhijian Liu (zhijian@mit.edu)

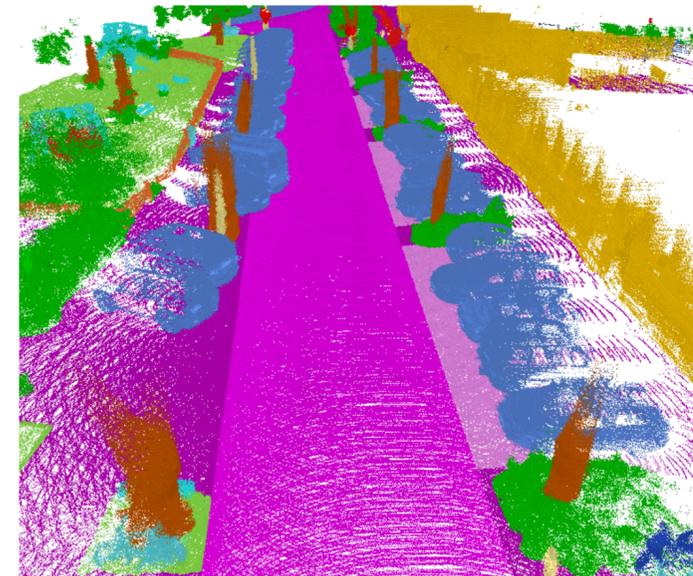
Massachusetts Institute of Technology

Collaborators: Alex Schuy, Jeff Krupa, Thomas Klijnsma, Phil Harris, Shih-Chieh Hsu, Nhan Tran

3D Deep Learning: Applications



Autonomous Driving

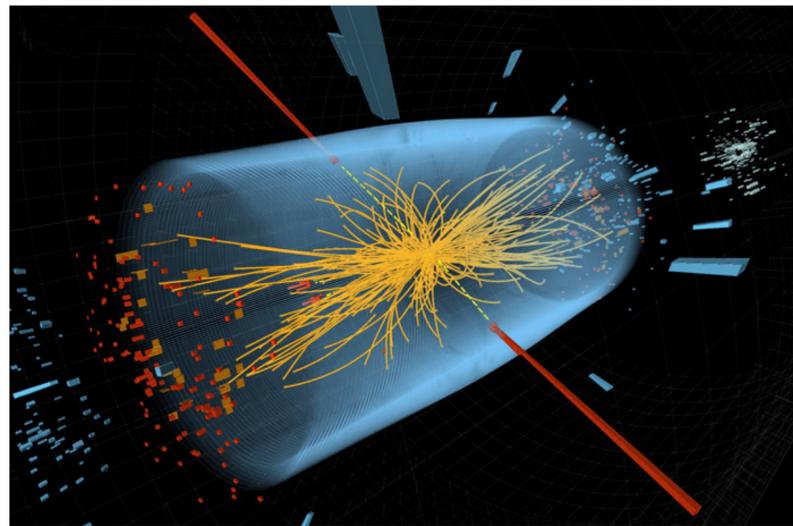


Scene Segmentation

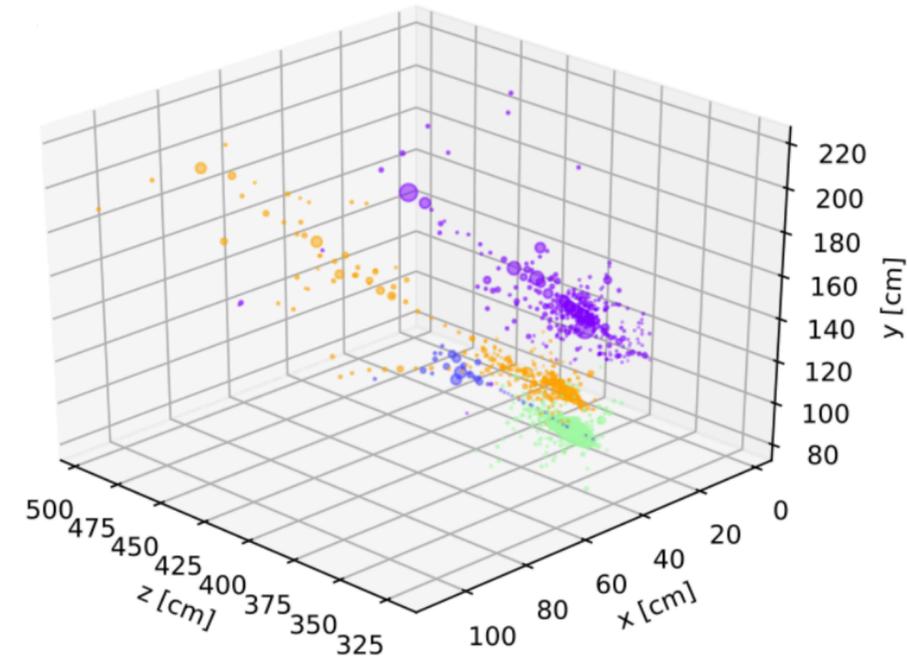
Input: (x, y, z, I)

Output: **classes**

3D Deep Learning: Applications



High-Energy Physics

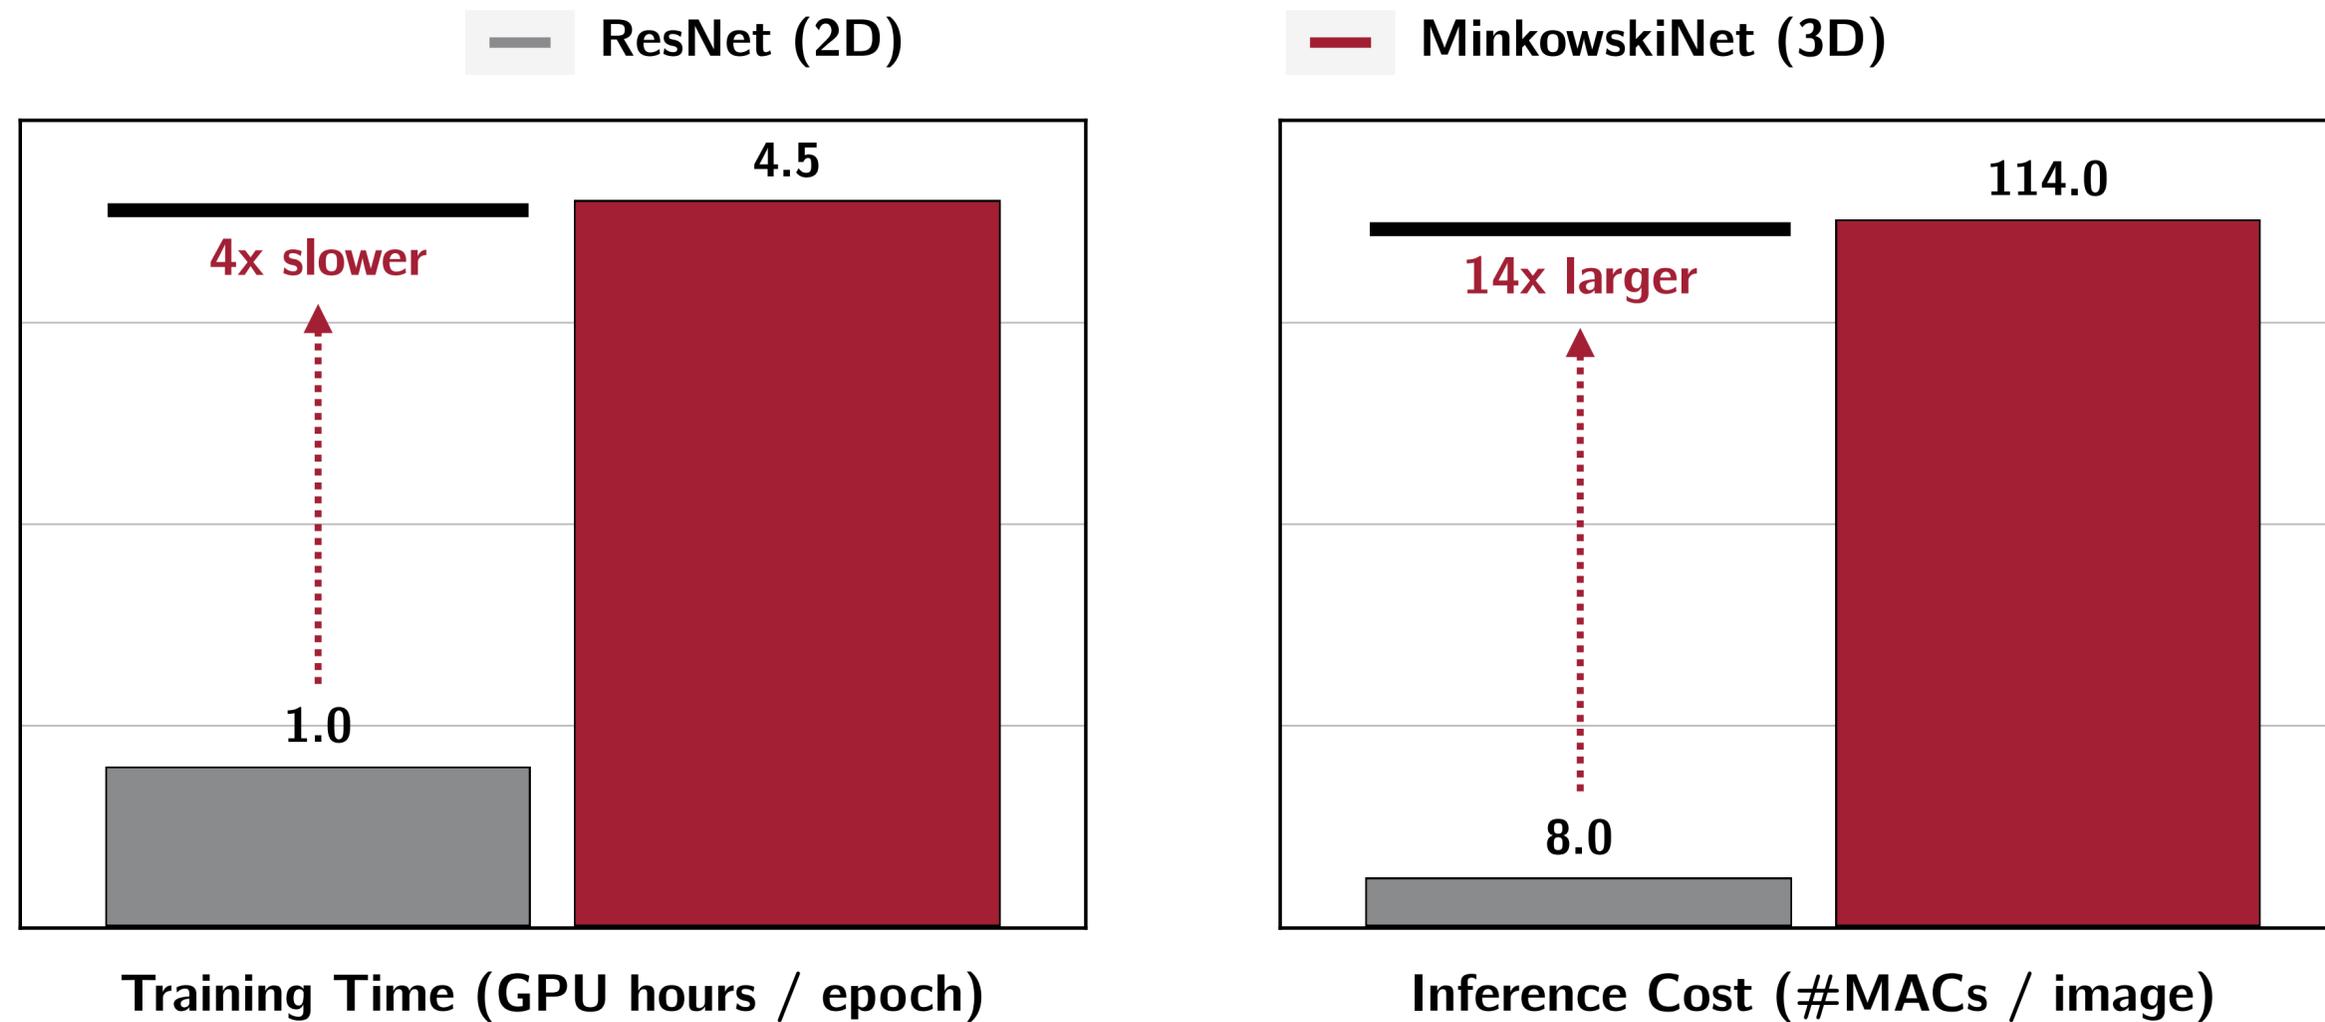


Particle Segmentation

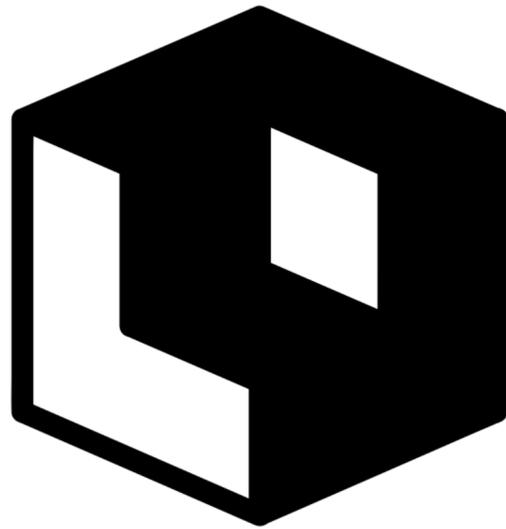
Input: (x, y, z, E, t)

Output: **particle types**

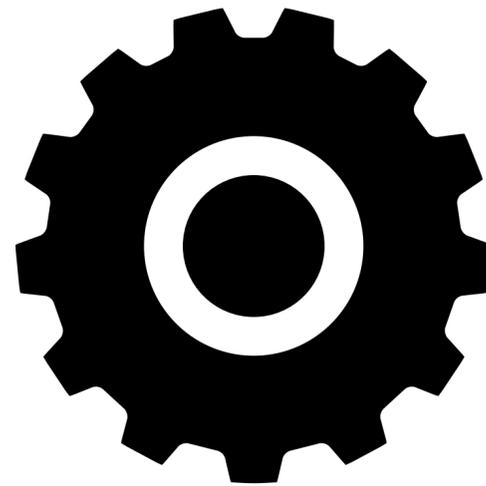
Efficient 3D Deep Learning: Challenges



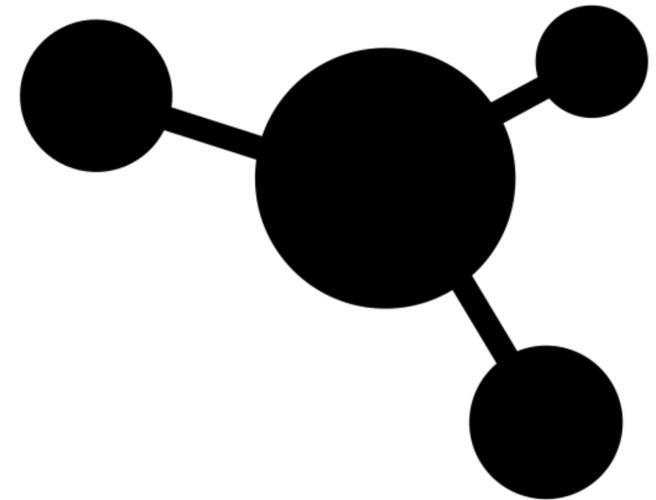
Efficient 3D Deep Learning: Roadmap



3D Building Block

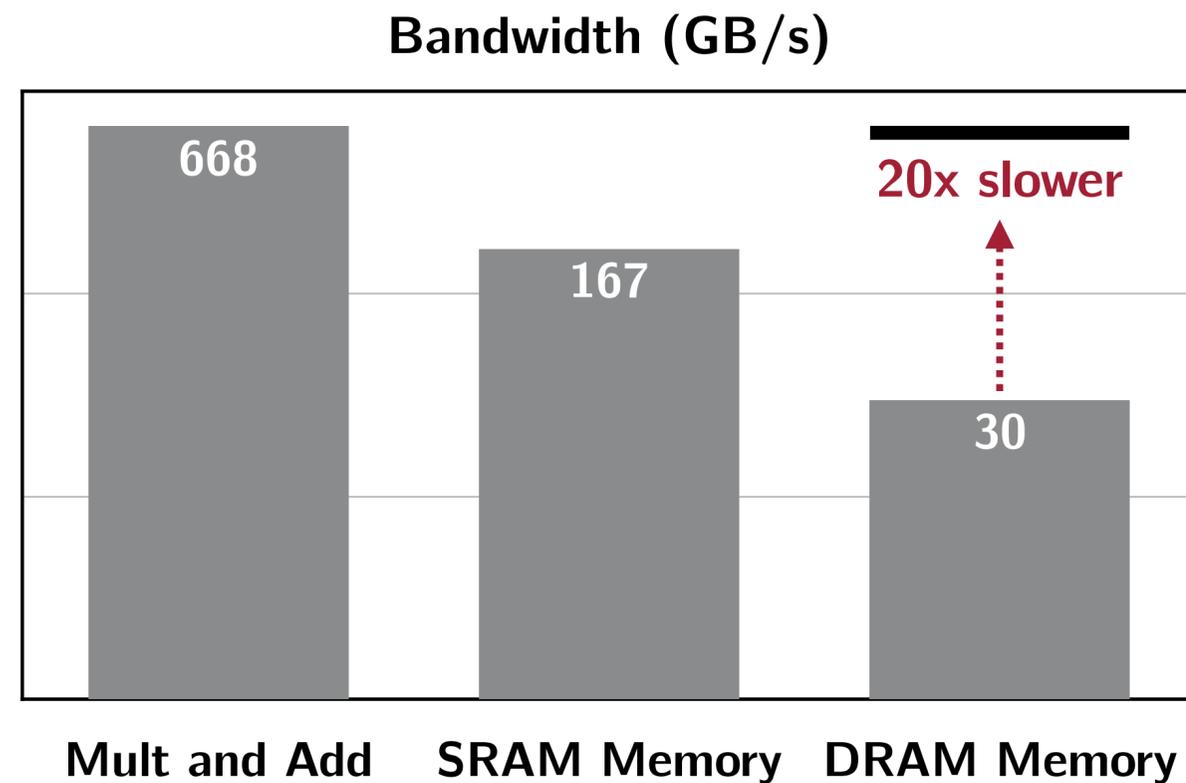


3D Kernel Optimization



3D Network Design

Efficient 3D Deep Learning: Bottlenecks



Off-chip DRAM access is much more expensive than arithmetic operation!



Sequential Memory Access

1	2	3	4	5	6	7	8
---	---	---	---	---	---	---	---

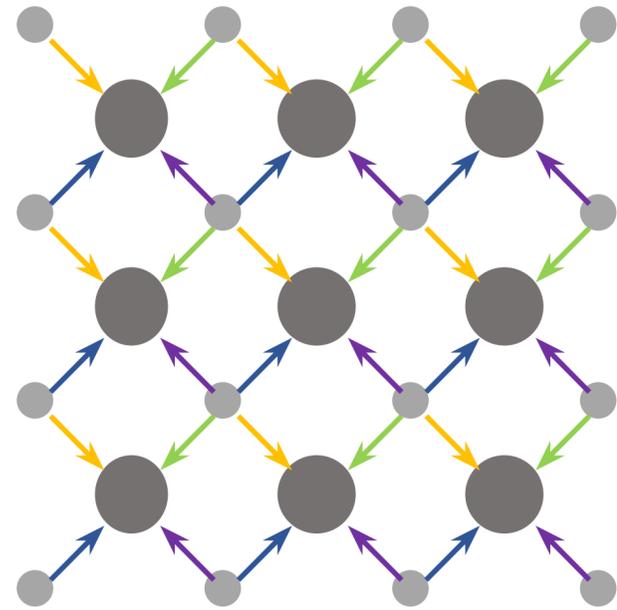


Random Memory Access

7	5	2	4	6	1	8	3
---	---	---	---	---	---	---	---

Random memory access is inefficient due to the potential bank conflicts!

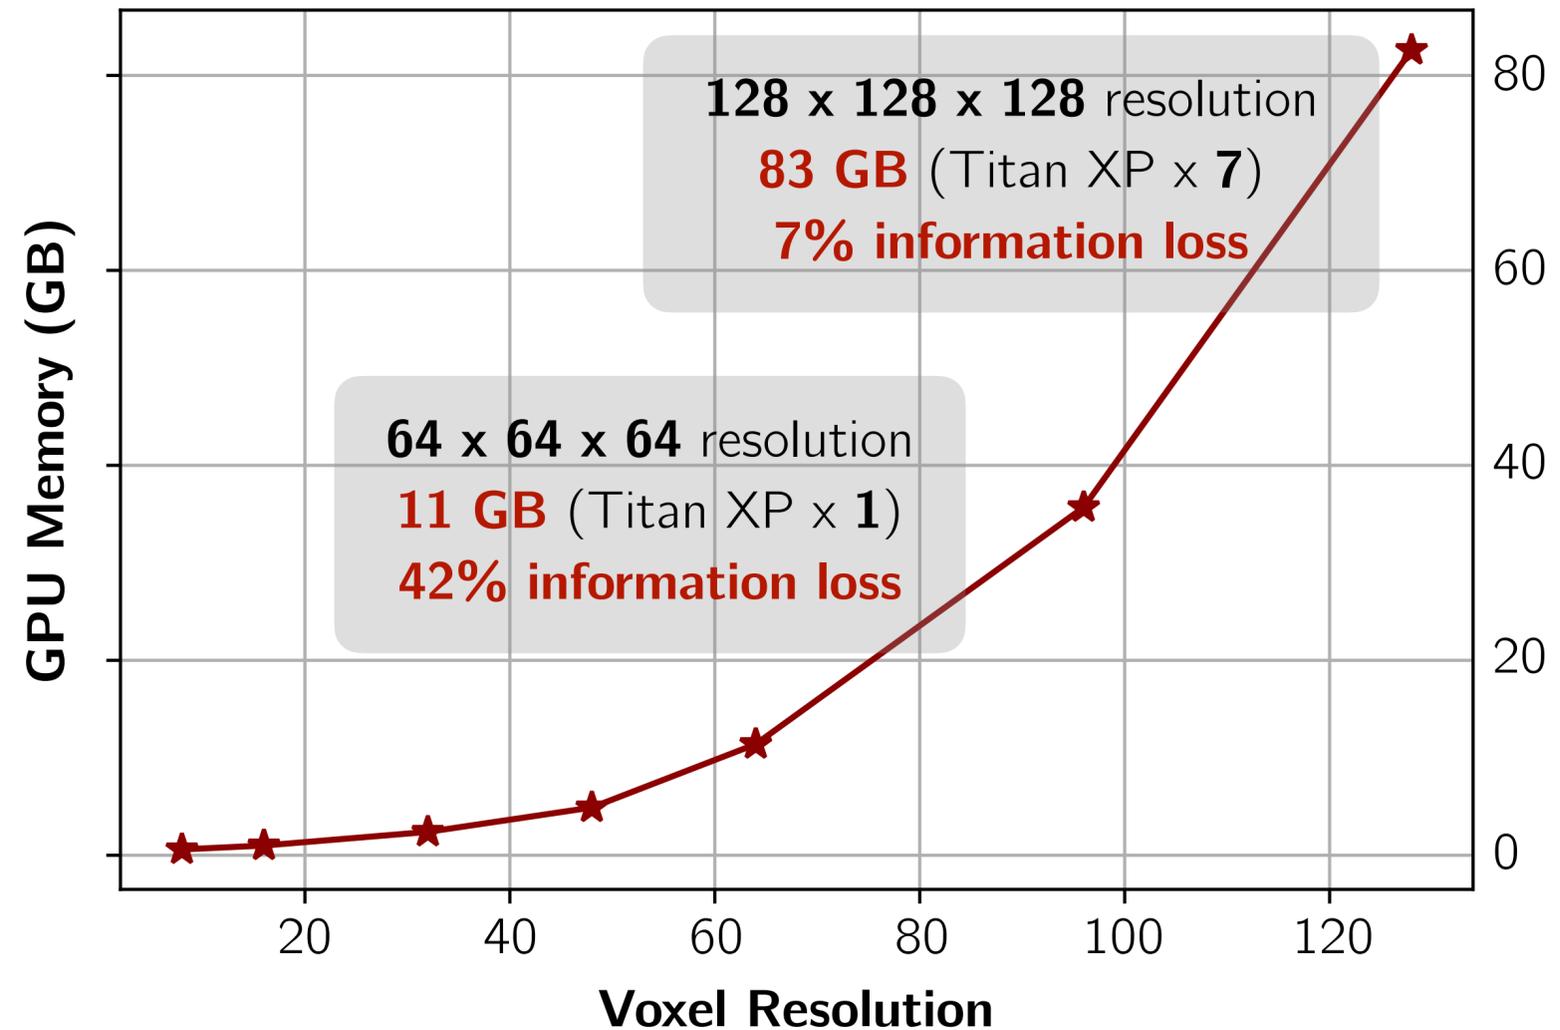
Voxel-Based Models: Cubically-Growing Memory



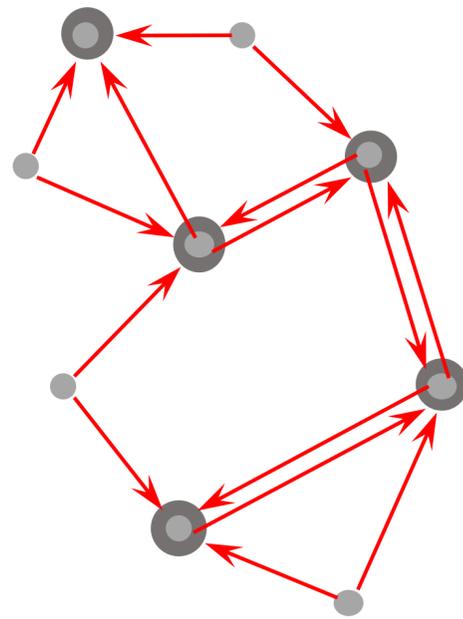
3D ShapeNets [CVPR'15]

VoxNet [IROS'15]

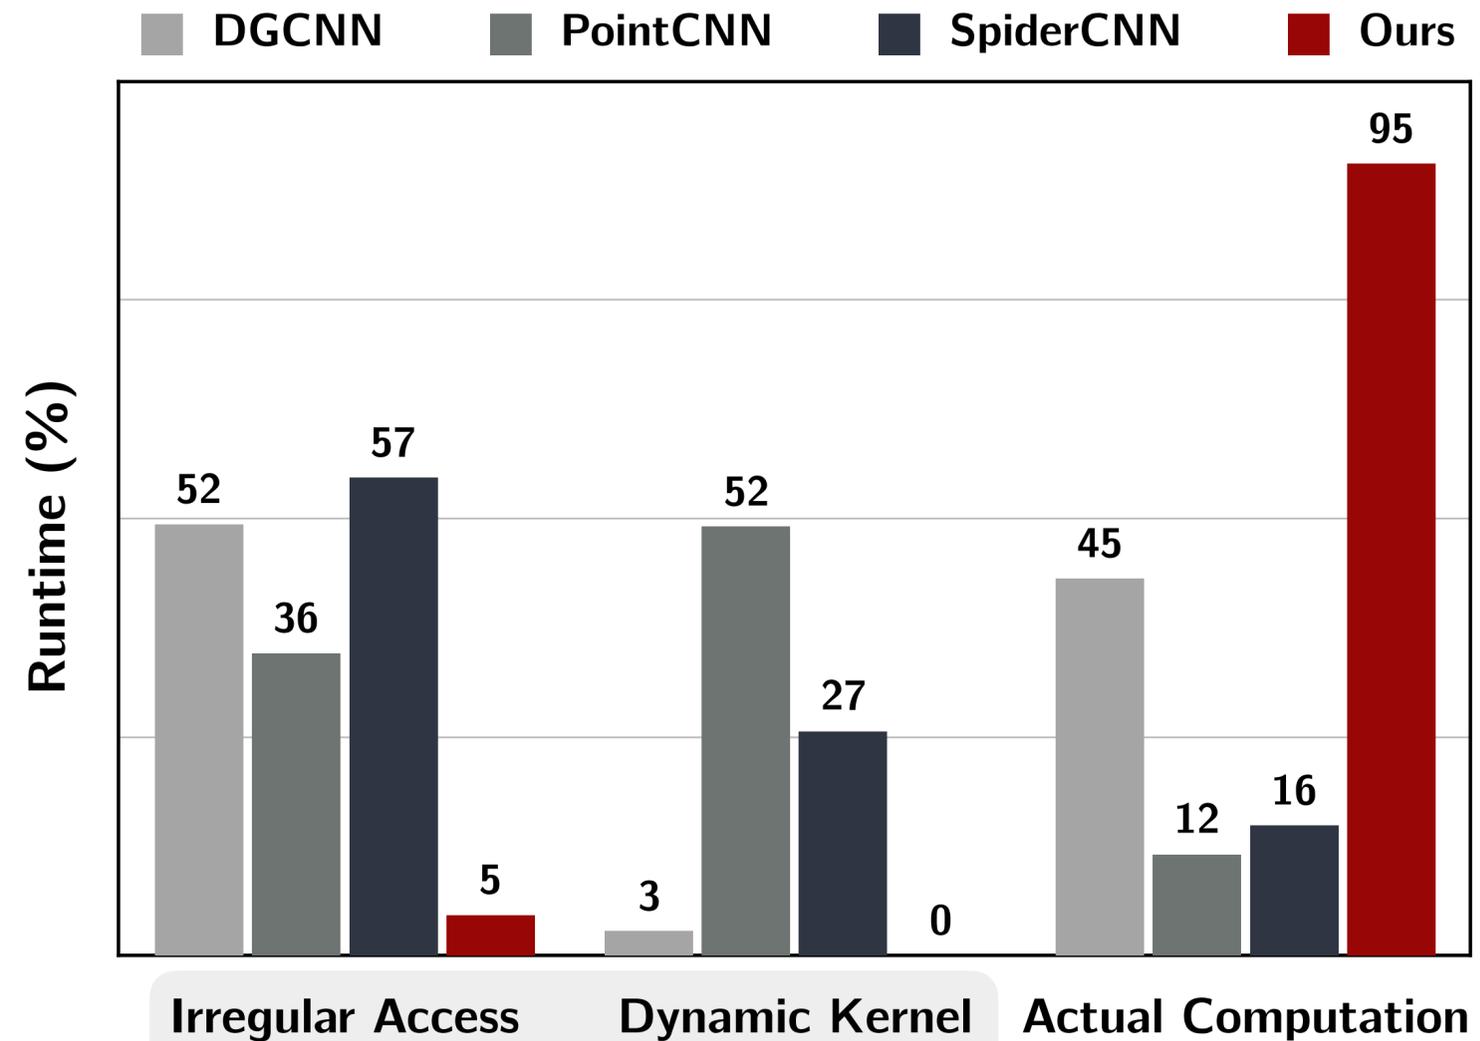
3D U-Net [MICCAI'16]



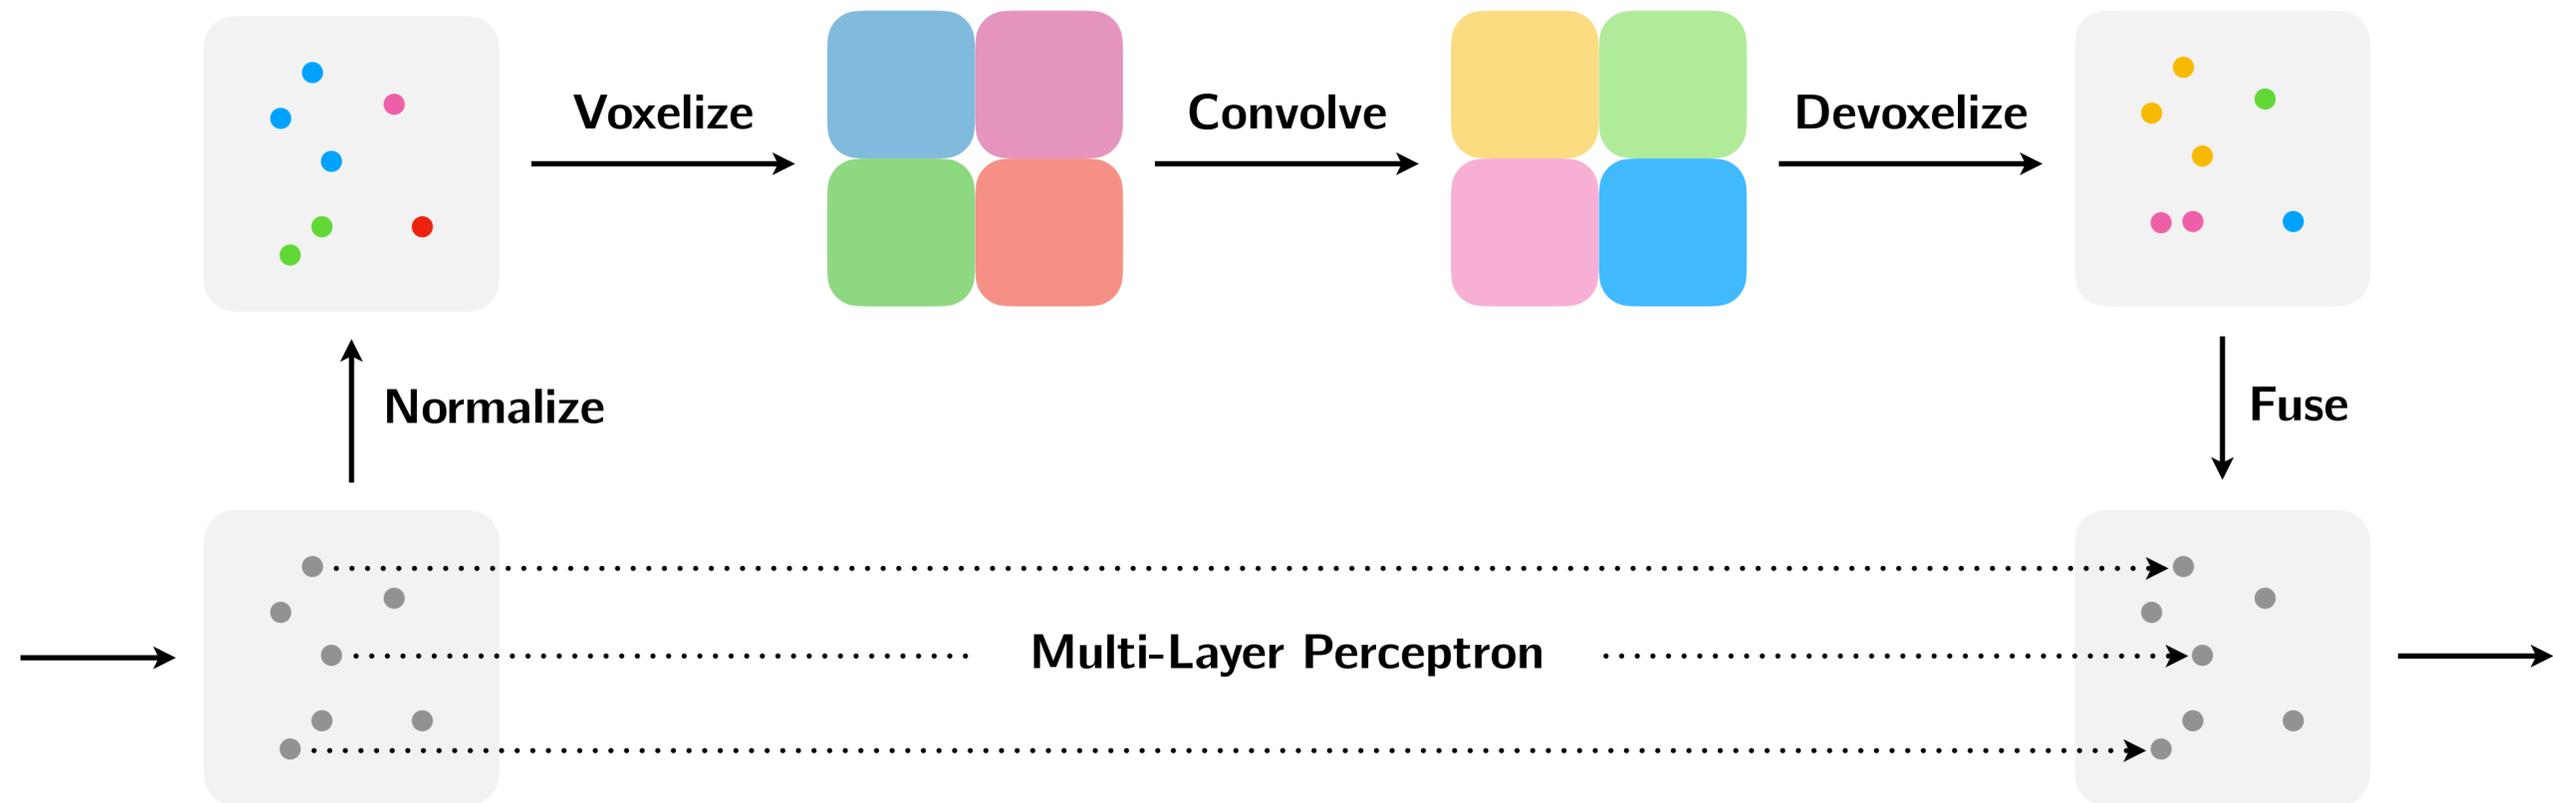
Point-Based Models: Sparsity Overheads



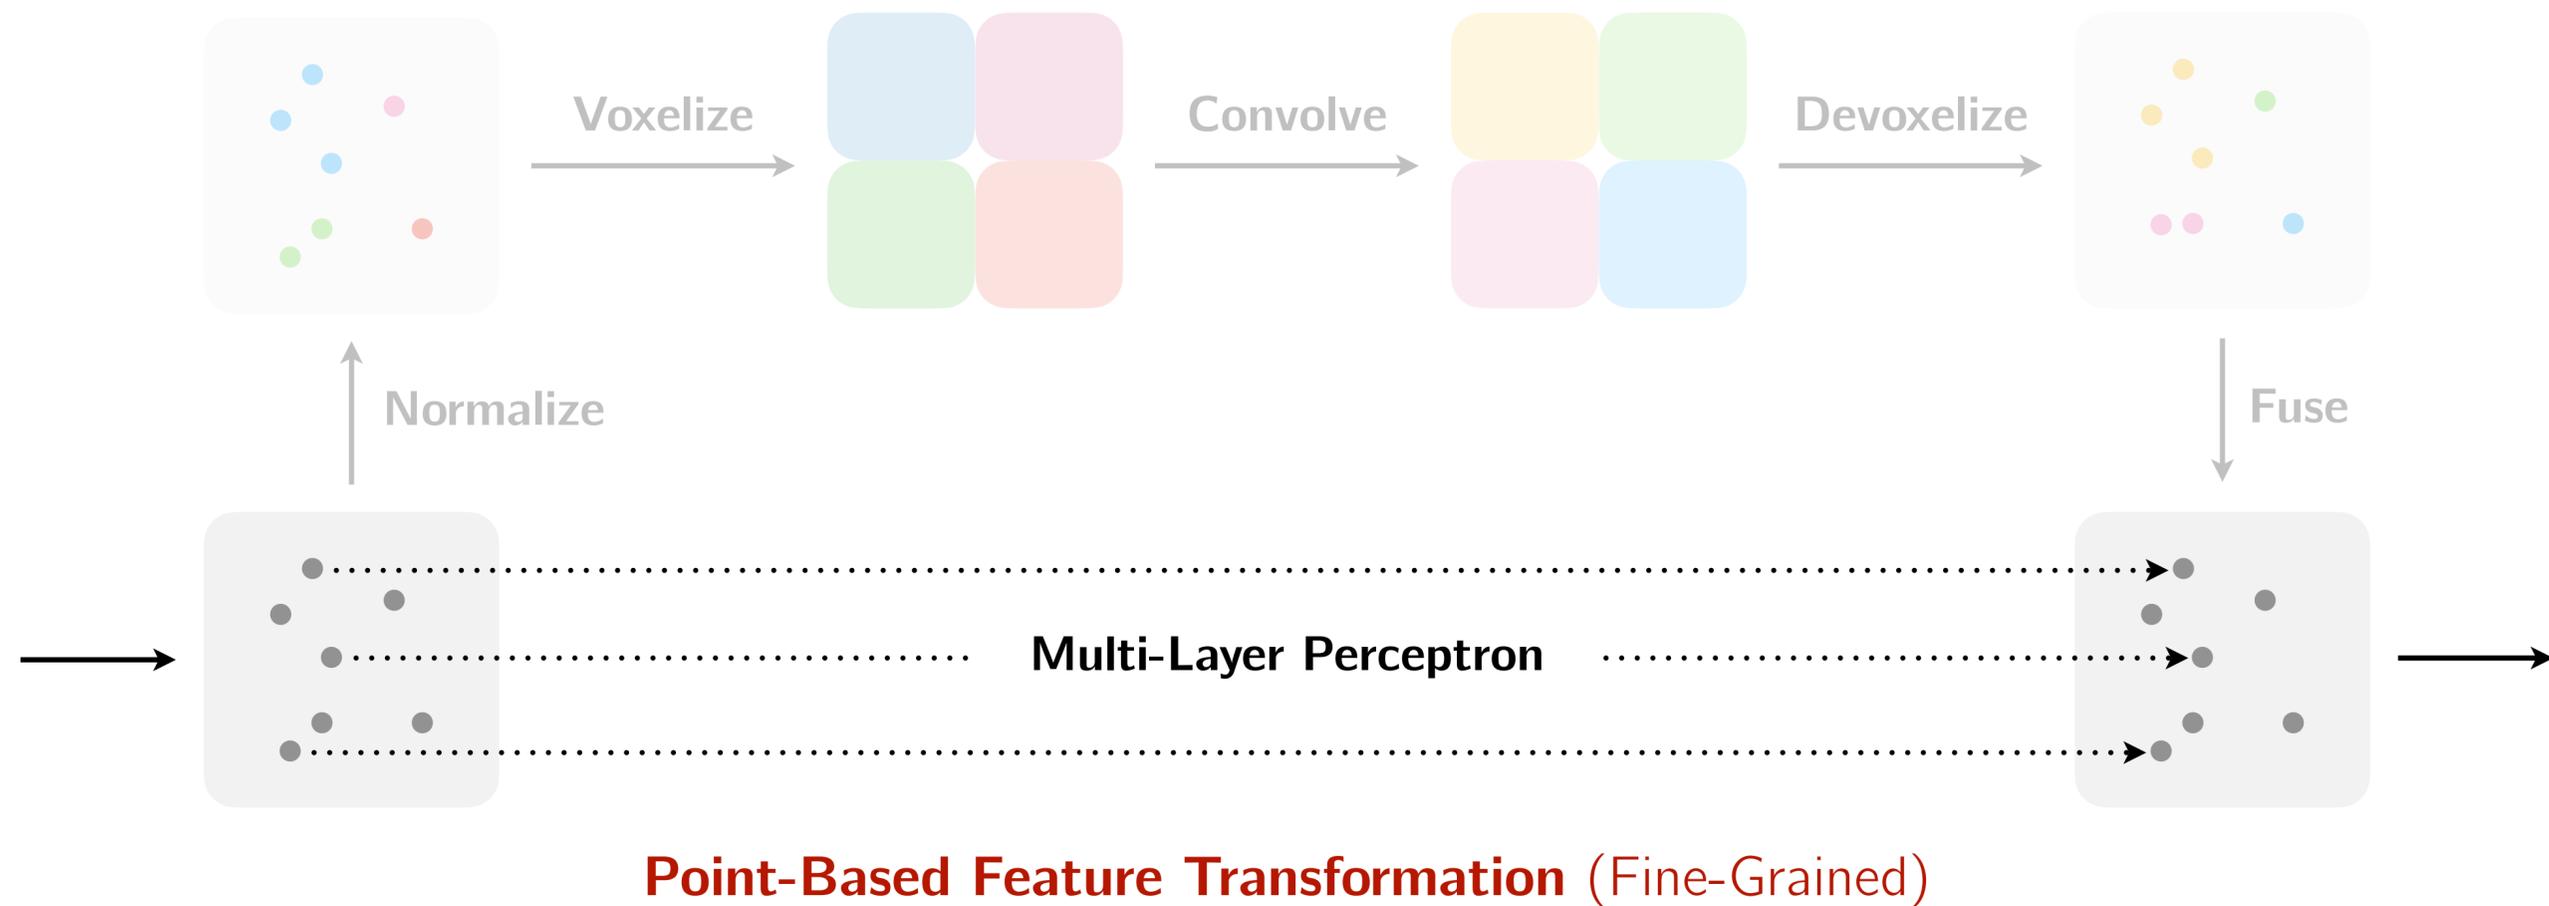
PointNet [CVPR'17]
PointCNN [NeurIPS'18]
DGCNN [SIGGRAPH'19]



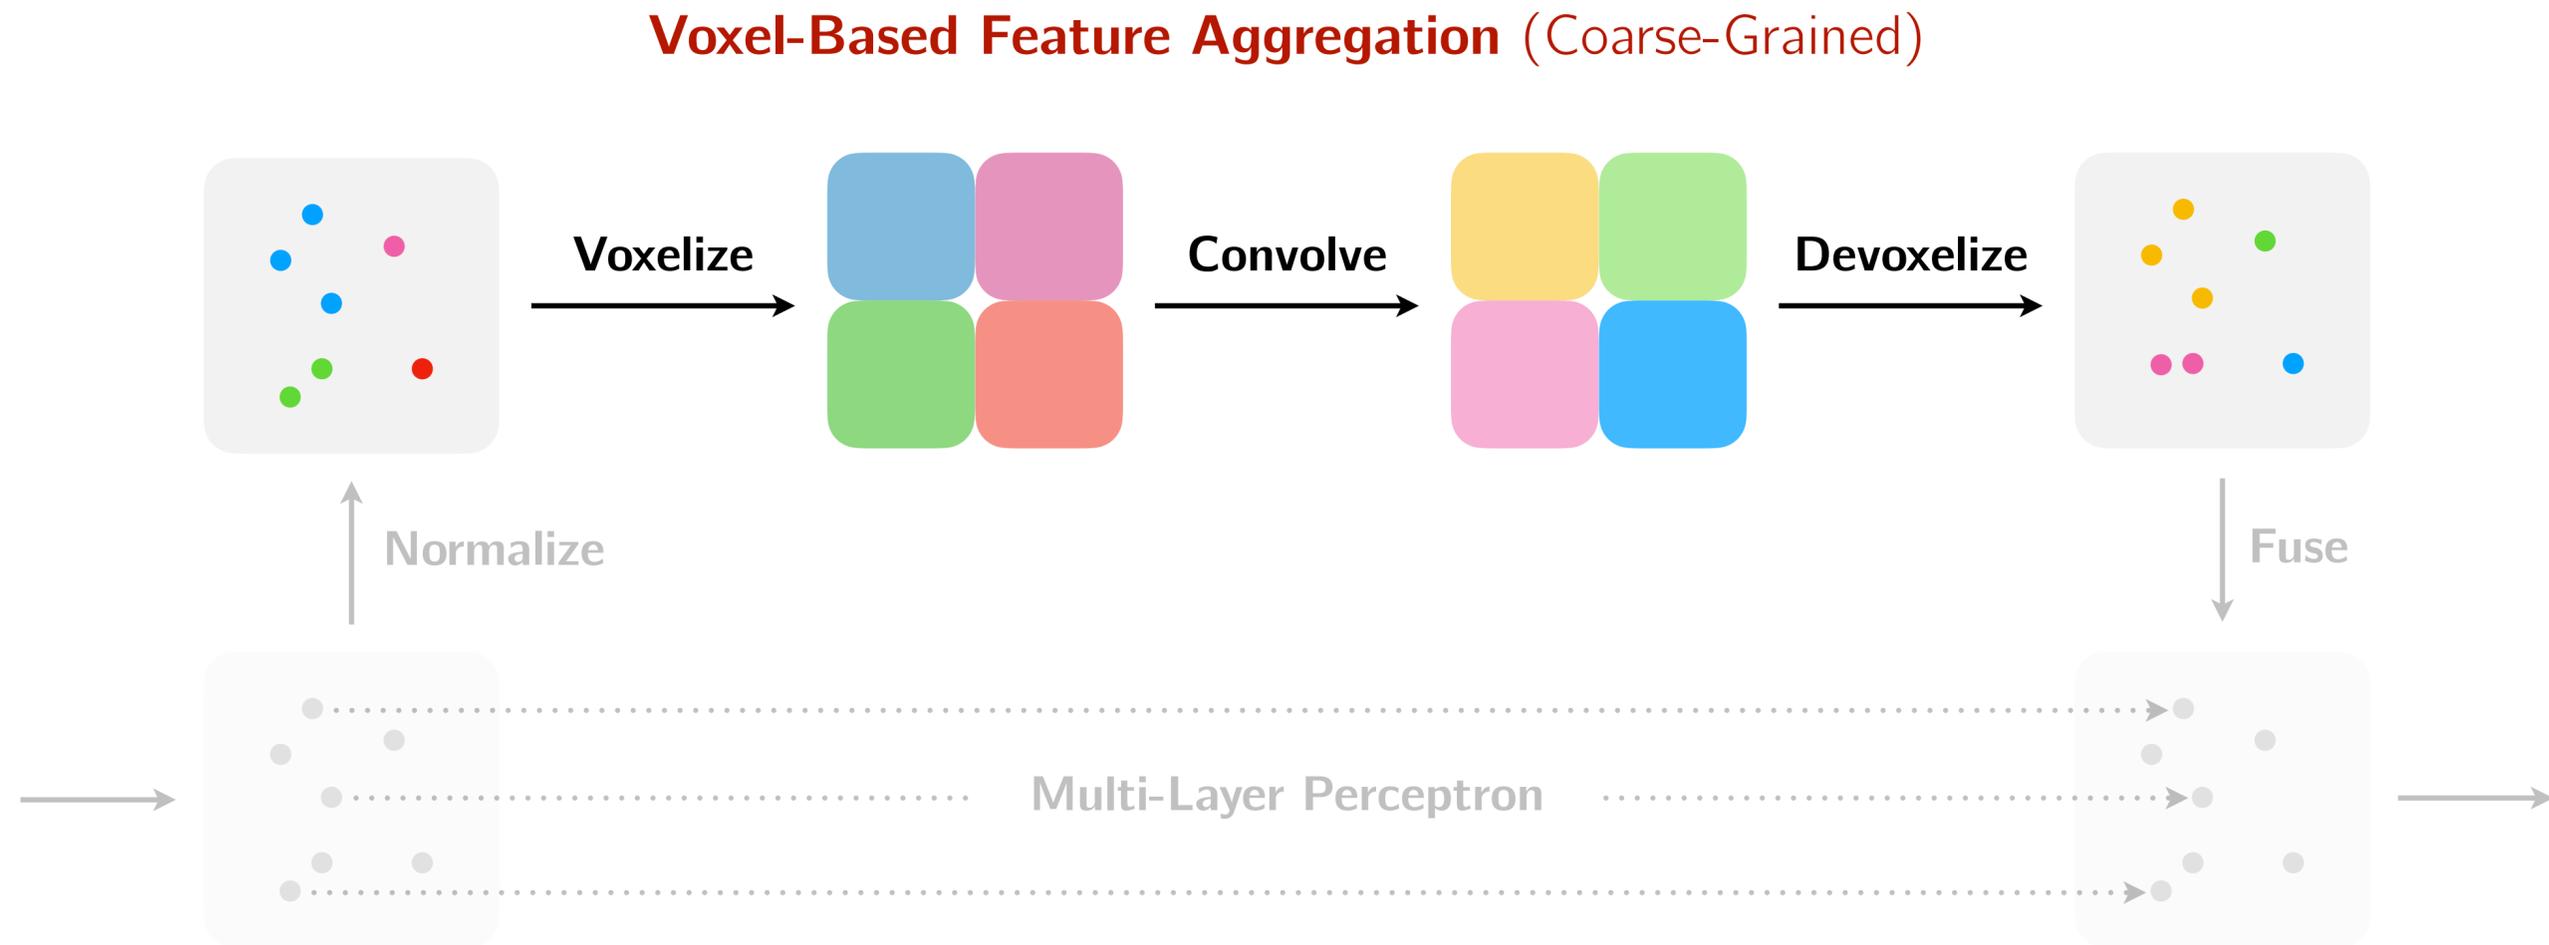
Point-Voxel Convolution (PVConv)



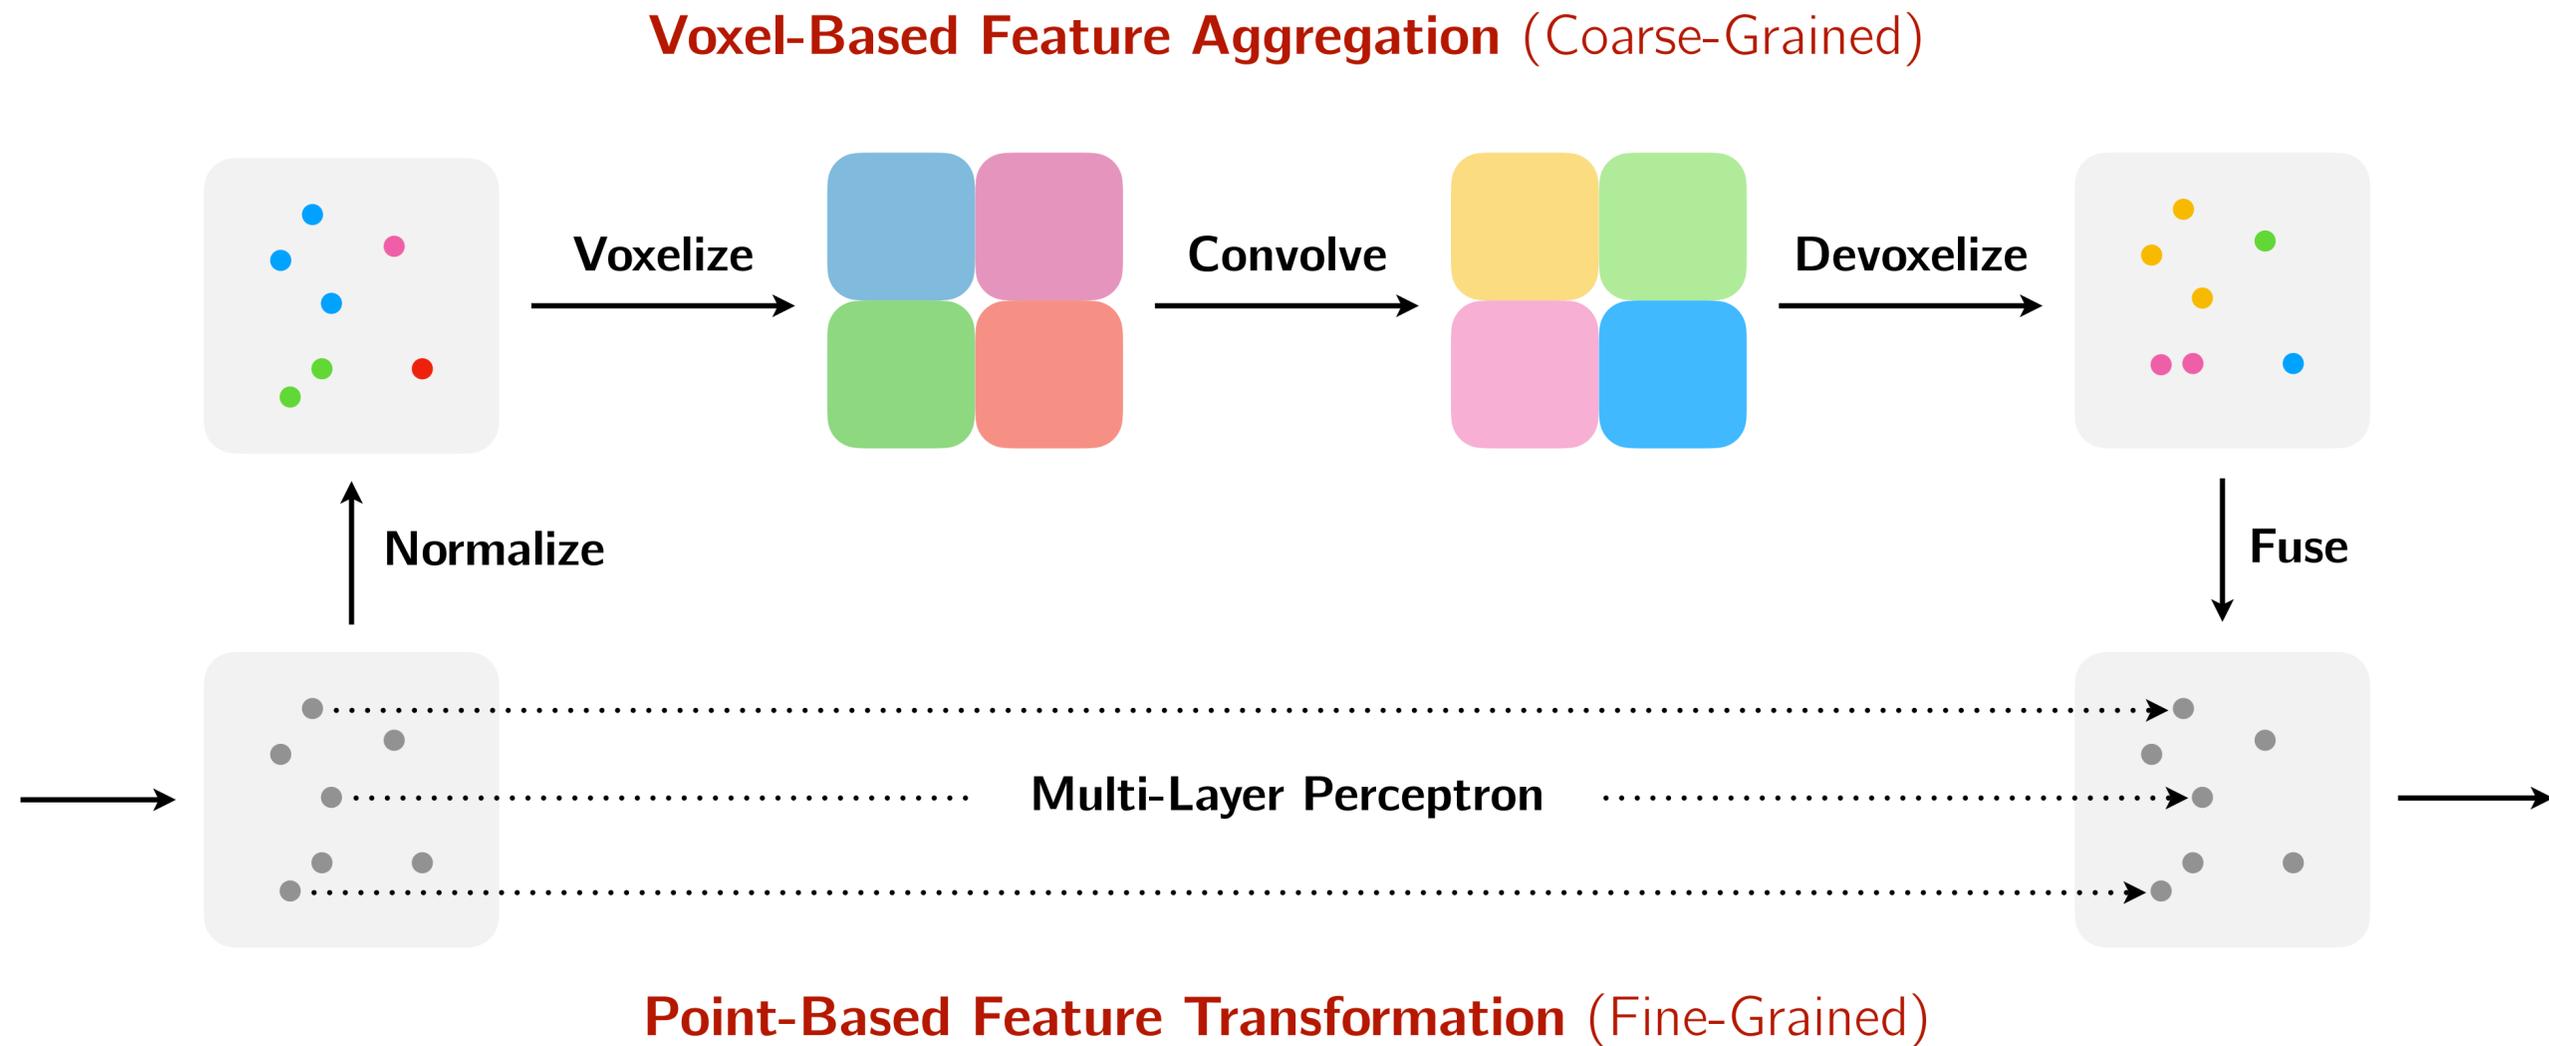
Point-Voxel Convolution (PVConv)



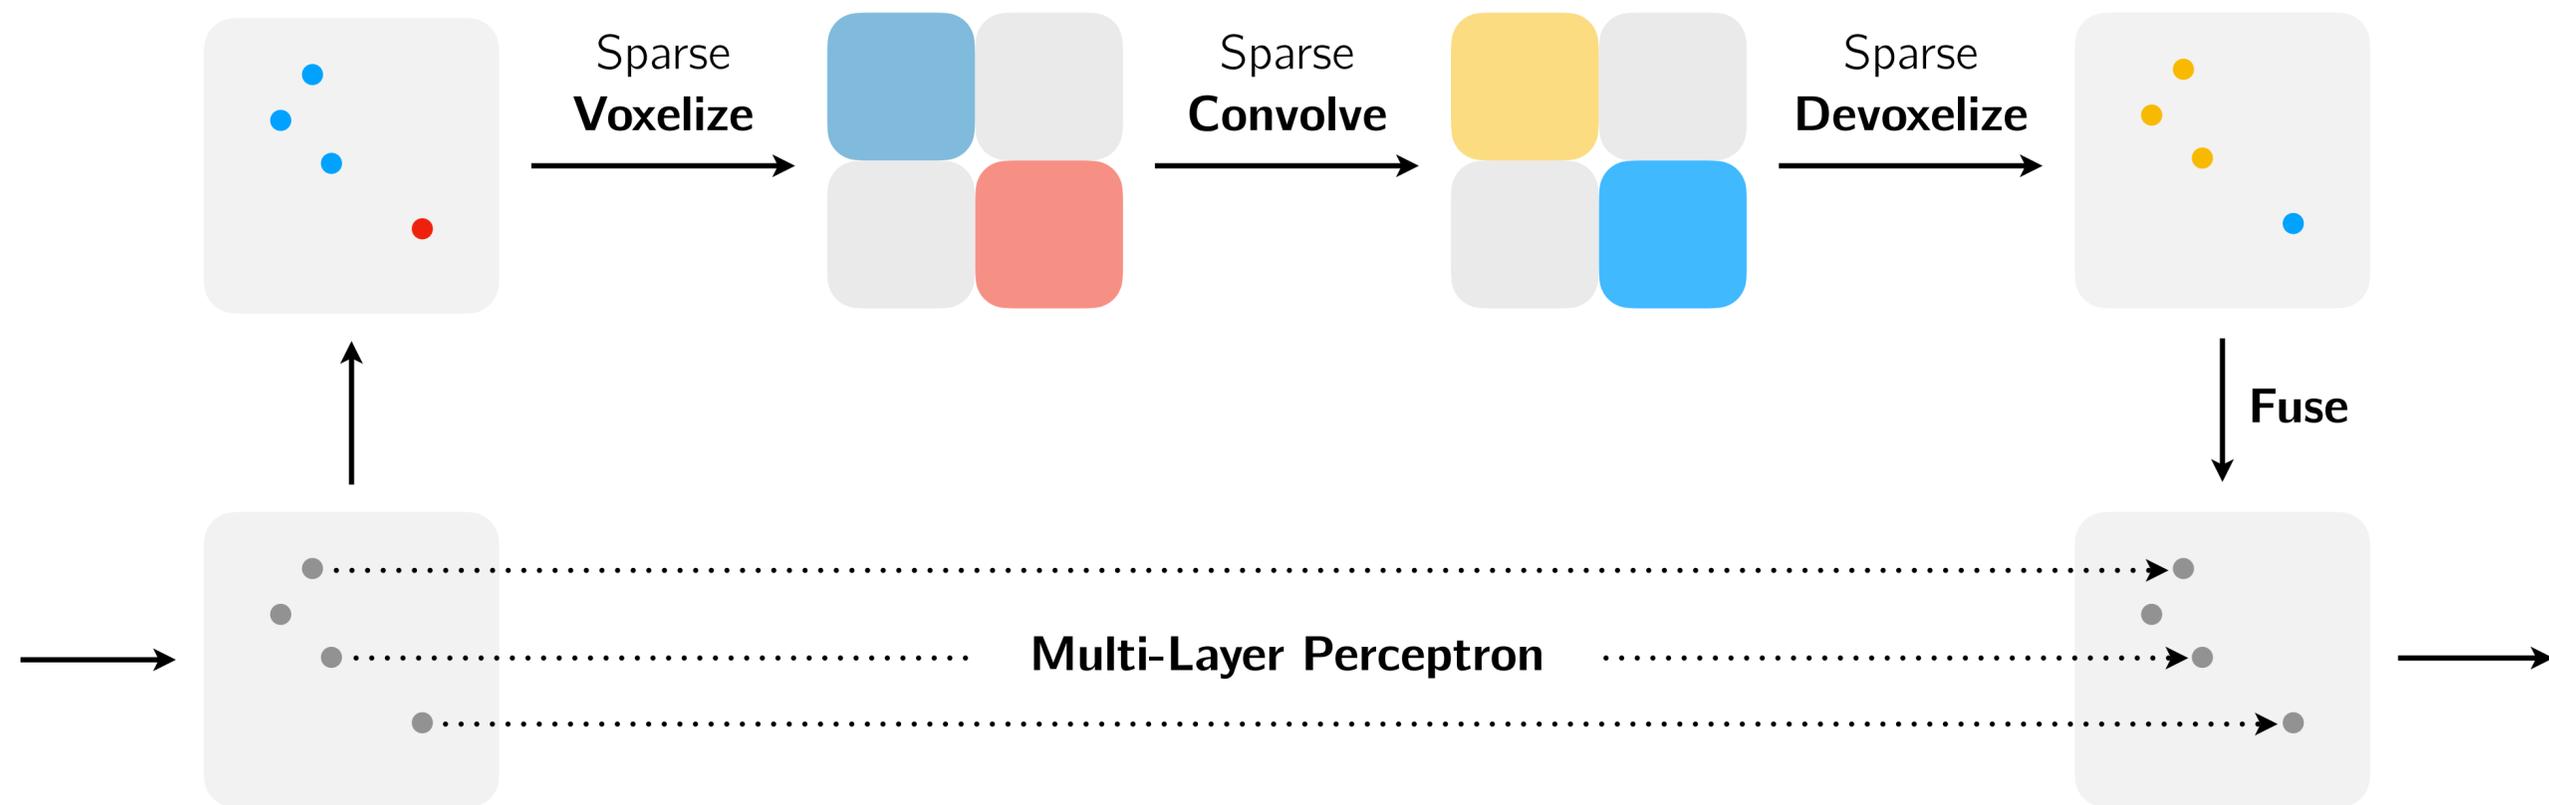
Point-Voxel Convolution (PVConv)



Point-Voxel Convolution (PVConv)

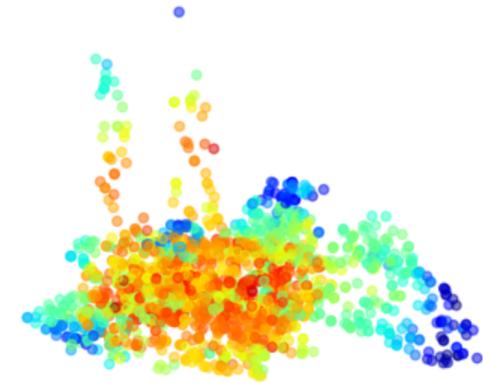
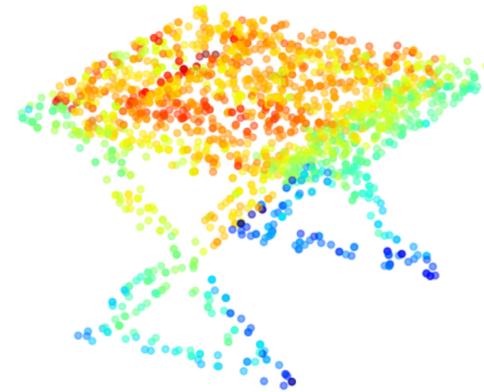
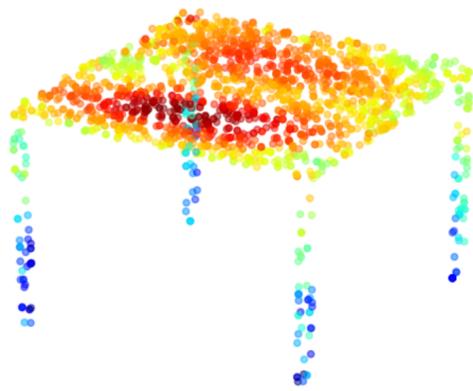


Sparse Point-Voxel Convolution (SPVConv)

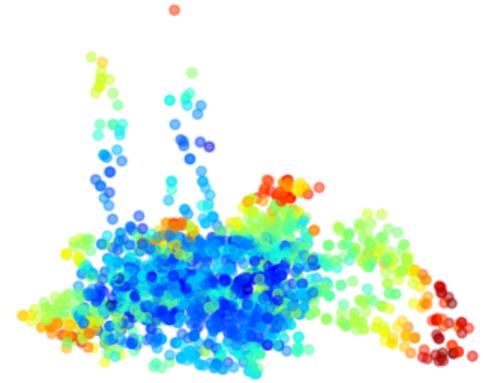
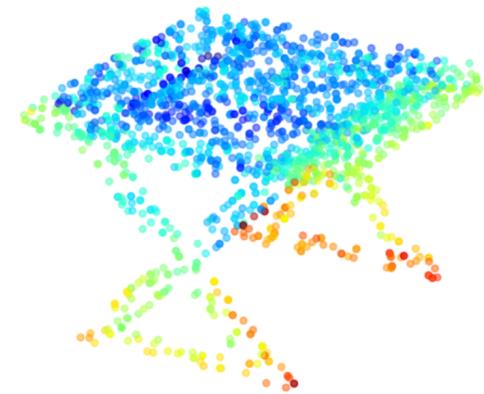
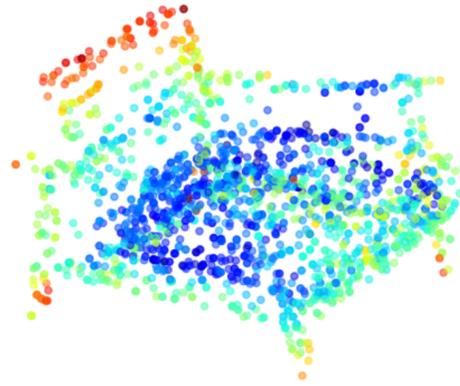
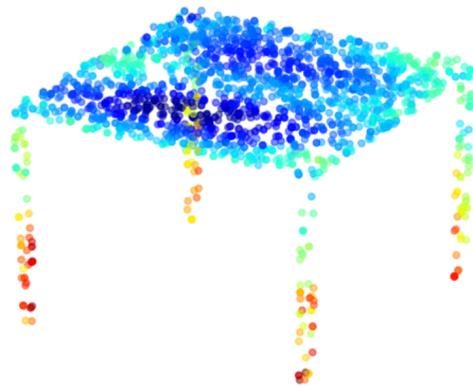


Point-Voxel Convolution (PVConv)

Features from **Voxel-Based Branch**:



Features from **Point-Based Branch**:



Accelerated 3D Sparse Convolution

examples	[Major] Add files.	19 days ago
torchsparse	[Major] Add files.	19 days ago
LICENSE	Update LICENSE.	7 days ago
README.md	[Minor] Update REAMD.md	17 days ago
setup.py	[Major] Add files.	19 days ago

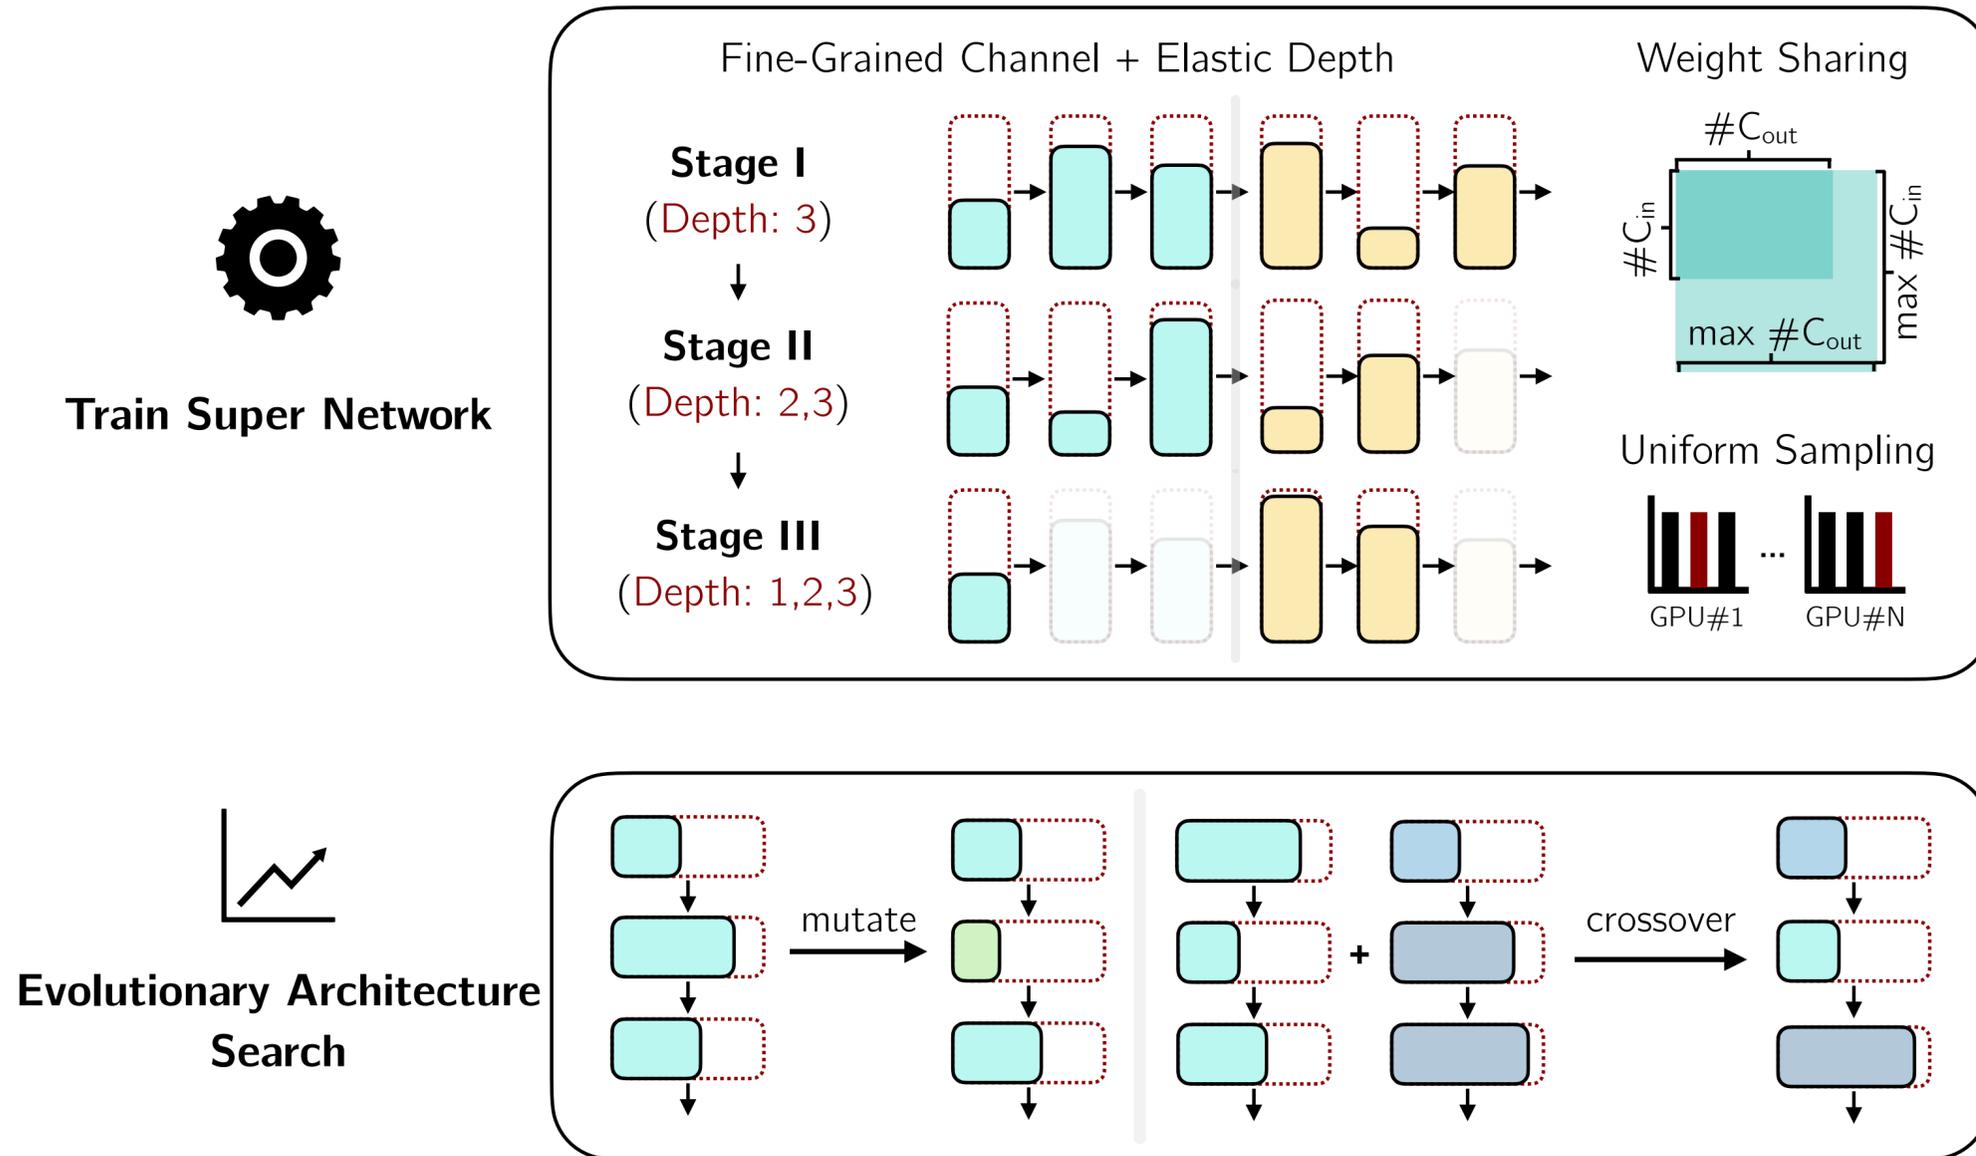
README.md



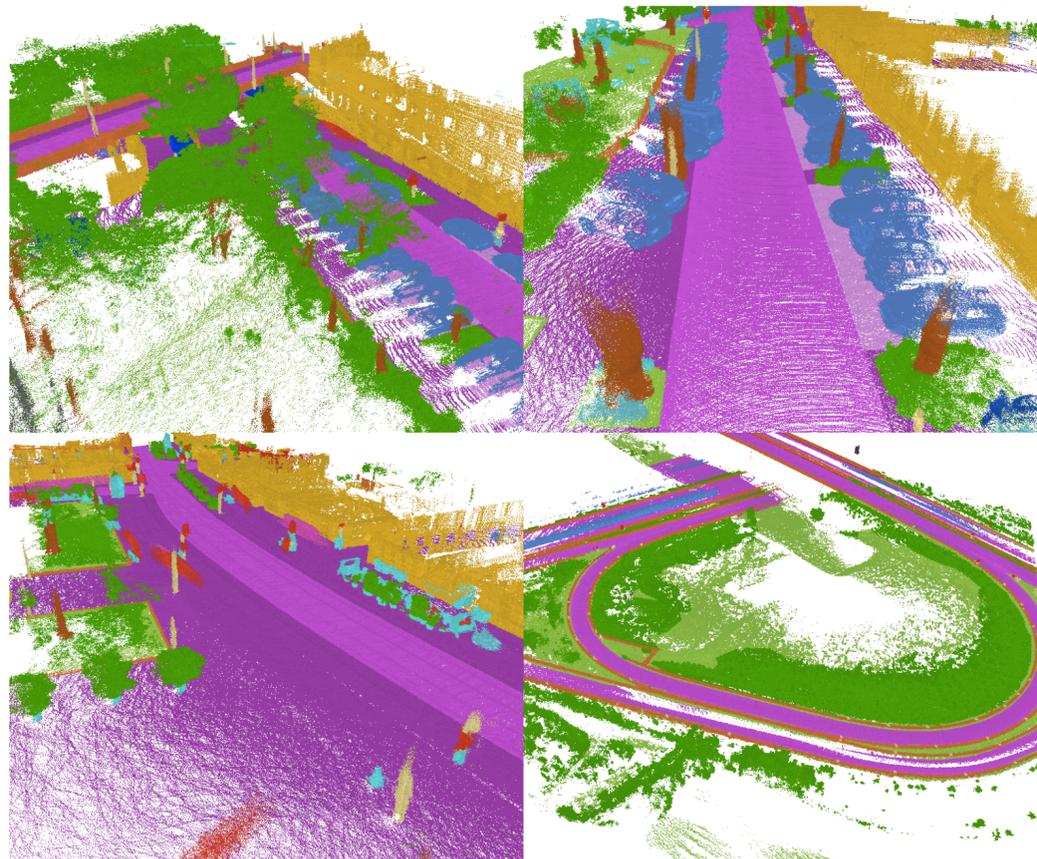
torchsparse: High-Performance Computing Library for Efficient 3D Sparse Convolution

Network	Latency (ME V0.4.3)	Latency (torchsparse V1.0.0)
MinkUNet18C (MACs / 10)	224.7	124.3
MinkUNet18C (MACs / 4)	244.3	160.9
MinkUNet18C (MACs / 2.5)	269.6	214.3
MinkUNet18C	323.5	294.0

3D Neural Architecture Search (3D-NAS)



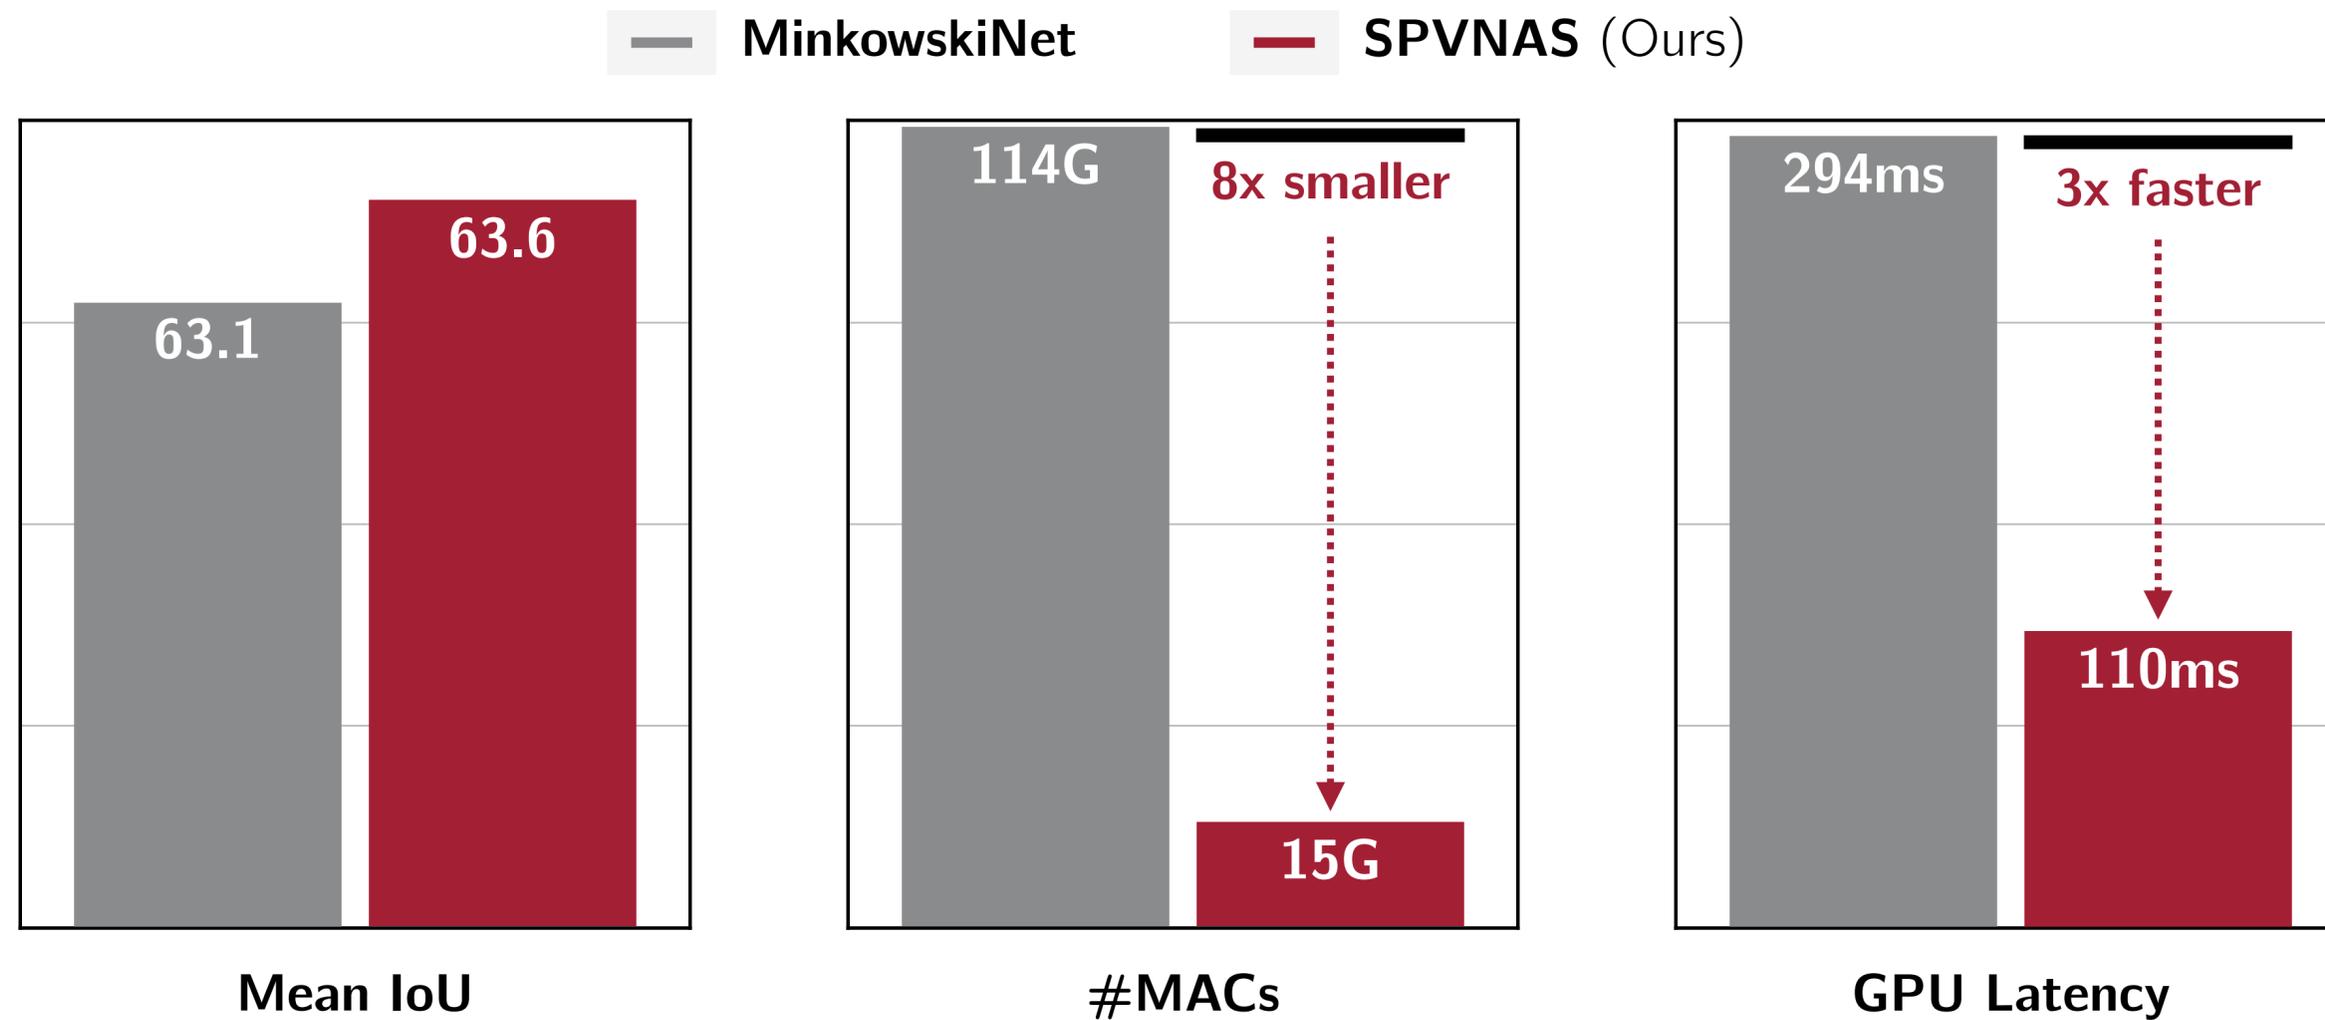
Results: 3D Outdoor Scene Segmentation



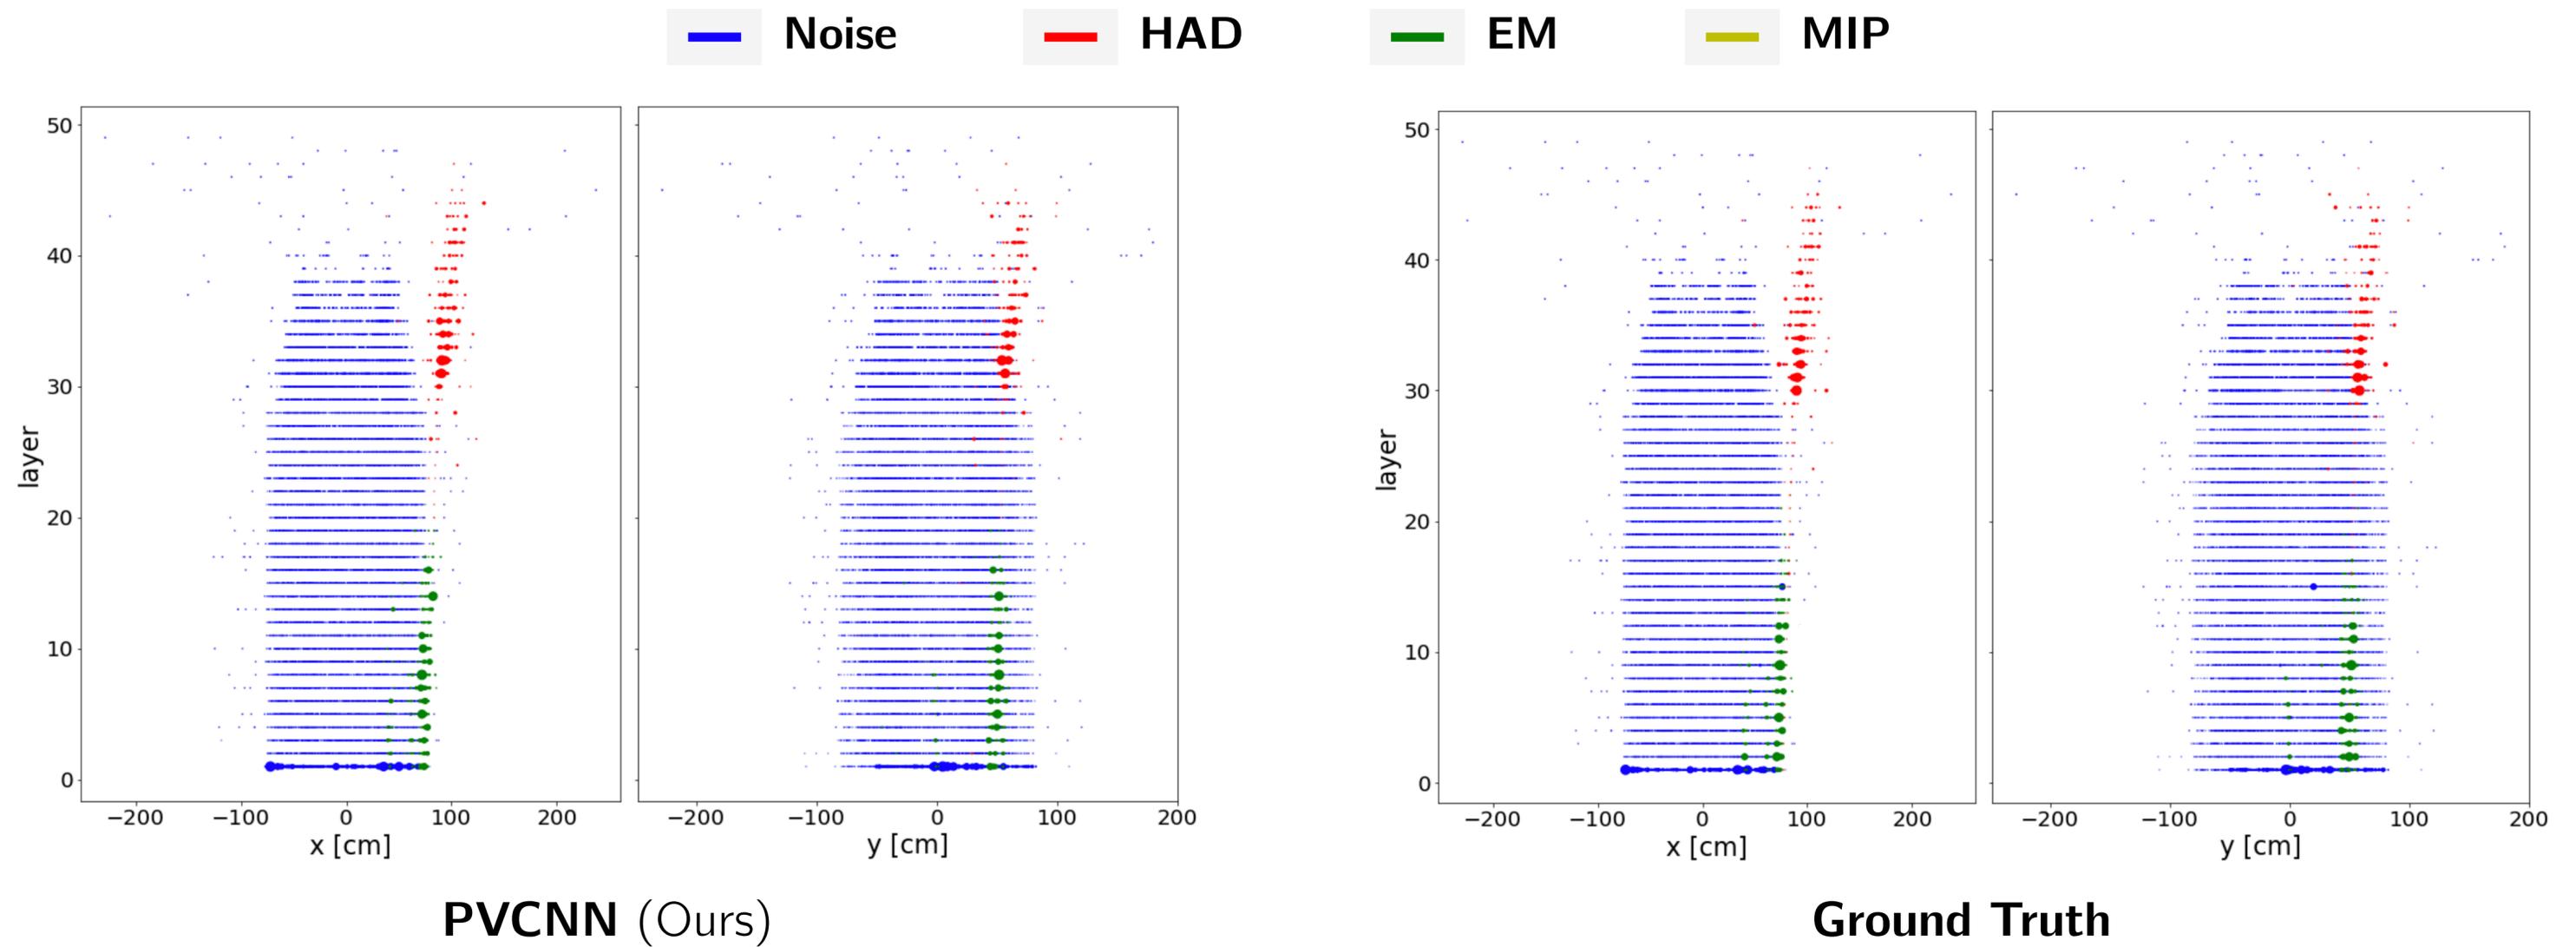
SemanticKITTI [ICCV 2019]

Approach	Paper	Code	mIoU	Classes (IoU)
SPVNAS			67.0	
TORNADONet			63.1	
KPRNet			63.1	
Cylinder3D			61.8	
FusionNet			61.3	
SalsaNext			59.5	
KPConv			58.8	
SqueezeSegV3			55.9	

Results: 3D Outdoor Scene Segmentation



Results: Semantic Segmentation on HGCAL

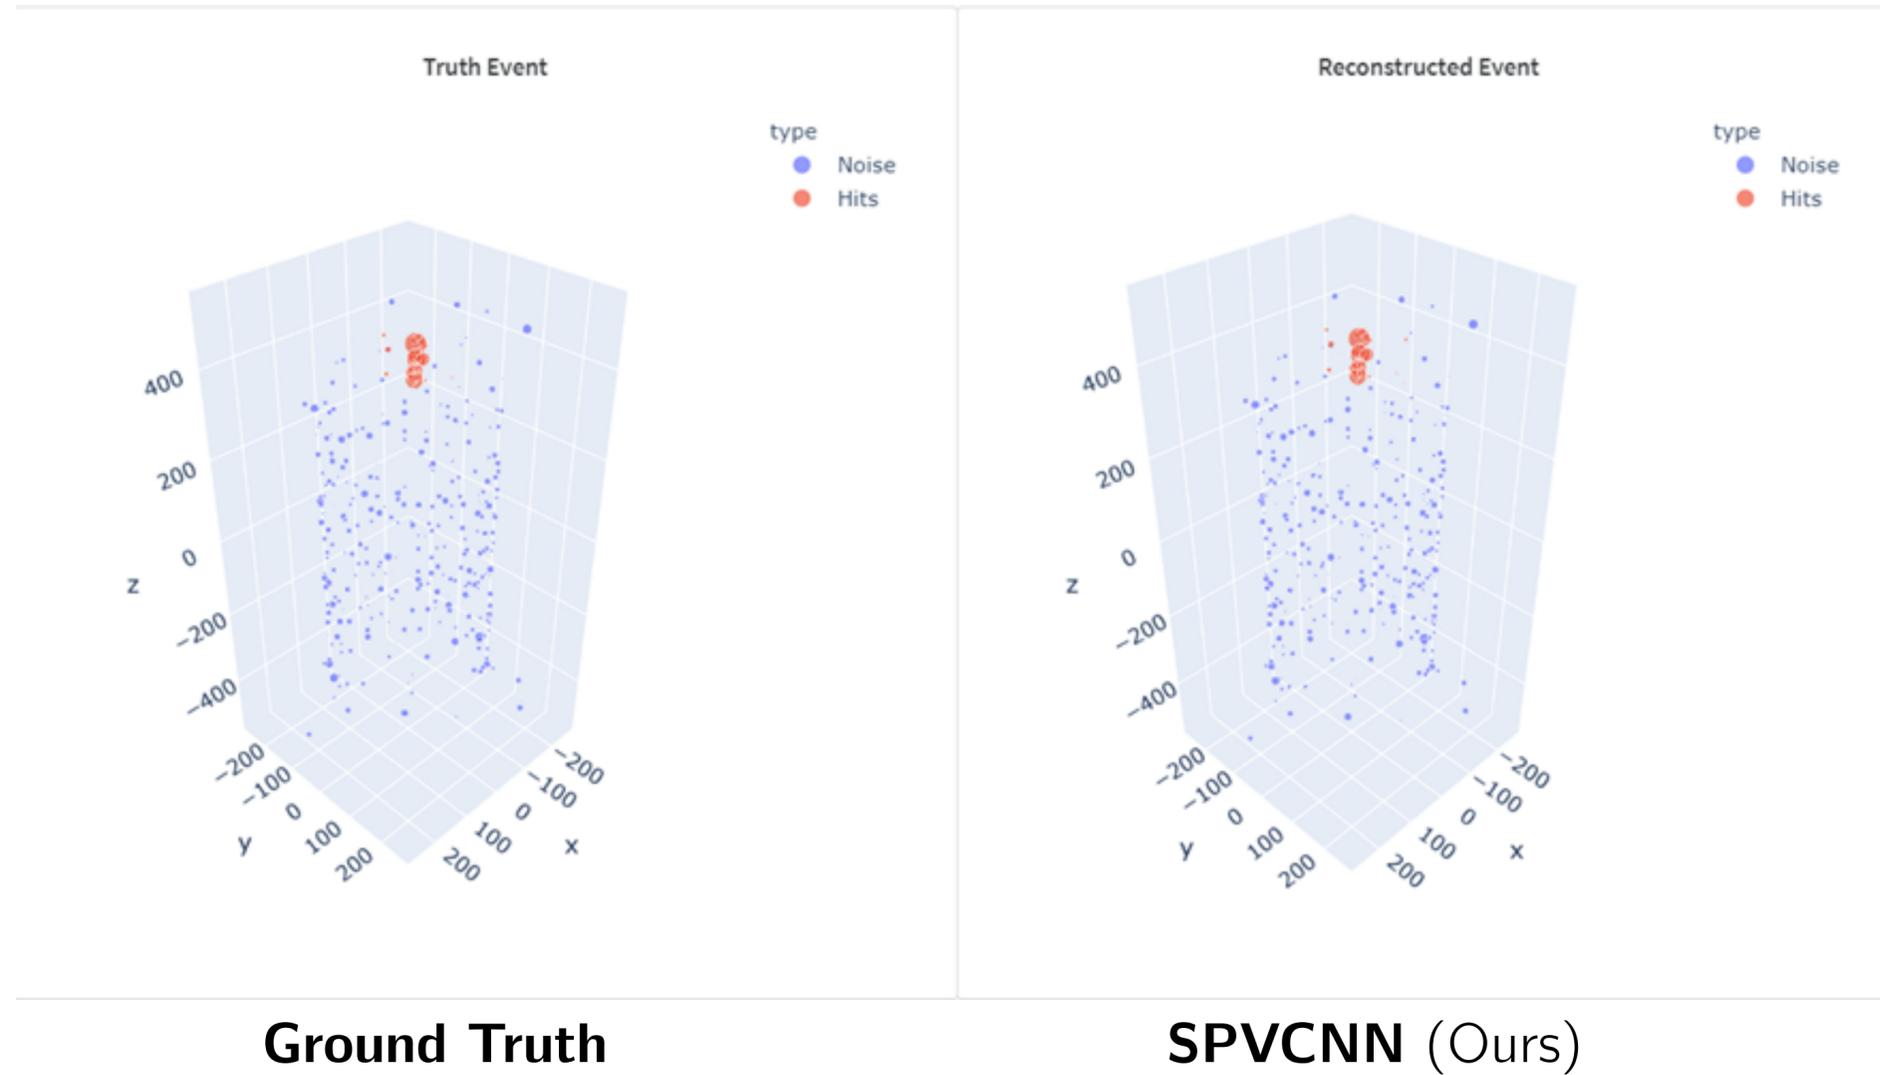


Results: Semantic Segmentation on HGCAL

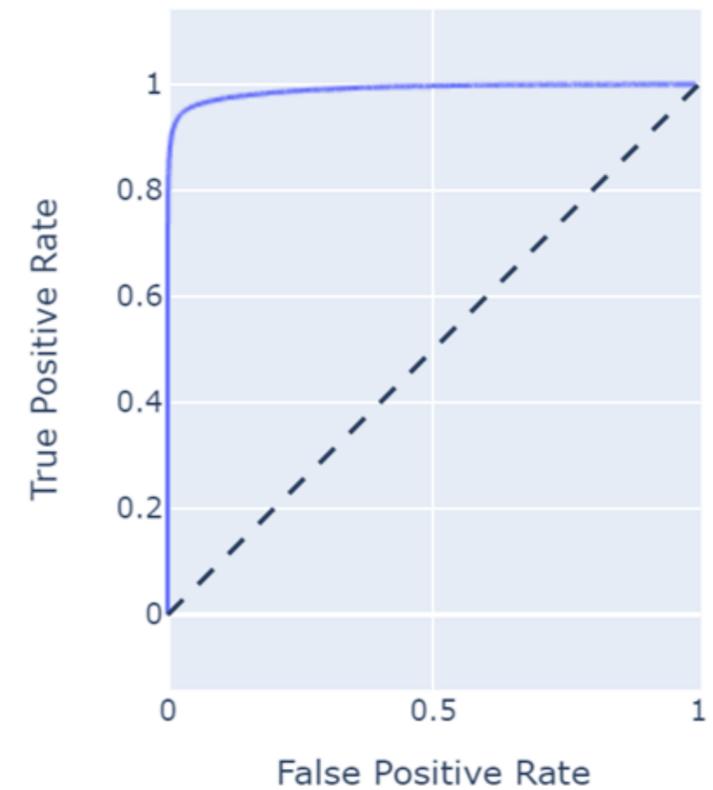
Very accurate in filtering particles from noise

	Noise	HAD	EM	MIP
Noise	<u>99.84%</u>	2.68%	2.75%	1.64%
HAD	0.09%	80.04%	14.87%	6.80%
EM	0.07%	17.15%	82.33%	0.89%
MIP	0.01%	0.14%	0.05%	90.68%

Results: Semantic Segmentation on HCAL

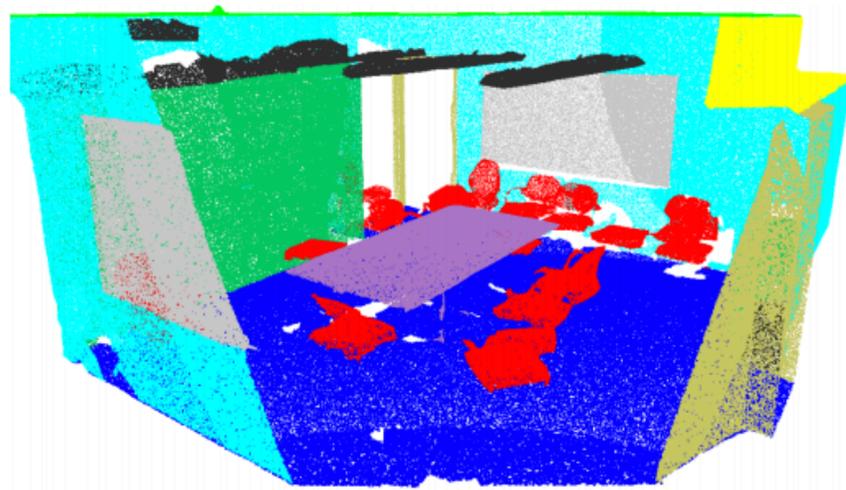


ROC Curve (AUC=0.9899)

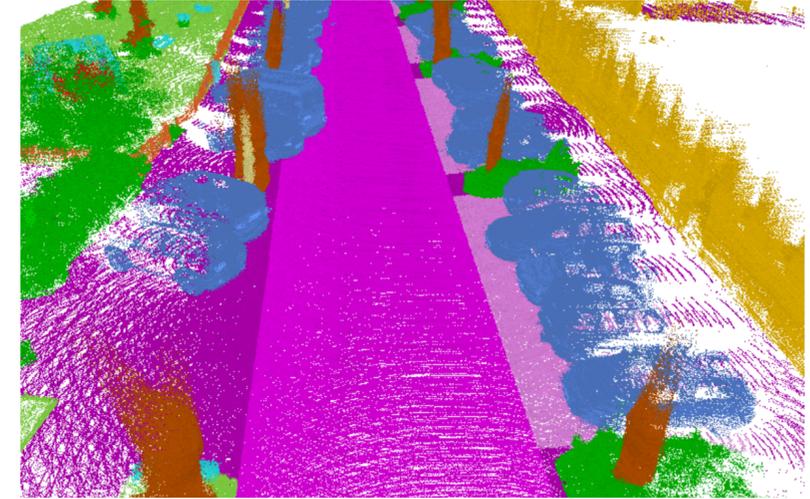




2.7x measured speedup
1.5x memory reduction



6.9x measured speedup
5.7x memory reduction

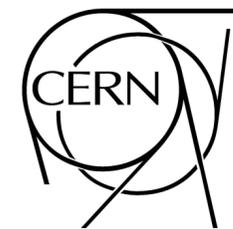


2.7x measured speedup
7.6x computation reduction

Autonomous Racing Vehicle:



High-Energy Particle Physics:



Media Coverage:



GitHub: <https://github.com/mit-han-lab/e3d>