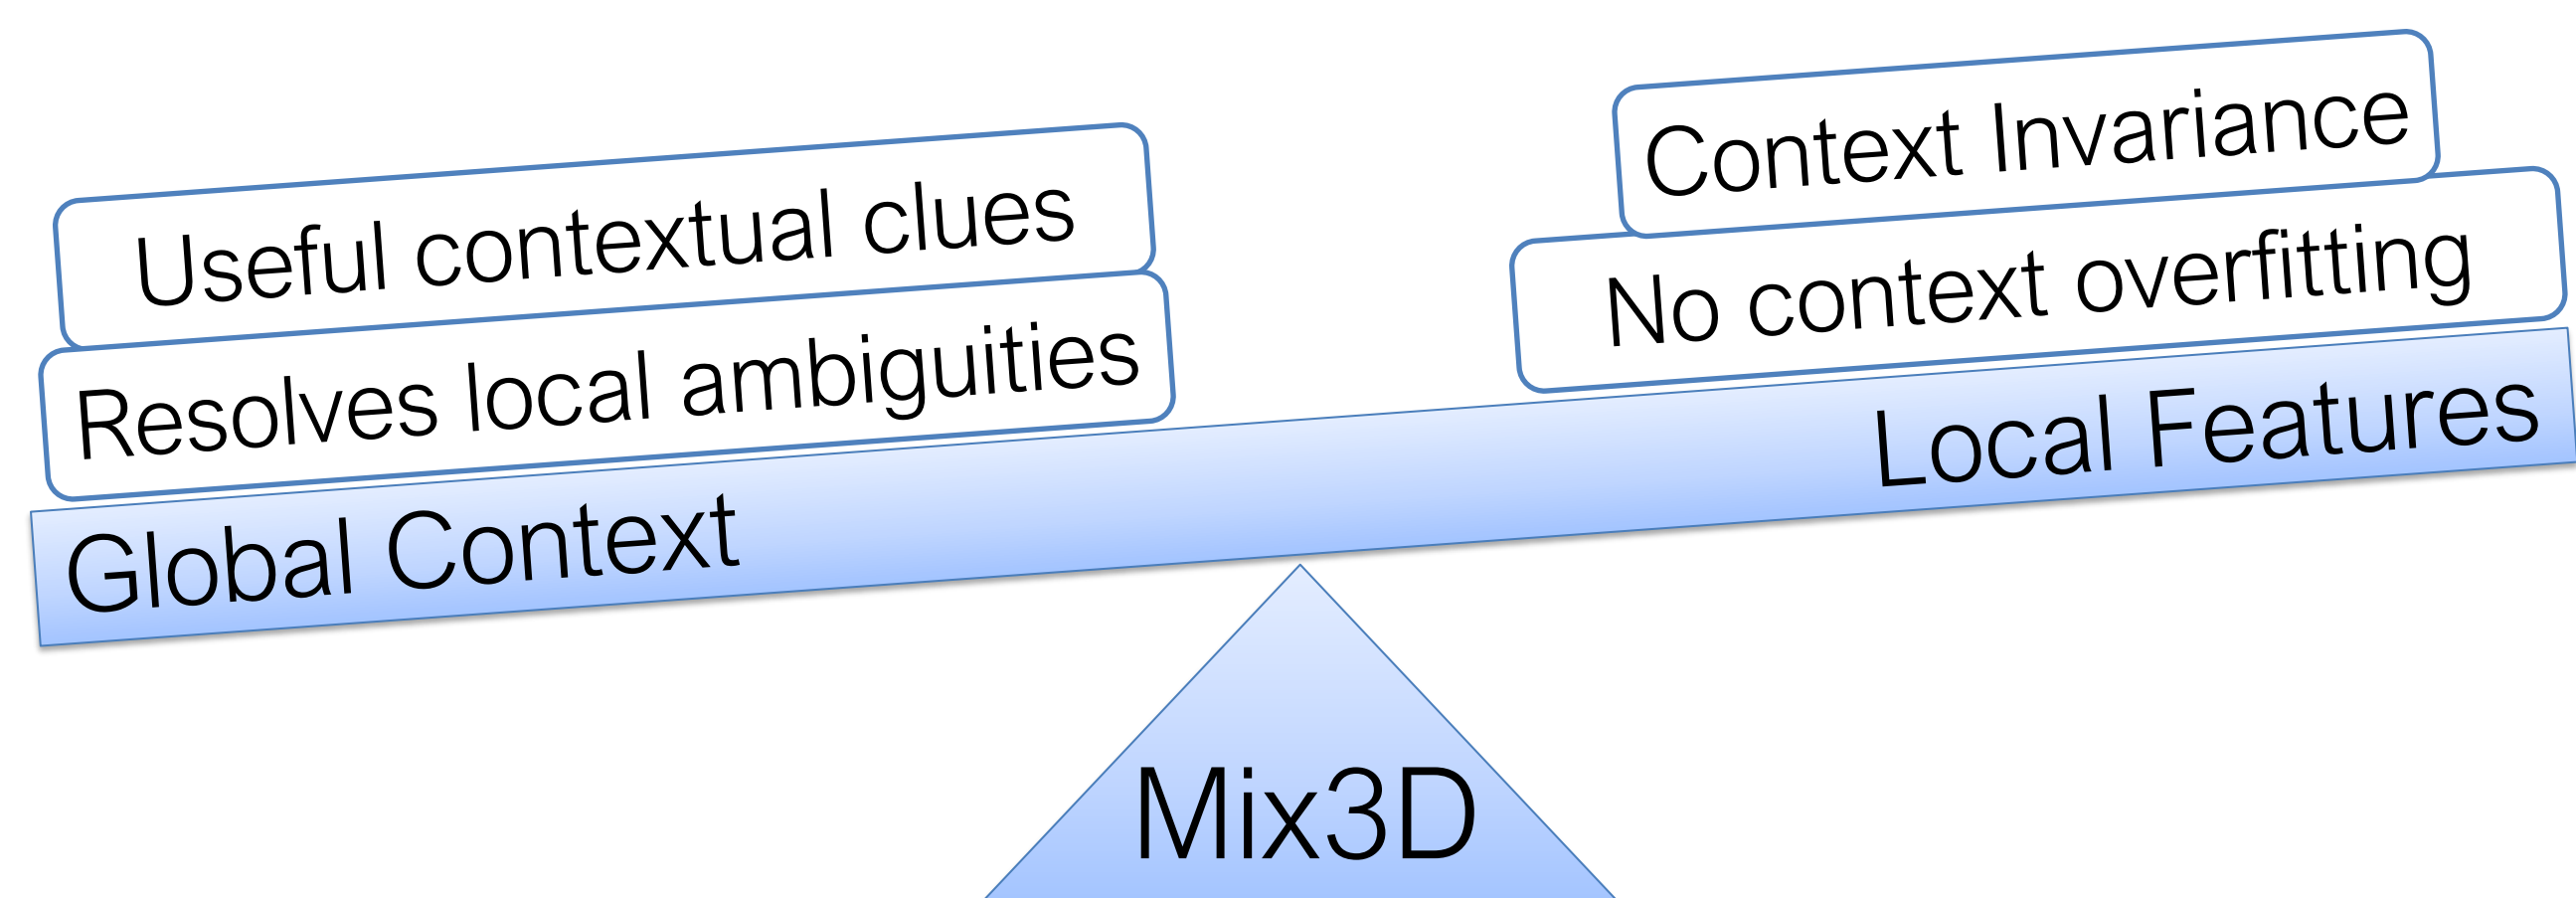




Mix3D balances global context and local geometry




Contributions

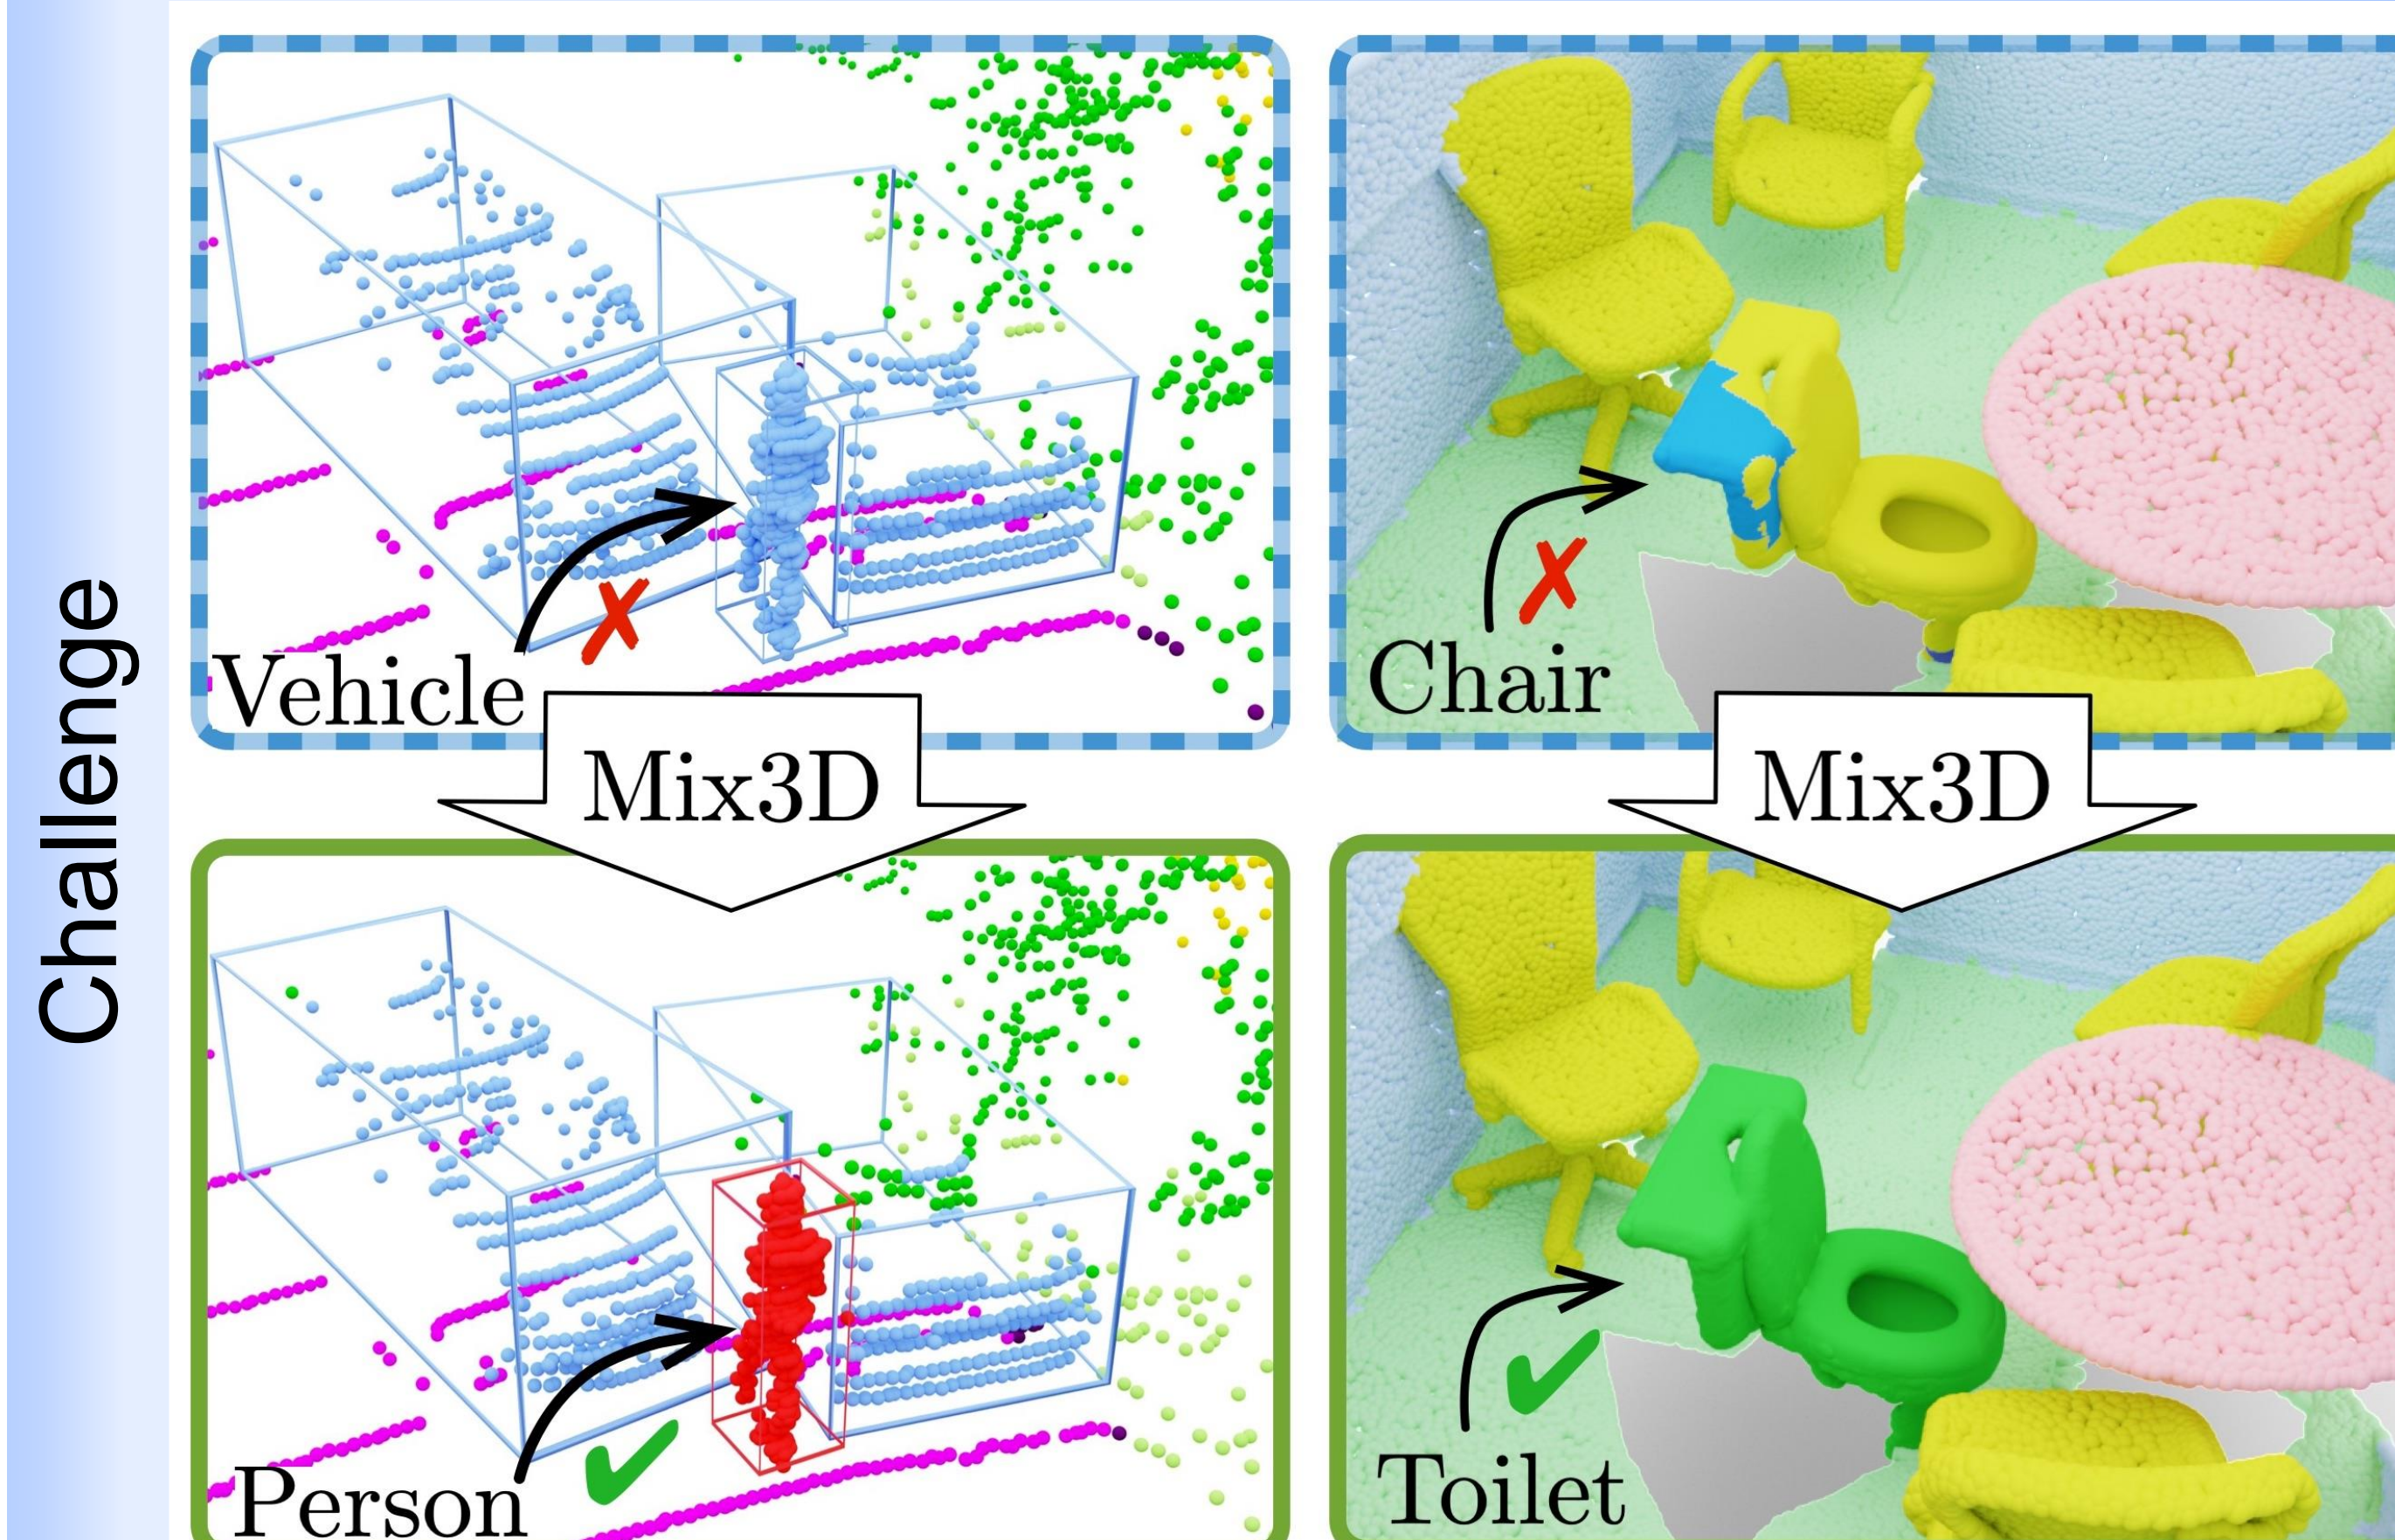
- Augmentation which balances global context and local geometry and therefore generalizes beyond the contextual priors of the training set
- SOTA methods trained with Mix3D show consistent improvements on indoor and outdoor datasets
- We report state-of-the-art on the popular ScanNet benchmark (78.1%)

ScanNet Semantic Segmentation
Benchmark Challenge (hidden test set)


Method	Input	mIoU	
Mix3D (ours)	3DV'21	3D	78.1
OccuSeg	CVPR'20	3D	76.4
O-Net	SIGGRAPH'17	3D	76.2
BpNet	CVPR'21	3D + 2D	74.9
VMNet	ICCV'21	3D + Mesh	74.6
Virtual MVFusion	ECCV'20	3D + 2D	74.6
MinkowskiNet	CVPR'19	3D	73.6
SparseConvNet	CVPR'18	3D	72.5
RFCR	CVPR'21	3D	70.2



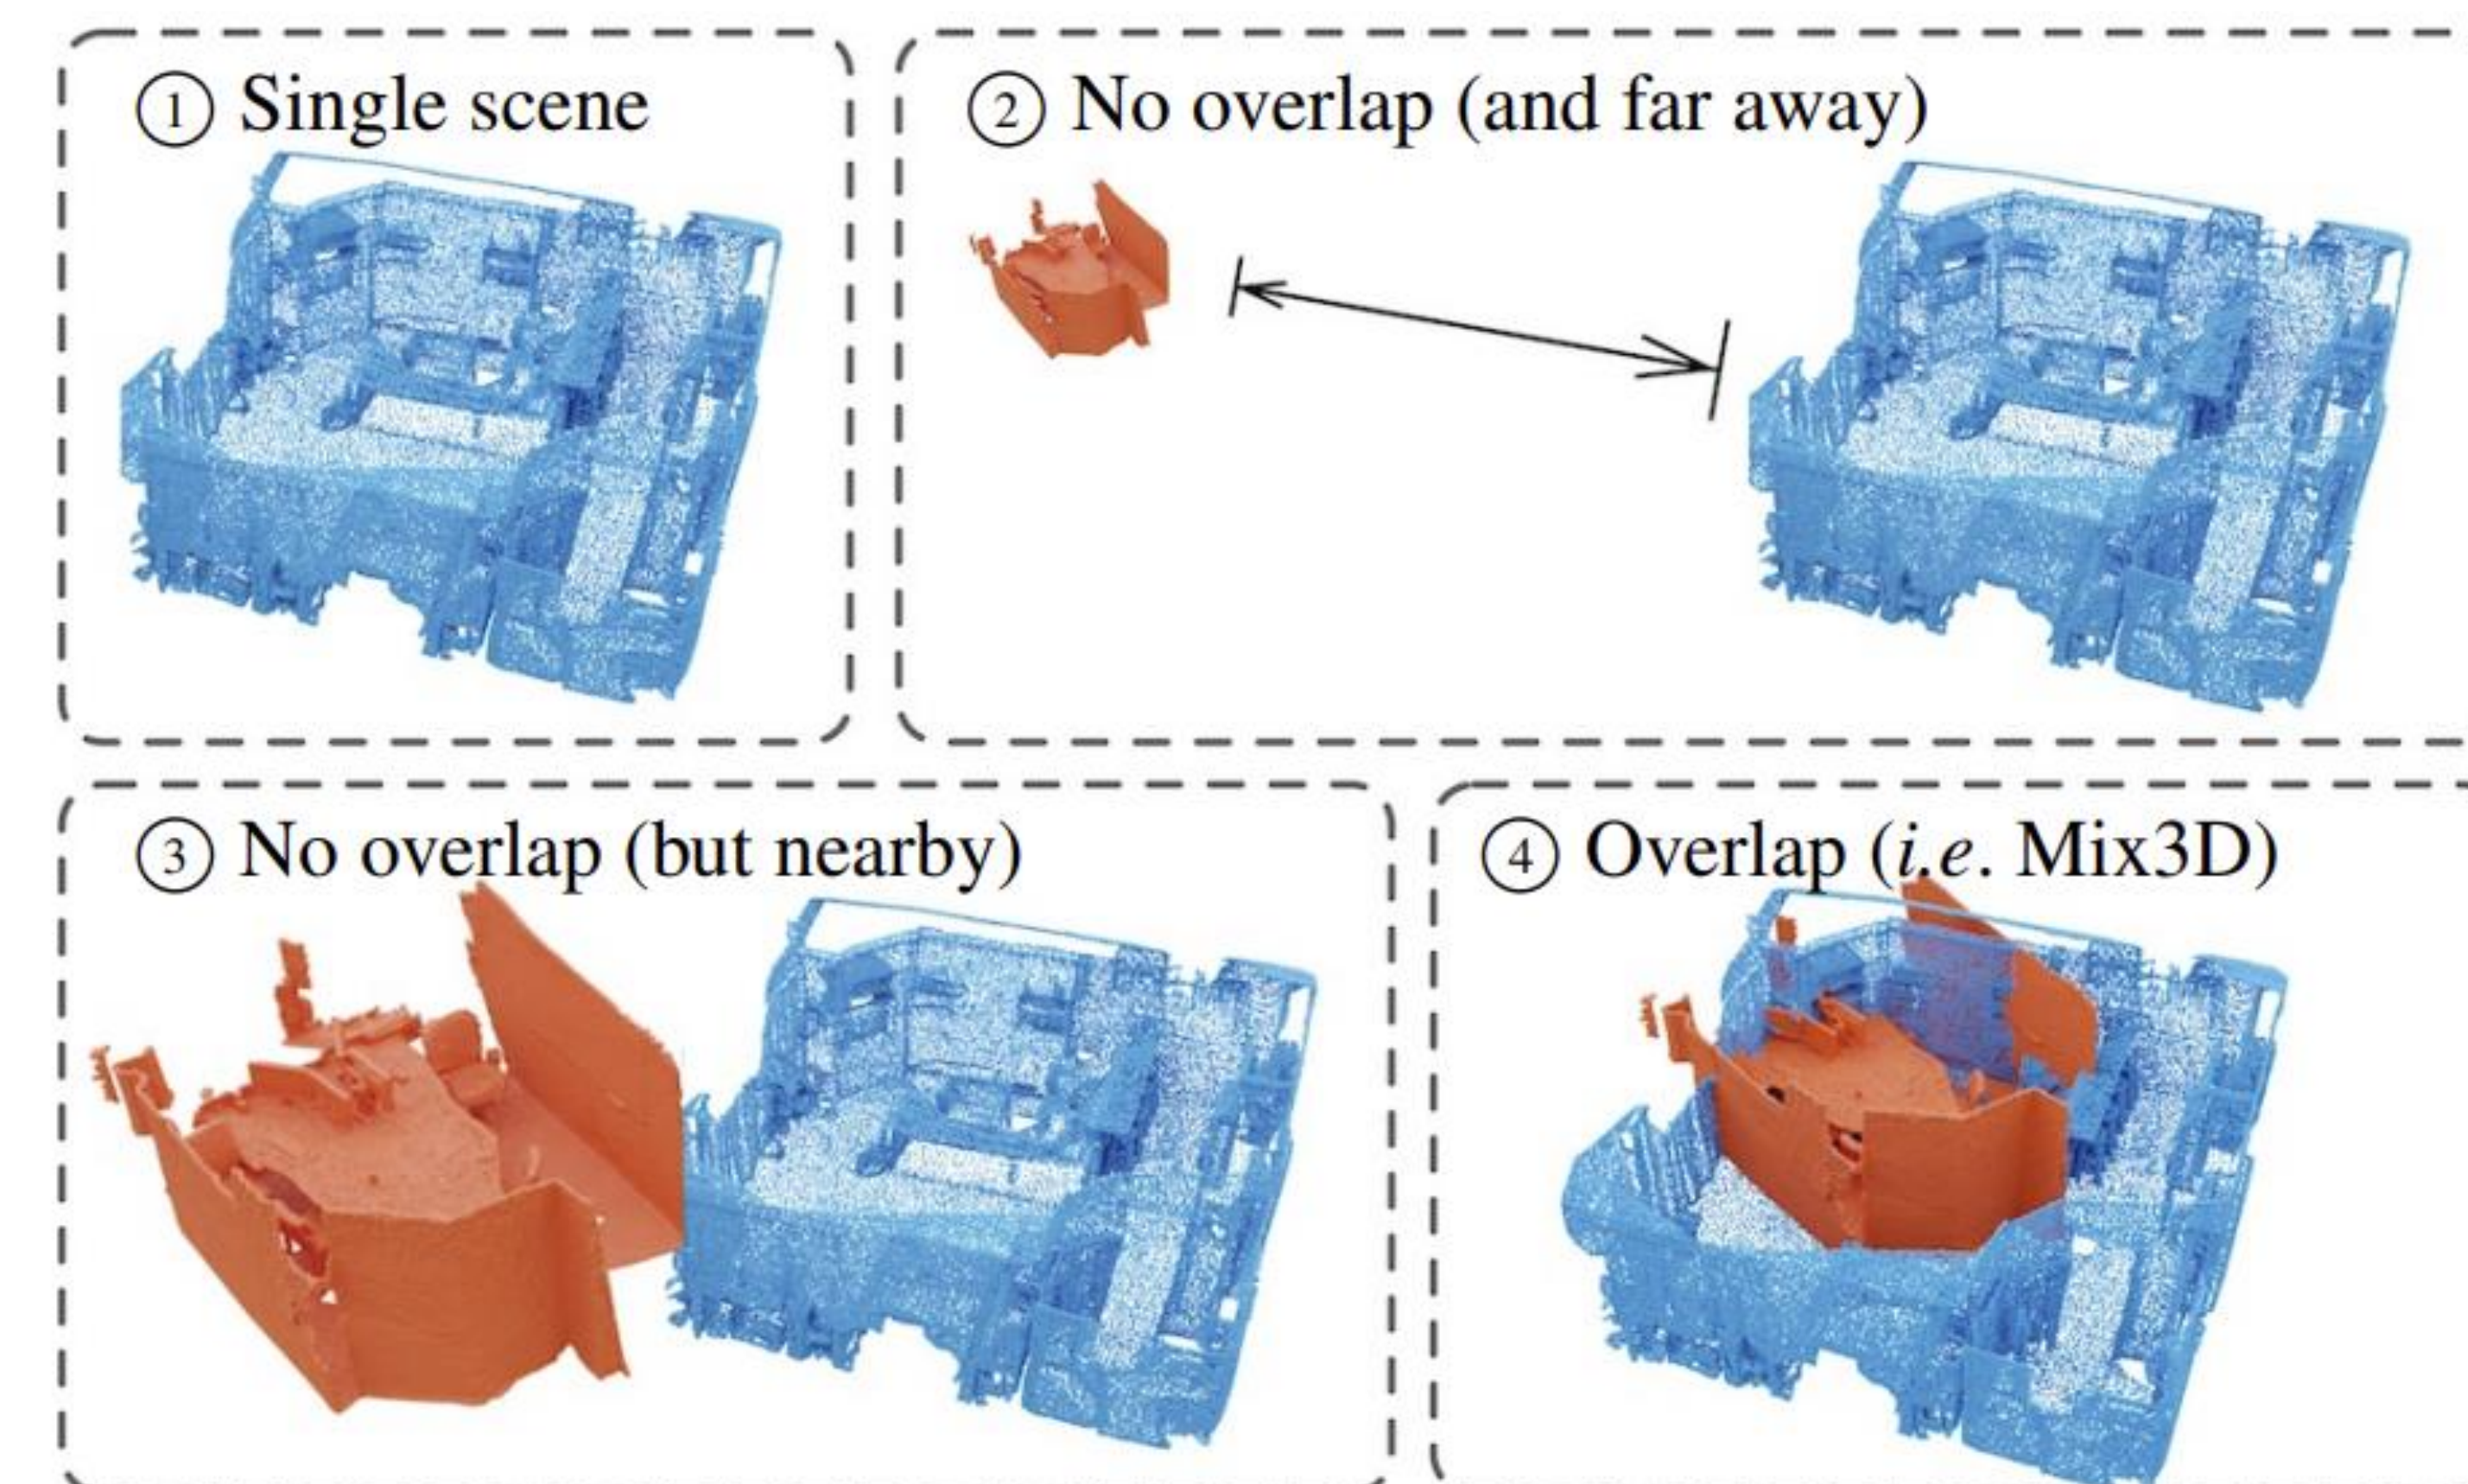
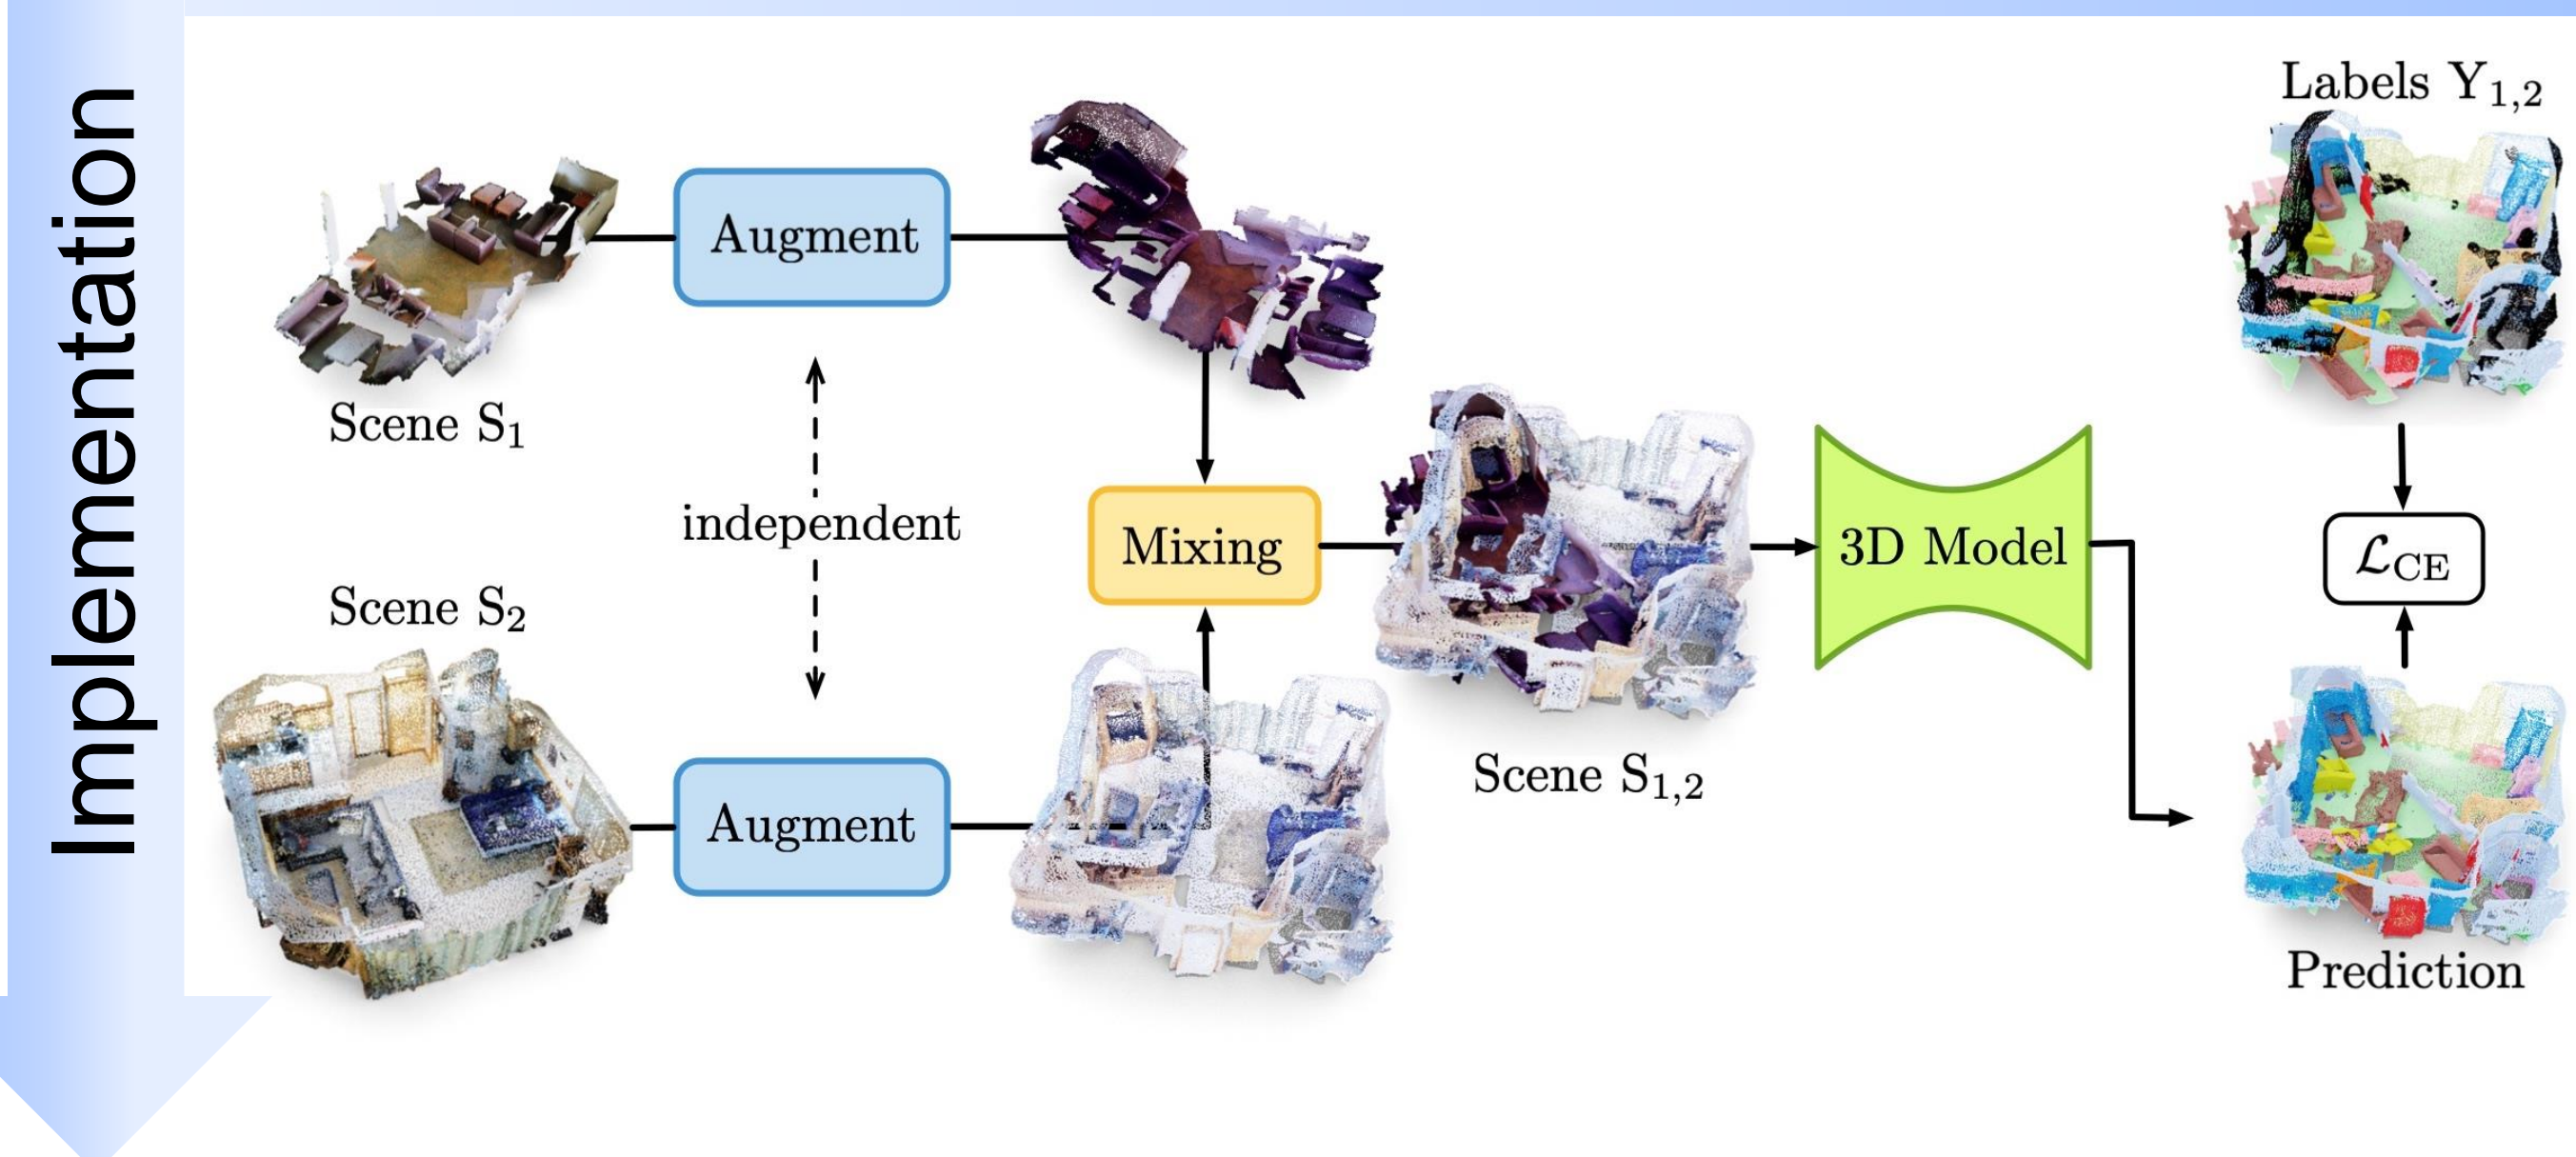
Generalization Beyond Contextual Priors



Training Sample Visualization



Data Augmentation Pipeline



Setup	mIoU
① No mixing (single scene only)	66.6
② No overlap (outside receptive field)	66.7
③ No overlap (inside receptive field)	67.9
④ Overlap (<i>i.e.</i> Mix3D)	69.0

Evaluation on single validation instances

Method	mIoU (on single instances)
MinkowskiNet	24.5 ± 0.88
MinkowskiNet + Mix3D (ours)	36.0 ± 0.57 $\uparrow +11.5$

