Mix3D: Out-of-Context Data Augmentation for 3D Scenes

Conference Paper**Author(s):**

Nekrasov, Alexey; Schult, Jonas; Litany, Or; Leibe, Bastian; [Engelmann, Francis](#)

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000529440>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/3dv53792.2021.00022>

Mix3D: Out-of-Context Data Augmentation for 3D Scenes

Alexey Nekrasov^{1*} Jonas Schult^{1*} Or Litany² Bastian Leibe¹ Francis Engelmann^{1,3}

¹RWTH Aachen University ²NVIDIA ³ETH AI Center

Abstract

We present *Mix3D*, a data augmentation technique for segmenting large-scale 3D scenes. Since scene context helps reasoning about object semantics, current works focus on models with large capacity and receptive fields that can fully capture the global context of an input 3D scene. However, strong contextual priors can have detrimental implications like mistaking a pedestrian crossing the street for a car. In this work, we focus on the importance of balancing global scene context and local geometry, with the goal of generalizing beyond the contextual priors in the training set. In particular, we propose a “mixing” technique which creates new training samples by combining two augmented scenes. By doing so, object instances are implicitly placed into novel out-of-context environments, therefore making it harder for models to rely on scene context alone, and instead infer semantics from local structure as well. We perform detailed analysis to understand the importance of global context, local structures and the effect of mixing scenes. In experiments, we show that models trained with *Mix3D* profit from a significant performance boost on indoor (ScanNet, S3DIS) and outdoor datasets (SemanticKITTI). *Mix3D* can be trivially used with any existing method, e.g., trained with *Mix3D*, MinkowskiNet outperforms all prior state-of-the-art methods by a significant margin on the ScanNet test benchmark (78.1% mIoU). Code is available at: <https://nekrasov.dev/mix3d/>

1. Introduction

In this work, we address the task of 3D semantic segmentation which assigns a semantic label to every point in a given 3D point cloud. Popularized by PointNet [38], the broader field of deep learning for 3D scene understanding has experienced tremendous growth. While these early approaches [38, 39] can classify individual objects or segment larger scenes by splitting them into cubical (or spherical) chunks, recent high-capacity models such as MinkowskiNets [9] and SparseConvNets [22] have large receptive fields and can directly process full rooms or even

Figure 1: Semantic segmentation improvements using Mix3D. Mix3D is a data augmentation technique that increases generalization beyond the contextual priors of the training scenes and improves predictions for rare events not well captured by the training data. Models trained with Mix3D (*bottom*) can handle critical situations where pedestrians walk in the middle of the road (*left*), and can correctly identify objects in unusual settings (*right*).

outdoor scenes without splitting them up. By doing so, they can capture the global *context* of a scene. Context refers to the strong configuration rules present in man-made environments [36]. These rules manifest in re-occurrences of object arrangements. For example, based on context, we can expect to see chairs around a table situated in a conference room. Importantly, context supports reasoning about object semantics, as confirmed by research in visual cognition, computer vision, and cognitive neuroscience [5, 25, 36].

However, while contextual information has proven useful, relying too heavily on global scene priors can be harmful when reasoning about object semantics. Models may overfit to contextual cues of the training set, resulting in poor generalization to unseen data, and in particular can lead to failure on long tail examples (Fig. 1, *right*). Out-of-context generalization is even more important in safety critical settings, such as autonomous driving, where rare but dangerous situations are often not well represented in training sets, e.g., pedestrians walking in the middle of the road or people appearing in the front of a vehicle (Fig. 1, *left*).

* Equal contribution.

3D datasets are usually significantly smaller and exhibit little in-data variance compared to 2D image datasets that are often collected “in the wild” [34]. High-capacity models [9, 22, 49] are particularly vulnerable to learn strong context-based priors when trained on small 3D datasets. As a result, modern network architectures for 3D processing often struggle with overfitting and typically expose large generalization gaps (Fig. 6).

To this end, we propose a new data augmentation technique, called Mix3D, that “mixes” pairs of training scenes. Specifically, mixing consists in randomly transforming pairs of point clouds such that they loosely align, followed by taking the union over the two point sets (see Fig. 2). This is implemented as concatenation of two point clouds as well as their corresponding training labels. Mixing two scenes essentially generates new training samples, that increase the training set, and reduces overfitting as demonstrated in our analysis. In contrast to standard augmentations that apply transformations to single training samples without changing the scene context, Mix3D exposes the objects of each original scene to the novel context of the mixed scenes. The result is an effective data augmentation strategy that can easily be incorporated into existing training pipelines, works well even without labeled data, and effectively reduces overfitting while increasing segmentation performance.

We evaluate Mix3D on three large-scale 3D semantic segmentation benchmarks: ScanNet [10], S3DIS [3] and SemanticKITTI [4]. We compare current state-of-the-art voxel-based MinkowskiNet [9] and point-based KP-Conv [49] trained with and without Mix3D. On all datasets and models, Mix3D consistently improves segmentation performance on both indoor and outdoor data. Our analyses show that mixing scenes is key to consistent improvements and verify that Mix3D effectively reduces overfitting to the contextual priors of the training scenes.

Our contributions are summarized as:

- We propose Mix3D, a data augmentation technique for large-scale 3D scenes and high-capacity models that can balance global context and local geometry. As a result, the trained models can generalize beyond the contextual priors of the training set.
- Applying Mix3D to recent point-based and voxel-based 3D models results in consistently improved segmentation performance on three large-scale indoor and outdoor datasets.
- New state-of-the-art semantic segmentation performance of 78.1% mean intersection over union (mIoU) on the ScanNet test benchmark challenge.
- In-depth analyses of Mix3D show the importance of global context, local geometry and mixing scenes, in terms of model generalization and segmentation performance.

2. Related Work

Mixed Sample Data Augmentation. In contrast to standard augmentations applied to single samples, MixUp pioneered by Zhang *et al.* [62] enforces smooth label transitions by linearly interpolating between pairs of input samples and their corresponding class labels [8, 23, 32, 52, 60, 62, 63]. Training on these virtual training samples encourages linear behavior outside the empirical train distribution, leading the classifier to show fewer label oscillations and, thus, improved generalization [62]. Numerous works extend MixUp: AdaMixUp [23] takes special provisions to prevent manifold intrusions of virtual training examples. Saliency guided approaches [30, 51] and (attentive) CutMix [53, 60] paste image crops into the image and thus increase localization accuracy. In contrast to mixing train samples in the input domain, Verma *et al.* [52] show that interpolating in the latent feature space leads to compressed representations which are theoretically linked to increased generalization capabilities. However, these methods are inherently limited to dense representations such as images, waveforms and spectrograms [62], as they rely on one-to-one correspondences between parts of the input signals, *e.g.*, pixel coordinates in an image. Sparse representations such as point clouds do not provide readily accessible optimal bijective correspondences. Therefore, PointMixUp [8] leverages the Earth Mover’s Distance [43] to interpolate between optimal pairwise assignments that lie on the shortest path between two point clouds. Inspired by (Attentive) CutMix [53, 60] and saliency guided approaches [30, 51], PointCutMix [63] improves upon this work by replacing subsets with corresponding points of other point clouds. Circumventing the problem of establishing bijective correspondences, Achituve *et al.* [1] explore combining point clouds for domain adaptation for single object semantic segmentation. Analogously to Achituve *et al.* [1], we treat augmented point clouds as their own entity, *i.e.*, point features and associated labels do not change in the presence of a second point cloud. Our approach is thus not limited by establishing optimal one-to-one point correspondences of equally sized sets. Moreover, we do not rely on heuristic preprocessing such as aligning the main symmetry axes of point clouds [8]. Thus, our approach is applicable to realistic tasks such as semantic segmentation of full scenes and is not limited to single objects of fixed point set sizes [8, 63].

Context Augmentations. Rosenfeld *et al.* [41] show that object detection models are vulnerable towards placing out-of-context instances in an image and therefore over-rely on context information. Shetty *et al.* [45] thus seek to alleviate the dependency to contextual variations by introducing a removal-based data augmentation where context or objects are cropped out of the image. Alternative methods propose to place novel instances into an image: On the one hand,

[2, 12, 50, 64] constitute a line of work which places objects in the scene at visually *plausible* locations. Dvornik *et al.* [12] show that merely placing additional instances in images in fact leads to decreased performances for 2D object detection. However, by placing instances at semantically plausible positions in an image, they were able to report improved performances. Similar to [12], Tripathi *et al.* [50] propose an adversarial learning strategy that estimates plausible positions for additional instances. Zhang *et al.* [64] improve upon [12, 50] by conditioning the placement of instances not only on context information but also on the instance itself. In the task of 3D semantic segmentation of outdoor scenes, Yan *et al.* [58] carefully populate scenes with car instances taking special care to avoid overlapping objects. A number of other methods ignores the requirement for realistic data and randomly place object instances into the scene [13, 19] providing realistic context information only on a local patch-level. This is comparable to Mix3D as mixing two 3D scenes can also introduce unrealistic merging of objects. In contrast to approaches adding instances into the input image, other approaches [7, 11, 46, 65] explore augmentation methods for removing information from the input image. The key idea is that rectangular regions of the input image are removed which reduces the risk to overfit to small train set specific context features [11]. Unlike works aiming to retain a visually plausible context, we argue that maintaining the context of augmented images amplifies the context bias. We therefore seek to regularize its effects by adding instances together with their respective contexts and voluntarily create context overlaps. Moreover, predicting plausible placements of objects demands complex augmentation systems, whereas our approach is fast, and can be added to existing approaches with a few lines of code (*c.f.* Fig. 3, *right*).

Context Aggregation of Deep Learning Models. Point-based approaches such as [33, 38, 39, 49, 55, 56] perform semantic segmentation by splitting scenes into smaller chunks, effectively restricting the model’s ability to learn from global context. Subsequent works [15, 16, 29, 59] report improvements by increasing the spatial context. By leveraging data efficient sparse convolutions, recent voxel-based methods *e.g.*, MinkowskiNet [9] and SparseConvNet [22] are capable of processing full scenes at once, thus, capturing the global context of the scene. 3D-MPA [14] and DOPS [42] extend on this idea for instance segmentation and object detection. They learn higher-level relationships among object proposals with the goal to learn global scene configuration rules [36]. Context is equally important in recent methods for 3D reconstruction [17, 37]. These developments allow to fully capture and rely on global scene context. However, our experiments show that this can have detrimental implications. By applying Mix3D, we seek to generalize beyond the contextual priors of the training set.

Figure 2: Illustration of two mixed 3D scenes. The Mix3D data augmentation consists in taking the union over two 3D scenes after random transformations while ensuring sufficient overlap. This results in the effect of objects appearing now in the novel context of both mixed scenes, *e.g.*, the classroom table directly faces the apartment door. The goal is to train a model that learns when to rely on context and when to focus on local geometry.

3. Method

We present a data augmentation technique for 3D deep learning on realistic large-scale scenes. The goal is to train 3D models that are less biased by misleading context priors, and learn to balance between local structures and global scene context without overfitting to training scenes.

Mix3D creates novel training examples by mixing two original scenes. The mixing of 3D scenes is illustrated in Fig. 2. In particular, we expose objects from a single input scene to the combined context of both mixed scenes. By doing so, the network does not only need to learn how to implicitly disentangle two mixed scene contexts, but also sees every object in a large variety of object arrangements that one would not normally encounter. Most notable, from initially N different scene contexts, with Mix3D we obtain $\mathcal{O}(N^2)$ novel contexts via pairwise combination of the existing scenes. In the following, we explain the details of the Mix3D training pipeline (Fig. 3, *left*).

Data Augmentation (Fig. 3, ■). When augmenting input scenes, we first translate them to the origin by subtracting the centroid from all point positions. By doing so, we make sure that the two scenes are overlapping in the following mixing stage. We randomly flip the point cloud in both horizontal directions, randomly rotate the scene along the up-right axis and along the other axis by $\text{Uniform}[-\frac{\pi}{64}, +\frac{\pi}{64}]$. We also apply random sub-sampling, elastic distortion [9, 40] and randomly scale the scene by $\text{Uniform}[0.9, 1.1]$. When the input features include color, we additionally apply random brightness and contrast augmentation as well as color-jitter.

Figure 3: (Left) Training pipeline. Mix3D is easy to incorporate into existing code bases: instead of feeding a single scene S_1 into the 3D model, a second input scene S_2 is augmented in parallel and mixed with the first one. The resulting mixed scene $S_{1,2}$ is input to the 3D model, which remains unchanged. For semantic segmentation, we compute the standard cross entropy loss $\mathcal{L}_{CE}$ on the predicted labels and the concatenated ground truth labels $Y_{1,2}$ of both scenes. **(Right) Exemplary implementation of Mix3D.** The individual scenes are augmented by centering at the origin, followed by random rotation and translation while ensuring overlap between both scenes. Additionally, we perform augmentations such as color-jitter, elastic distortion etc. (Details in Sec. 3). Finally, since the order of the points does not change during augmentation, the ground truth semantic labels are obtained using concatenation as well.

Mixing (Fig. 3, ■). After augmenting, both scenes are mixed together, which is simply the union of both augmented point clouds. Technically, we implemented this by concatenating batch entries in pairs. Moreover, since the order of the points is not modified during augmentation, the ground truth labels of the mixed point cloud are also obtained by concatenation (Fig. 3, *right*).

3D Model (Fig. 3, ■). Mix3D is independent of the underlying 3D deep learning model. In experiments, we use MinkowskiNet [9] and KPConv [49] as state-of-the-art representatives of both voxel-based and point-based models. The exact experimental setup is described in more detail in Sec. 4 and the supplementary.

Loss function. We supervise the model for semantic segmentation using the standard cross-entropy loss $\mathcal{L}_{CE}$ which we apply to the predictions of the mixed scene $S_{1,2}$ and the concatenated ground truth labels $Y_{1,2}$ of both scenes.

Discussion. Crucially, Mix3D differs from existing point cloud mixing-based techniques in two main aspects. Firstly, we are the first to explore mixed sample augmentations beyond single object classification for tackling large-scale scene segmentation. We see this as an important step towards improving real-world 3D scene understanding. Second, inspired by MixUp [62], recent techniques propose to interpolate labels [32, 63] and even input samples [8]. While appropriate for a holistic classification task, here we keep the mixed samples and labels in their original form. We motivate this in the following manner: consider the task of predicting the convex combination of two inputs

$(aS_1, (1-a)S_2)$ as $aY_1 + (1-a)Y_2$. This implies that the prediction of Y_1 changes in the presence of Y_2 . Yet, this is the opposite of our aim here, where we wish to keep the prediction of Y_1 agnostic to the existence of Y_2 . In particular, we preserve the complete context information of each mixed sample, whereas PointMixUp [8] distorts point clouds by interpolation and RSMix [32] only preserves the spatial structure of locally restricted chunks. We provide pseudo code in Figure 3 (*right*).

4. Experiments

In this section, we first compare a variety of state-of-the-art models trained with and without Mix3D on three large-scale indoor and outdoor 3D semantic segmentation benchmarks (Sec. 4.1). We provide detailed analysis experiments to motivate and understand the importance of scene context, local geometry and mixing scenes (Sec. 4.2). Finally, we show additional qualitative results. More analysis and results are found in the supplementary.

4.1. Comparison with State-of-the-art Methods

Datasets. S3DIS [3] consists of dense 3D point clouds from 6 large-scale areas including 271 rooms from 3 different buildings. The annotations cover 13 semantic classes. We follow the common evaluation protocol [38, 48, 55] and train on all areas except Area 5, which serves as test-set. ScanNet V2 [10] contains 3D scenes of a wide variety of indoor rooms annotated with 20 semantic classes. We follow the public training, validation, and test split of 1201, 312 and 100 scans, respectively. SemanticKITTI [4]

Model	Mix3D	ScanNet [10]	S3DIS [3]	SemanticKITTI [4]	
		Validation	Area-5	Validation	Test
MinkowskiNet [9]	✗	72.4 ± 0.18	64.7 ± 0.34	56.7	53.2
	✓	73.6 ± 0.31	65.4 ± 0.37	59.9	58.1
Rigid KPConv [49]	✗	68.8 ± 0.10	65.0 ± 0.31	65.1	63.1
	✓	69.5 ± 0.33	66.5 ± 0.44	66.6	63.6
Deformable KPConv [49]	✗	69.3 ± 0.10	65.8 ± 0.47	—	—
	✓	70.3 ± 0.06	67.2 ± 0.28	—	—

Table 1: Effect of Mix3D on models for 3D semantic segmentation. We compare MinkowskiNet (voxel-based) and KPConv (point-based), with and without Mix3D on large-scale indoor scenes (ScanNet, S3DIS) and outdoor scenes (SemanticKITTI).

Method	Mix3D	mIoU	wall	floor	cabinet	bed	chair	sofa	table	door	wndw	books	pic	counter	desk	curtain	fridge	shower	toilet	sink	bathtub	otherf.
MinkNet [9]	✗	72.4	85.6	96.5	64.7	82.1	91.0	84.4	74.5	65.0	62.8	79.5	32.4	64.4	63.7	75.5	51.6	69.0	93.0	67.6	87.6	56.3
(2cm)	✓	73.6	86.0	96.6	66.3	82.1	91.9	86.1	75.7	64.7	64.7	79.6	36.3	67.7	67.0	74.7	59.3	68.8	93.9	68.6	87.9	55.1
KPConv [49]	✗	69.3	82.4	94.4	64.5	79.2	88.5	77.2	73.0	60.5	59.1	79.8	28.4	59.9	63.7	71.6	53.1	54.1	91.5	63.3	86.0	56.4
(Deformable)	✓	70.3	82.6	94.4	65.8	79.5	90.2	79.9	72.5	61.7	58.2	79.6	31.5	62.1	63.2	72.0	58.1	56.8	91.8	64.6	85.6	55.4

Table 2: Semantic segmentation scores on ScanNet validation [10]. We report the mean per-class scores over three trained models.

shows large-scale outdoor traffic scenes recorded with a Velodyne-64 laser scanner. It consists of 19130 training scenes, recorded over 10 different driving sequences, and a validation set of 4071 samples. There are 28 annotated semantic classes from which 19 are evaluated. We report scores in the single scan setting, *i.e.*, individually per frame.

Models in comparison. We apply Mix3D to MinkowskiNet [9] and KPConv [49], as the recent best-performing point-based and voxel-based methods. MinkowskiNet [9] is a voxel-based model that implements the idea of sparse convolutional networks originally presented by Graham *et al.* [22]. KPConv [49] represent the currently best performing point-based approach. Its kernel weights are located in Euclidean space which is more flexible than fixed grid positions. The kernel point positions are either deformable or defined on a rigid pattern. In our experiments, we evaluate both rigid and deformable KPConvs. For all models, we use the official code. For MinkowskiNets on S3DIS and SemanticKITTI, we use our own training pipeline.

Results are presented in Tab. 1 and Tab. 2. For each model, we report the mean (and standard deviation) over three training runs. Training with Mix3D consistently improves previous state-of-the-art models (MinkowskiNet [9], KPConv [49]) across 3 different indoor and outdoor datasets by a statistically significant margin. The most notable improvement is obtained on SemanticKITTI with up to +4.9 mIoU. Importantly, these results are obtained by only adding Mix3D, while keeping training hyper-parameters unchanged, including training time and effective batch size (same number of points per batch) for all datasets. Tab. 2

shows that the class “fridge”, with one of the fewest training samples, improves significantly by +5 mIoU with KPConv and +7.7 mIoU (MinkowskiNet). We also train a MinkowskiNet ensemble with Mix3D and report state-of-the-art scores on the ScanNet benchmark challenge [10] of **78.1 mIoU** (Tab. 3). Each model is trained for 120k iterations, *i.e.*, the same number of training iterations as suggested by Choy *et al.* [9], on both the train and validation split, and we follow [22, 44, 49] for test time augmentation. More architectural and training details, as well as per-class results on all datasets are provided in the supplementary.

ScanNet Semantic Segmentation Benchmark Challenge (hidden test set)			
Method	Input		mIoU
Mix3D (ours)	3DV’21	3D	78.1
OccuSeg [24]	CVPR’20	3D	76.4
O-Net [54]	SIGGRAPH’17	3D	76.2
BPNet [26]	CVPR’21	3D + 2D	74.9
VMNet [27]	ICCV’21	3D + Mesh	74.6
Virtual MVFusion [31]	ECCV’20	3D + 2D	74.6
MinkowskiNet [9]	CVPR’19	3D	73.6
SparseConvNet [22]	CVPR’18	3D	72.5
RFCR [20]	CVPR’21	3D	70.2
JSENet [28]	ECCV’20	3D	69.9
FusionNet [61]	ECCV’20	3D	68.8
KP-Conv [49]	ICCV’19	3D	68.4
PointConv [56]	CVPR’18	3D	66.6
PointASNL [57]	CVPR’20	3D	66.6
DCM-Net [44]	CVPR’20	3D + Mesh	65.8

Table 3: Semantic segmentation results on ScanNet test. We include methods that additionally rely on 2D images and polygon meshes as input. Benchmark accessed on 5th October 2021.

Figure 4: Influence of context during training. Increasing context during training improves semantic segmentation performance on ScanNet validation. However only up to a certain point at which overfitting to scene context limits performance.

4.2. Motivation and Analysis Experiments

In this section, we experimentally motivate and analyze Mix3D. Unless stated otherwise, the model is MinkowskiNet (5 cm voxels) trained on the ScanNet train split and evaluated on the validation split for semantic segmentation.

What is the effect of context? Recent developments focus on models with ever larger receptive fields, such that they can capture larger scene context. The intuition is that global context provides additional information that is helpful for making local decisions. To verify this hypothesis, we artificially reduce the scene context. In particular, we train different models on smaller regions of a scene, *i.e.*, on crops of smaller sizes. The key idea is that, during training, the model sees only a fraction of the original scene, and therefore it cannot rely on global scene context (only on reduced context within the crop). In this experiment, we train MinkowskiNet [9] models on increasing chunk sizes, as well as KPConv [49] models on spheric crops of increasing radius. Indeed, the results in Fig. 4 show that increasing the crop size (and thus the context) is in general helpful. However, we also observe that MinkowskiNet models trained on $1/4$ -fractions of a scene perform on-par to models trained on full scenes. This suggests that relying too much on global context can reduce performance. A similar effect is observed for KPConv where the performance does not significantly increase after 2 m radius. In conclusion, global context does indeed improve performance, however only up to a certain point at which overfitting to context limits the performance.

Does Mix3D help to focus more on local geometry? In the previous experiment, we have seen that increased global context is not necessarily beneficial as overfitting to scene priors can occur. The goal of the next experiment is to demonstrate that models trained with Mix3D are less reliant on context information and are able to make better use of local geometry and structures. For this experiment, the

setup during training remains unchanged, *i.e.*, we train the model with full scenes from the ScanNet training split. At test time, however, we simulate missing context by showing only isolated objects to the trained network. The individual objects are cropped out using the object mask annotations available in the ScanNet validation split. This experiment shows how much the model depends on context, and how much it can make use of local geometry. In Tab. 4, we show semantic segmentation scores for *individual* instances (mean IoU) for two types of models, one trained with, and one trained without Mix3D.

Method	mIoU (on single instances)
MinkowskiNet [9]	24.5 ± 0.88
MinkowskiNet + Mix3D (ours)	36.0 ± 0.57 $\nearrow +11.5$

Table 4: Evaluation on single validation instances. Using the ground truth instance masks, we evaluate on cropped out single instances. Both models are trained on full scenes. The Mix3D model improves over the model trained without Mix3D by +11.5 mIoU. We report mean IoU scores and standard deviations of three trained models each.

While the overall scores are significantly lower than in full-scene validation experiments – showing the importance of global context – we also see that the Mix3D models perform significantly better (+11.5 mIoU) on single instances than models trained without Mix3D. In particular, this means that Mix3D models are notably less dependent on global context to make more accurate predictions given only local geometry. As such, we argue that they are less prone to overfitting on strong context priors and more importantly, Mix3D models can extract more discriminative information from local geometry alone. We conclude: although exposed to larger crops, models trained without Mix3D are limited in their ability to leverage context information fully (*c.f.* Fig. 4). As seen in this experiment, Mix3D trained models, however, balance better between global context and local geometry and are thus less prone to overfit to context cues, resulting in better performance.

Do alternative out-of-context augmentations work? We have already seen that Mix3D influences the scene context during training so that the model overfits less on global context and relies more on local geometry by exposing objects of one scene to the context of another scene, and vice versa. Next, we want to see if it is possible to find a simpler approach that achieves a similar effect as Mix3D. First, we look at random noise as a means of obscuring scene context as suggested in [6]. We compare two different noise patterns: first, we add random points near the surface of 20% of original points, *i.e.*, within a radius of 0.5 m around the original point. The intuition is that this noise pattern only affects

instances locally by obscuring their context. The second approach consists in adding random, uniformly distributed points within the bounding volume of the scene, which can enable non-local information flow. For each 60cm^3 cube we add a point with a 0.1 m random offset. Tab. 5, ① shows the resulting scores: adding these random noise patterns drastically reduces the performance by more than 6.5% mIoU compared to the baseline. Next, we adapt CutOut [11] from the image domain and remove points that fall into randomly sampled cubical chunks. This resembles a *thinning-out* effect [47] on the context level which reduces the risk of overfitting to small training set-specific context features [11]. We evaluate varying cuboid sizes and cutting frequencies (see Tab. 5, ②). The experiments show that CutOut is sensitive to the parameter choice and that Mix3D still outperforms the best performing CutOut model. Last, we mix a second scene without ground truth labels, *i.e.*, the semantic segmentation loss is computed only for points from one scene. While this model (Tab. 5, ③) outperforms the baseline significantly, it does not reach the performance of the fully supervised model. Nevertheless, this reveals an additional strength of Mix3D: the ability to use unlabeled scenes for mixing demands minimal acquisition effort since no human effort for labelling is required.

Augmentation			mIoU
Baseline			66.6
① Noise	Setting	Frequency	
	near surface	20%	51.4
	uniform	$1/60\text{ cm}^3$	60.1
② CutOut	Size	Cut frequency	
	in meters	#cuts/ 10^4 points	
	0.05–2.0	1	67.4
	0.05–2.0	2	67.6
	0.05–1.0	5	67.9
	0.05–1.0	10	66.2
	0.05–0.5	20	66.4
③ Mixing			
	without labels		68.4
with labels (<i>i.e.</i> Mix3D)			69.0

Table 5: Alternative Context Changes. We compare random noise patterns ①, cutting out random chunks ② and mixing scenes with and without annotated labels ③. Mix3D with labels performs best, while mixing without labels is still a viable approach when large amounts of unlabeled data are available or too costly to label.

How important is the scene overlap in Mix3D? The previous experiment established that mixing scenes is key to the improved performance of Mix3D. In the next experiment, we take a closer look at the mixing itself, *i.e.*, the combination of two scenes with increasing overlap. The baseline is a MinkowskiNet trained without Mix3D, *i.e.*, a single scene constitutes a training sample (*c.f.* Tab. 6, ①). As Mix3D technically concatenates batch entries in pairs, we conduct an experiment (*c.f.* Tab. 6, ②) to see the influence of the altered batching, *i.e.*, the concatenation of two scenes that are placed very far away, such that the receptive field of the model does not enable any information flow between the two scenes. The scores are reported in Tab. 6. Perhaps not surprisingly, this setup shows no significant difference to the baseline, as contexts of the scenes do not mutually influence each other. In the next setup (*c.f.* Tab. 6, ③), the scenes are placed nearby, such that information can flow via the receptive field, however without physical overlap of the scenes. We observe an improvement of at least +1.2 mIoU compared to the baseline, as now each scene is influenced indirectly by the context of the neighboring scene via the receptive field of the model. In the last experiment (*c.f.* Tab. 6, ④), we also include physical overlap of the scenes (this corresponds to Mix3D) and we observe the biggest improvement (+2.4 mIoU) compared to the baseline. We conclude that the context overlap caused by mixing scenes is the decisive aspect which explains the performance of Mix3D. In particular, objects of one scene are additionally exposed to the context of the other scene, and vice versa.

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

Table 6: Influence of scene overlap. Overlap of scene-context is the decisive aspect of Mix3D. Placing scenes next to each other without overlaps or too far away for the model’s receptive field to enable mutual information exchange between both scenes.

Figure 5: Qualitative results. We show qualitative results of models trained with and without Mix3D on ScanNet (left), SemanticKITTI (center) and S3DIS (right). Compared to the original model, in (a) Mix3D helps to tell apart two stacked washing machines labeled as “other furniture” from “fridge”. In (b), the MinkowskiNet trained without Mix3D wrongly classifies the “bicycle” on the sidewalk next to the vegetation as “other-object”. For S3DIS, models trained with Mix3D particularly profit in the “window” class, as in (c).

Figure 6: Learning curves with and without Mix3D. We show the mean and standard deviation over three training runs of a MinkowskiNet trained with and without Mix3D evaluated on the ScanNet validation set. The MinkowskiNet is trained on 5 cm voxels including the standard data augmentations: rotation, translation, color-jitter and elastic distortion. The training loss is notably lower without Mix3D (left). The validation loss increases much later for models trained with Mix3D (center) which indicates less overfitting, and results in overall improved performance (right).

Does Mix3D improve generalization? Poor generalization can be caused by overfitting to the training set. Overfitting and large generalization gaps are a considerable problem for 3D scene understanding tasks where training samples are few, and in particular, when combined with high-capacity models [9, 22, 49] that can overfit by memorizing the global context of the training scenes. Learning curves are a useful tool to diagnose overfitting, specifically the generalization error on previously unseen data. Fig. 6 shows the training loss $\mathcal{L}_{\text{train}}$, validation loss $\mathcal{L}_{\text{val}}$ and the mean IoU per training epoch for models trained with and without Mix3D. Mix3D reduces the generalization gap between $\mathcal{L}_{\text{train}}$ and $\mathcal{L}_{\text{val}}$, while $\mathcal{L}_{\text{val}}$ of the Mix3D models increases much later than for models trained without Mix3D. This shows less overfitting to the training set and better generalization on the validation set in terms of semantic segmentation perfor-

mance (mIoU). Increasing the size of the training set is also a simple approach to reduce the generalization error [21]. Mix3D can be seen as a technique to create additional training examples by pairwise combination of existing samples.

Qualitative Results. We show qualitative comparisons of models trained with and without Mix3D. Specifically, in Fig. 1 we show out-of-context scenarios, and Fig. 5 shows predictions on indoor and outdoor datasets.

5. Discussion and Conclusion

In this work, we have introduced Mix3D, a simple yet effective data augmentation technique for large-scale 3D scene segmentation. Models trained with Mix3D learn to better balance global context priors and local structure, as indicated by our analysis. In experiments, we have shown that Mix3D significantly boosts the semantic segmentation performance of multiple state-of-the-art 3D deep learning models on both indoor and outdoor datasets. Importantly, Mix3D is easy to implement and can directly be incorporated into existing training pipelines. It remains an open research question how Mix3D can be applied to approaches combining 3D scans and corresponding 2D images. Further, increased realism could improve the performance, since Mix3D generates unrealistic scenes. Similarly, other tasks, such as object detection or instance segmentation, can be investigated. In future research, we intend to investigate how to adapt the strong generalization abilities of Mix3D for weakly- and self-supervised learning, and we expect to see further developments along this line of work.

Acknowledgments: This work is supported by the ERC Consolidator Grant DeeViSe (ERC-2017-CoG-773161) and compute resources from RWTH Aachen University (rwth0629, thes0775). Contributions from the first author are part of his master thesis.

References

- [1] I. Achituve, H. Maron, and G. Chechik. Self-Supervised Learning for Domain Adaptation on Point Clouds. In *Winter Conference on Applications of Computer Vision (WACV)*, 2021. 2
- [2] H. Alhaija, S. Mustikovela, L. Mescheder, A. Geiger, and C. Rother. Augmented Reality Meets Computer Vision: Efficient Data Generation for Urban Driving Scenes. In *International Journal of Computer Vision (IJCV)*, 2018. 3
- [3] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. Savarese. 3D Semantic Parsing of Large-Scale Indoor Spaces. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 4, 5, 8, 12, 15
- [4] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 4, 5, 8, 12, 15
- [5] I. Biederman. Perceiving Real-World Scenes. *Science*, 1972. 1
- [6] Chappelle, Olivier and Weston, Jason and Bottou, Léon and Vapnik, Vladimir. Vicinal Risk Minimization. In *Neural Information Processing Systems (NIPS)*, 2001. 6
- [7] P. Chen, S. Liu, H. Zhao, and J. Jia. GridMask Data Augmentation. *arXiv preprint arXiv:2001.04086*, 2020. 3
- [8] Y. Chen, V. T. Hu, E. Gavves, T. Mensink, P. Mettes, P. Yang, and C. Snoek. PointMixup: Data Augmentation for Point Clouds. In *European Conference on Computer Vision (ECCV)*, 2020. 2, 4, 12
- [9] C. Choy, J. Y. Gwak, and S. Savarese. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 3, 4, 5, 6, 8, 12, 15
- [10] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 2, 4, 5, 8, 12, 13, 15
- [11] T. DeVries and G. W. Taylor. Improved Regularization of Convolutional Neural Networks with Cutout. *arXiv preprint arXiv:1708.04552*, 2017. 3, 7
- [12] N. Dvornik, J. Mairal, and C. Schmid. Modeling Visual Context is Key to Augmenting Object Detection Datasets. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [13] D. Dwibedi, I. Misra, and M. Hebert. Cut, Paste and Learn: Surprisingly Easy Synthesis for Instance Detection. In *International Conference on Computer Vision (ICCV)*, 2017. 3, 13
- [14] F. Engelmann, M. Bokeloh, A. Fathi, B. Leibe, and M. Nießner. 3D-MPA: Multi Proposal Aggregation for 3D Semantic Instance Segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [15] F. Engelmann, T. Kontogianni, A. Hermans, and B. Leibe. Exploring Spatial Context for 3D Semantic Segmentation of Point Clouds. In *International Conference on Computer Vision Workshop (ICCV'W)*, 2017. 3
- [16] F. Engelmann, T. Kontogianni, and B. Leibe. Dilated Point Convolutions: On the Receptive Field Size of Point Convolutions on 3D Point Clouds. In *International Conference on Robotics and Automation (ICRA)*, 2020. 3
- [17] F. Engelmann, K. Rematas, B. Leibe, and V. Ferrari. From Points to Multi-Object 3D Reconstruction. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [18] P. F. Felzenszwalb and D. P. Huttenlocher. Efficient Graph-based Image Segmentation. *International Journal on Computer Vision (IJCV)*, 2004. 12
- [19] G. Ghiasi, Y. Cui, A. Srinivas, R. Qian, T.-Y. Lin, E. D. Cubuk, Q. V. Le, and B. Zoph. Simple Copy-Paste is a Strong Data Augmentation Method for Instance Segmentation. *arXiv preprint arXiv:2012.07177*, 2020. 3, 13
- [20] J. Gong, J. Xu, X. Tan, H. Song, Y. Qu, Y. Xie, and L. Ma. Omni-supervised Point Cloud Segmentation via Gradual Receptive Field Component Reasoning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5
- [21] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. www.deeplearningbook.org. 8
- [22] B. Graham, M. Engelcke, and L. van der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2, 3, 5, 8, 12
- [23] H. Guo, Y. Mao, and R. Zhang. MixUp as Locally Linear Out-of-Manifold Regularization. In *Conference on Artificial Intelligence (AAAI)*, 2019. 2
- [24] L. Han, T. Zheng, L. Xu, and L. Fang. OccuSeg: Occupancy-aware 3D Instance Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5, 12
- [25] H. S. Hock, G. P. Gordon, and R. Whitehurst. Contextual Relations: The Influence of Familiarity, Physical Plausibility, and Belongingness. *Perception and Psychophysics*, 1974. 1
- [26] W. Hu, H. Zhao, L. Jiang, J. Jia, and T.-T. Wong. Bidirectional Projection Network for Cross Dimension Scene Understanding. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 5
- [27] Z. Hu, X. Bai, J. Shang, R. Zhang, J. Dong, X. Wang, G. Sun, H. Fu, and C.-L. Tai. VMNet: Voxel-Mesh Network for Geodesic-Aware 3D Semantic Segmentation. In *International Conference on Computer Vision (ICCV)*, 2021. 5
- [28] Z. Hu, M. Zhen, X. Bai, H. Fu, and C.-l. Tai. JSENet: Joint Semantic Segmentation and Edge Detection Network for 3D Point Clouds. In *European Conference on Computer Vision (ECCV)*, 2020. 5
- [29] Q. Huang, W. Wang, and U. Neumann. Recurrent Slice Networks for 3D Segmentation of Point Clouds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [30] J.-H. Kim, W. Choo, and H. O. Song. Puzzle Mix: Exploiting Saliency and Local Statistics for Optimal Mixup. In *International Conference on Machine Learning (ICML)*, 2020. 2
- [31] A. Kundu, X. Yin, A. Fathi, D. Ross, B. Brewington, T. Funkhouser, and C. Pantofaru. Virtual Multi-view Fusion for 3D Semantic Segmentation. In *European Conference on Computer Vision (ECCV)*, 2020. 5
- [32] D. Lee, J. Lee, J. Lee, H. Lee, M. Lee, S. Woo, and S. Lee. Regularization Strategy for Point Cloud via Rigidly Mixed Sample. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 4

- [33] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. PointCNN: Convolution On X-Transformed Points. In *Neural Information Processing Systems (NIPS)*, 2018. 3
- [34] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollar, and L. Zitnick. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision (ECCV)*, 2014. 2
- [35] Z. Liu, X. Qi, and C.-W. Fu. One Thing One Click: A Self-Training Approach for Weakly Supervised 3D Semantic Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 12
- [36] A. Oliva and A. Torralba. The Role of Context in Object Recognition. In *Trends in Cognitive Sciences*, 2007. 1, 3
- [37] S. Popov, P. Bauszat, and V. Ferrari. CoReNet: Coherent 3D Scene Reconstruction from a Single RGB Image. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [38] C. R. Qi, H. Su, K. Mo, and L. J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 3, 4
- [39] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Neural Information Processing Systems (NIPS)*, 2017. 1, 3
- [40] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, 2015. 3
- [41] A. Rosenfeld, R. Zemel, and J. K. Tsotsos. The Elephant in the Room. *arXiv preprint arXiv:1808.03305*, 2018. 2
- [42] D. A. R. Ross, G. L. Lai, Z. L. Lu, A. K. Kundu, A. F. Fathi, T. F. Funkhouser, C. P. Pantofaru, M. N. Najibi, L. S. D. Davis, and V. R. Rathod. DOPS: Learning to Detect 3D Objects and Predict Their 3D Shapes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [43] Y. Rubner, C. Tomasi, and L. J. Guibas. The Earth Mover's Distance as a Metric for Image Retrieval. In *International Journal of Computer Vision (IJCV)*, 2000. 2
- [44] J. Schult*, F. Engelmann*, T. Kontogianni, and B. Leibe. DualConvMesh-Net: Joint Geodesic and Euclidean Convolutions on 3D Meshes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5, 12
- [45] R. Shetty, B. Schiele, and M. Fritz. Not Using the Car to See the Sidewalk Quantifying and Controlling the Effects of Context in Classification and Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [46] K. K. Singh, H. Yu, A. Sarmasi, G. Pradeep, and Y. Lee. Hide-and-Seek: A Data Augmentation Technique for Weakly-Supervised Localization and Beyond. *arXiv preprint arXiv:1811.02545*, 2018. 3
- [47] N. Srivastava, G. E. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks From Overfitting. In *Journal of Machine Learning Research (JMLR)*, 2014. 7
- [48] L. P. Tchapmi, C. B. Choy, I. Armeni, J. Gwak, and S. Savarese. SEGCloud: Semantic Segmentation of 3D Point Clouds. In *International Conference on 3D Vision (3DV)*, 2017. 4
- [49] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. In *International Conference on Computer Vision (ICCV)*, 2019. 2, 3, 4, 5, 6, 8, 12, 13, 15
- [50] S. Tripathi, S. Chandra, A. Agrawal, A. Tyagi, J. M. Rehg, and V. Chari. Learning to Generate Synthetic Data via Compositing. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [51] A. F. M. S. Uddin, M. S. Monira, W. Shin, T. Chung, and S.-H. Bae. SaliencyMix: A Saliency Guided Data Augmentation Strategy for Better Regularization. In *International Conference on Learning Representations (ICLR)*, 2021. 2
- [52] V. Verma, A. Lamb, C. Beckham, A. Najafi, I. Mitliagkas, D. Lopez-Paz, and Y. Bengio. Manifold Mixup: Better Representations by Interpolating Hidden States. In *International Conference on Machine Learning (ICML)*, 2019. 2, 12
- [53] D. Walawalkar, Z. Shen, Z. Liu, and M. Savvides. Attentive CutMix: An Enhanced Data Augmentation Approach for Deep Learning Based Image Classification. *arXiv preprint arXiv:2003.13048*, 2020. 2
- [54] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong. O-CNN: Octree-based Convolutional Neural Networks for 3D Shape Analysis. In *ACM Transactions on Graphics (SIGGRAPH)*, 2017. 5
- [55] S. Wang, S. Suo, W. Ma, A. Pokrovsky, and R. Urtasun. Deep Parametric Continuous Convolutional Neural Networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3, 4
- [56] W. Wu, Z. Qi, and F. Li. PointConv: Deep Convolutional Networks on 3D Point Clouds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3, 5
- [57] X. Yan, C. Zheng, Z. Li, S. Wang, and S. Cui. PointASNL: Robust Point Clouds Processing Using Nonlocal Neural Networks with Adaptive Sampling. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [58] Y. Yan, Y. Mao, and B. Li. SECOND: Sparsely Embedded Convolutional Detection. In *Sensors*, 2018. 3, 13
- [59] X. Ye, J. Li, H. Huang, L. Du, and X. Zhang. 3D Recurrent Neural Networks with Context Fusion for Point Cloud Semantic Segmentation. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [60] S. Yun, D. Han, S. J. Oh, S. Chun, J. Choe, and Y. Yoo. CutMix: Regularization Strategy to Train Strong Classifiers with Localizable Features. In *International Conference on Computer Vision (ICCV)*, 2019. 2
- [61] F. Zhang, J. Fang, B. Wah, and P. Torr. Deep Fusionnet for Point Cloud Semantic Segmentation. In *European Conference on Computer Vision (ECCV)*, 2020. 5
- [62] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations (ICLR)*, 2018. 2, 4, 12
- [63] J. Zhang, L. Chen, B. Ouyang, B. Liu, J. Zhu, Y. Chen, Y. Meng, and D. Wu. PointCutMix: Regularization Strategy for Point Cloud Classification. *arXiv preprint arXiv:2101.01461*, 2021. 2, 4

- [64] L. Zhang, T. Wen, J. Min, J. Wang, D. Han, and J. Shi. Learning Object Placement by Inpainting for Compositional Data Augmentation. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [65] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random Erasing Data Augmentation. In *Conference on Artificial Intelligence (AAAI)*, 2020. 3

Mix3D: Out-of-Context Data Augmentation for 3D Scenes

Supplementary Material

Abstract

This supplementary material provides more details on the network architectures, including training and inference details (Sec. A). In Sec. 4.2, we provide additional analysis, experiments and results. Finally, in Sec. C we report per-class semantic segmentation scores on all three datasets (ScanNet, S3DIS, SemanticKITTI).

A. Details on Architectures and Training

MinkowskiNet [9] In Tab. 7, we provide the architectural details of the MinkowskiNet [9] model used for the reported scores on the datasets ScanNet [10], S3DIS [3] and SemanticKITTI [4]. We use the same model definitions as provided in the official code release of MinkowskiNet¹. For ScanNet, we use 2 cm voxels for comparing to state-of-the-art methods and for the benchmark results (see Tab. 1-3). All other analysis experiments are performed on 5 cm voxels for faster training.

KPConv [49] We use the official TensorFlow code release² for rigid and deformable KPConv on ScanNet and S3DIS. For rigid KPConv on SemanticKITTI, we utilize the PyTorch reimplementation³ by the original authors. Code for deformable KPConv on SemanticKITTI is not directly available. In particular, we follow the same evaluation procedure as KPConv. That is, we perform 100 evaluation runs for ScanNet and S3DIS (20 runs for SemanticKITTI) on augmented validation samples and conduct majority voting on the resulting predictions. Note that we do not apply Mix3D during test time, *i.e.*, during inference the trained networks operate one single scenes.

Training details. As Mix3D combines scenes, the total number of points in a single training sample consists of the sum over the number of all points of each mixed scene. To ensure a fair comparison between MinkowskiNets, we appropriately reduce the batch size of Mix3D for keeping the total number of points in a training iteration equal to the one of the baseline method. For instance, we train the 2 cm baseline MinkowskiNet on a batch size of 6. When training with Mix3D, we reduce the batch size to 3 as each training sample consists of 2 mixed scenes. By doing so, we ensure that each model is trained on the same number of labeled points and therefore the results are comparable. For KPConv, we resort to approximately keeping the total num-

ber of points in a train iteration equal. Here, the original codebase operates on a variable batch size with a variable number of points in each sample. We therefore occasionally need to reject samples for Mix3D in order to mix the specified number of spheres. To ensure comparability to the baseline method, we follow the same training protocol proposed by Thomas *et al.* [49] and evaluate after the same number of epochs.

Inference settings. For the ScanNet benchmark submission, we train an ensemble of four MinkowskiNet [9] models with Mix3D. In particular, the four models are Res16UNet34C and Res16UNet34A, each with the first convolution kernel size set to 3 and 5. We train each model for $12 \cdot 10^4$ iterations on both the official training and validation splits of ScanNet [10]. Similar to [22, 44, 49], we perform test-time-augmentation (translation, rotation, scaling, color-jitter, random brightness and contrast) and average the predictions of 40 runs for each model. To obtain the final predictions, we employ a graph-based over-segmentation scheme [18] (similar to OccuSeg [24] and [35]) with the goal to further smooth the predicted labels.

B. Additional Analysis

Influence of the number of mixed scenes. In this experiment, we vary the number of mixed scenes to evaluate if mixing more than two scenes further improves the segmentation performance. Tab. 8 shows that the validation performance peaks at two mixed scenes and noticeably decreases with more scenes. This is in line with other mixed sample data augmentation works [8, 52, 62].

Dataset	Voxel Size	Model	Input Features
ScanNet [10]	2 cm	Res16UNet34A	RGB
	5 cm	Res16UNet34C	RGB
S3DIS [3]	5 cm	Res16UNet34C	RGB
Sem.KITTI [4]	15 cm	Res16UNet14A	Reflect., Dist.

Table 7: MinkowskiNet [9] architectures. We provide architectural details about the MinkowskiNetworks used on ScanNet, S3DIS and SemanticKITTI. In particular, we use standard MinkowskiNets provided in the official repository and did not change the model definitions. Our ablation study provided in the main paper is performed on 5 cm voxels.

¹www.github.com/chrischoy/SpatioTemporalSegmentation-ScanNet

²www.github.com/HuguesTHOMAS/KPConv

³www.github.com/HuguesTHOMAS/KPConv-PyTorch

# non-mixed spheres	mIoU ($\pm \sigma$)
0	69.7 $\pm$ 0.11
1	70.1 $\pm$ 0.35
2	70.3 $\pm$ 0.06
3	69.9 $\pm$ 0.27
4	70.3 $\pm$ 0.11
baseline	69.3 $\pm$ 0.10

Table 9: Influence of non-mixed spheres in a batch. We observe that exposing the network also to non-mixed training samples has a significant impact on Deformable KPConv’s performance on the ScanNet validation set [10]. Exposing the model only to mixed scenes achieves an improvement of +0.4 mIoU, whereas having 2 or 4 non-mixed batch entries yields an improvement of +1.0 mIoU. Each batch contains about 10 spheres, depending on the number of points in each batch (see [49] for details).

Dataset split	Baseline	+Mix3D	<i>abs.</i> Δ	<i>rel.</i> Δ
5 %	43.6	45.6	+2.0	+4.6 %
25 %	59.3	62.1	+2.8	+4.7 %
50 %	63.0	65.7	+2.7	+4.3 %
75 %	65.2	67.3	+2.1	+3.2 %
100 %	66.6	69.0	+2.4	+3.6 %

Table 10: Consistent improvements wrt. dataset size. We report consistent absolute performance gains in the range of +[2.0, 2.8] mIoU for all subsets of the ScanNet train set.

# Scenes	1	2	3	4	7	8
mIoU	66.6	69.0	68.6	68.7	68.3	65.7

Table 8: Varying number of mixed point clouds. While individual objects on two mixed point clouds (b) can still be easily distinguished, mixing eight point clouds (d) results in scenes where individual objects are hardly visually distinguishable.

Influence of non-mixed spheres in a batch. On ScanNet, Thomas *et al.* [49] propose to train KPConv on spheres with 2 m radii cropped from the original point cloud. As Mix3D is implemented as concatenating subsequent pairs of batch entries, we will always obtain fully overlapping spheres in each novel training sample. However, when training MinkowskiNets, we mix full *scenes*. This results in overlapping areas, but also in unmixed areas (*c.f.* Fig. 2). We therefore perform an ablation study with KPConv to measure the performance influence of adding non-mixed spheres to a batch, effectively imitating un-mixed regions. In Tab. 9, we observe that exposing the network also to non-mixed training samples improves the performance upon the initial Mix3D implementation by up to 0.6 mIoU.

Consistent improvements wrt. dataset size. In Tab. 10, we simulate a scarce data setting by training MinkowskiNet on randomly sampled subsets of the ScanNet train set and evaluate it on the full validation set. We observe that Mix3D achieves consistent absolute improvements (*abs.* Δ) over all dataset sizes in the range of [2.0, 2.8] mIoU, while the relative improvement (*rel.* Δ) becomes increasingly larger for smaller datasets. This shows the increased benefit of Mix3D when only few training samples are available.

Comparison to single-instance mixing. In Tab. 11, we investigate an alternative formulation of Mix3D. Instead of mixing complete scenes, we consider mixing single instances randomly sampled from an instance database. This approach explicitly aligns with the key motivation of Mix3D – reducing the context bias by placing instances in novel out-of-context environments. Instance augmentation as proposed in 2D methods [13, 19] or 3D object detection [58] has not yet been investigated for the task of 3D semantic segmentation. As such, we suggest that mixing single instances constitutes an interesting avenue worth investigating. To this end, we create a database consisting of instances from the training set. Next, we place randomly sampled instances from this database in the scene. In particular, we evaluate two strategies for instance placing: ① creating overlapping instances by placing randomly sampled instances exactly at the center position of arbitrary instances of the scene and ② randomly placing instances within the bounding scene volume. We evaluate two sampling ratios. For a sampling ratio of +1.0x, we sample exactly one other random instance from the instance database for each original instance in the scene. For a sampling ratio of +0.5x, we sample half the number of original instances from the instance set. We observe that single-instance mixing’s performance is in the interval of three Mix3D runs and is therefore comparable to Mix3D. Perhaps unsurpris-

	placing method	sampling ratio	mIoU
①	overlapping instances	+1.0x	67.7
		+0.5x	67.9
②	free placing	+1.0x	68.4
		+0.5x	69.2
	Baseline		66.6 ± 0.09
	Mix3D		69.0 ± 0.32

Table 11: Mixing with single instances. For every instance originally occurring in the scene we add 1x or 0.5x as many instances. We place these newly sampled instances either at positions of already existing instances, *i.e.* *overlapping*, or freely in the volume of the scene, *i.e.* *free space*, where we might create overlaps with existing instances. Above, we show an augmented scene with instances freely placed in the bounding volume with a sampling ratio of +1.0x. We observe that single-instance mixing’s performance is in the interval of three Mix3D runs.

ingly, single-instance mixing also improves significantly over the baseline, as the key idea of instance mixing aligns with the motivation of Mix3D. However, Mix3D is easier to implement into existing pipelines, and avoids the following fundamental disadvantages of single-instance mixing: (1) In contrast to Mix3D, single-instance mixing requires a database of object instances available during training to augment the scenes. (2) The construction of such a database requires instance level annotations which are not necessarily available in semantic segmentation datasets. As an alternative, one could rely on synthetic CAD models (*e.g.* from ShapeNet). Mix3D does not require additional annotations or other datasets. (3) Single-Instance mixing requires a dedicated instance placing procedure introducing potentially sensitive hyper-parameters, while Mix3D simply concatenates overlapping scenes. (4) Finally, by mixing the scene with instances, we effectively train on more data points, *i.e.*, we introduce a compounding factor of exposing the model to instances of the training set more often as we do for Mix3D. This makes a fair comparison harder.

C. Per-Class Quantitative Evaluation

Tab. 12-14 show per-class semantic segmentation results on all three datasets and different models trained with and without Mix3D. Moreover, in Tab. 15, we extend Tab. 4 and provide the per-class segmentation scores on isolated instances of the ScanNet validation set. When no scene context is available, a model trained with Mix3D outperforms the same model trained without Mix3D on 18 out of 20 semantic classes, indicating that Mix3D enables models to extract more meaningful information from local geometry and structures alone.

Method	Mix3D	mIoU	wall	floor	cabinet	bed	chair	sofa	table	door	wndw	books	pic	counter	desk	curtain	fridge	shower	toilet	sink	bathtub	otherf.
MinkNet [9]	✗	66.7	81.0	93.8	61.0	79.7	88.6	79.6	71.9	55.6	57.54	75.4	25.4	55.0	61.3	61.3	44.2	61.6	85.9	61.8	79.3	53.8
(5cm)	✓	69.1	82.1	93.9	63.9	79.2	89.2	81.1	70.8	59.3	59.8	78.1	25.6	62.2	62.1	67.0	54.1	64.0	89.8	64.1	81.1	54.3
MinkNet [9]	✗	72.4	85.6	96.5	64.7	82.1	91.0	84.4	74.5	65.0	62.8	79.5	32.4	64.4	63.7	75.5	51.6	69.0	93.0	67.6	87.6	56.3
(2cm)	✓	73.6	86.0	96.6	66.3	82.1	91.9	86.1	75.7	64.7	64.7	79.6	36.3	67.7	67.0	74.7	59.3	68.8	93.9	68.6	87.9	55.1
KPConv [49]	✗	68.8	82.2	94.4	63.4	77.9	88.3	76.6	72.8	60.0	58.4	80.1	28.6	59.3	62.7	69.9	53.1	51.5	90.6	63.4	86.4	55.3
(Rigid)	✓	69.5	82.5	94.6	65.3	79.8	89.7	79.0	72.2	61.0	58.7	80.1	30.6	60.0	62.8	70.7	55.3	50.2	91.5	65.0	85.3	55.6
KPConv [49]	✗	69.3	82.4	94.4	64.5	79.2	88.5	77.2	73.0	60.5	59.1	79.8	28.4	59.9	63.7	71.6	53.1	54.1	91.5	63.3	86.0	56.4
(Deformable)	✓	70.3	82.6	94.4	65.8	79.5	90.2	79.9	72.5	61.7	58.2	79.6	31.5	62.1	63.2	72.0	58.1	56.8	91.8	64.6	85.6	55.4

Table 12: ScanNet [10] validation set. Per-class semantic segmentation IoU scores. We report the mean performance on the validation set over 3 trained models.

Method	Mix3D	mIoU	ceiling	floor	wall	beam	column	window	door	chair	table	bookshelf	sofa	board	clutter
MinkNet [9]	✗	64.8	92.6	96.4	82.0	0.2	28.1	53.3	64.1	87.7	78.5	73.8	65.5	63.1	56.4
(5cm)	✓	65.5	93.5	97.5	82.8	0.0	24.3	57.5	64.8	88.6	80.7	72.9	62.6	68.3	57.5
KPConv [49]	✗	65.0	93.6	98.3	81.9	0.0	19.4	54.8	64.7	91.1	81.4	74.4	64.2	63.0	58.0
(Rigid)	✓	66.5	93.9	98.4	82.9	0.0	19.4	59.4	69.0	91.3	81.6	75.6	65.7	68.5	59.0
KPConv [49]	✗	65.8	93.8	98.4	82.2	0.0	20.7	55.8	68.9	91.3	81.9	75.3	64.9	63.5	58.8
(Deformable)	✓	67.2	93.9	98.4	83.5	0.0	22.3	60.8	66.3	91.7	82.3	76.2	66.7	70.7	60.5

Table 13: S3DIS [3] Area 5. Per-class semantic segmentation IoU scores. We report the mean over 3 trained models.

Method	Mix3D	mIoU	car	bike	mbike	truck	vehicle	person	cyclist	mcyclist	road	parking	sidewalk	other-gr	building	fence	veget	trunk	terrain	pole	sign
MinkNet [9]	✗	53.2	94.0	26.4	24.5	27.5	18.4	40.5	46.7	13.5	88.4	57.1	71.4	22.6	90.4	62.5	83.5	65.3	65.8	54.0	59.1
	✓	58.1	95.1	30.6	38.3	43.5	35.3	50.8	51.8	21.9	89.4	59.5	73.0	23.2	90.9	65.2	83.3	68.0	65.4	57.6	61.7
KPConv [49]	✗	63.1	95.9	45.2	50.9	44.4	45.4	62.9	64.7	36.6	89.4	65.1	73.5	31.3	91.7	67.5	84.4	69.4	68.4	57.9	54.4
	✓	63.6	96.4	49.1	59.7	40.5	52.1	69.6	69.3	18.2	89.4	66.2	74.3	28.0	91.5	66.2	84.4	71.0	67.9	58.7	54.9

Table 14: SemanticKITTI [4] test set. Per-class semantic segmentation IoU scores.

Mix3D	mIoU	wall	floor	cabinet	bed	chair	sofa	table	door	wndw	books	pic	counter	desk	curtain	fridge	shower	toilet	sink	bathtub	otherf.
✗	24.58	57.95	85.76	6.29	34.64	65.90	66.28	13.91	3.26	20.61	52.65	0.78	0.26	8.63	19.72	3.89	0.11	10.65	7.51	16.28	16.47
	± 0.89	± 0.97	± 0.43	± 0.56	± 4.46	± 5.36	± 3.70	± 6.40	± 2.11	± 3.46	± 2.96	± 0.74	± 0.34	± 3.92	± 1.23	± 5.50	± 0.10	± 6.59	± 5.66	± 4.30	± 2.38
✓	36.00	65.09	87.40	10.73	43.13	80.91	79.49	33.89	6.49	27.94	59.35	6.35	0.51	19.28	40.71	3.08	0.00	52.91	17.55	60.42	24.76
	± 0.58	± 0.33	± 1.91	± 0.81	± 1.74	± 0.92	± 1.42	± 4.98	± 0.50	± 4.46	± 0.98	± 4.36	± 0.49	± 6.51	± 1.27	± 1.30	± 0.00	± 8.65	± 8.90	± 6.27	± 3.23

Table 15: Isolated instances. Per-class semantic segmentation IoU scores on the ScanNet validation set [10]. We extend Table 2 of the main paper by providing per-class semantic segmentation scores. We report performance improvements on 18 out of 20 semantic classes. The study is performed with MinkowskiNet on 5 cm voxels and we train 3 models to obtain the means and standard deviations.