

PointMixup: Augmentation for Point Clouds

Yunlu Chen^{*1}, Vincent Tao Hu^{*1}, Efstratios Gavves¹, Thomas Mensink^{2,1},
Pascal Mettes¹, Pengwan Yang^{1,3}, and Cees G. M. Snoek¹

¹ University of Amsterdam

² Google Research, Amsterdam

³ Peking University

Abstract. This paper introduces data augmentation for point clouds by interpolation between examples. Data augmentation by interpolation has shown to be a simple and effective approach in the image domain. Such a mixup is however not directly transferable to point clouds, as we do not have a one-to-one correspondence between the points of two different objects. In this paper, we define data augmentation between point clouds as a shortest path linear interpolation. To that end, we introduce PointMixup, an interpolation method that generates new examples through an optimal assignment of the path function between two point clouds. We prove that our PointMixup finds the shortest path between two point clouds and that the interpolation is assignment invariant and linear. With the definition of interpolation, PointMixup allows to introduce strong interpolation-based regularizers such as mixup and manifold mixup to the point cloud domain. Experimentally, we show the potential of PointMixup for point cloud classification, especially when examples are scarce, as well as increased robustness to noise and geometric transformations to points. The code for PointMixup and the experimental details are publicly available¹.

Keywords: interpolation, point cloud classification, data augmentation

1 Introduction

The goal of this paper is to classify a cloud of points into their semantic category, be it an airplane, a bathtub or a chair. Point cloud classification is challenging, as they are sets and hence invariant to point permutations. Building on the pioneering PointNet by Qi *et al.* [15], multiple works have proposed deep learning solutions to point cloud classification [16, 12, 29, 30, 36, 23]. Given the progress in point cloud network architectures, as well as the importance of data augmentation in improving classification accuracy and robustness, we study how could data augmentation be naturally extended to support also point cloud data, especially considering the often smaller size of point clouds datasets (*e.g.* ModelNet40 [31]). In this work, we propose point cloud data augmentation by interpolation of existing training point clouds.

^{*} Equal contribution. {y.chen3, t.hu}@uva.nl

¹ Code is available at: <https://github.com/yunlu-chen/PointMixup/>

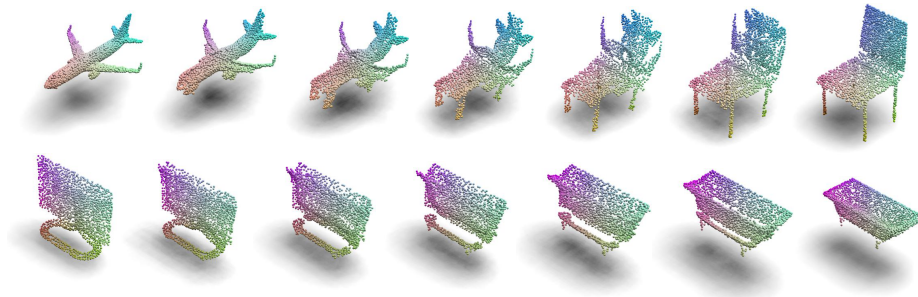


Fig. 1: **Interpolation between point clouds.** We show the interpolation between examples from different classes (airplane/chair, and monitor/bathtub) with multiple ratios λ . The interpolants are learned to be classified as $(1 - \lambda)$ the first class and λ the second class. The interpolation is not obtained by learning, but induced by solving the optimal bijective correspondence which allows the minimum overall distance that each point in one point cloud moves to the assigned point in the other point cloud.

To perform data augmentation by interpolation, we take inspiration from augmentation in the image domain. Several works have shown that generating new training examples, by interpolating images and their corresponding labels, leads to improved network regularization and generalization, *e.g.*, [8, 24, 34, 26]. Such a mixup is feasible in the image domain, due to the regular structure of images and one-to-one correspondences between pixels. However, this setup does not generalize to the point cloud domain, since there is no one-to-one correspondence and ordering between points. To that end, we seek to find a method to enable interpolation between permutation invariant point sets.

In this work, we make three contributions. First, we introduce data augmentation for point clouds through interpolation and we define the augmentation as a shortest path interpolation. Second, we propose PointMixup, an interpolation between point clouds that computes the optimal assignment as a path function between two point clouds, or the latent representations in terms of point cloud. The proposed interpolation strategy therefore allows usage of successful regularizers of Mixup and Manifold Mixup [26] on point cloud. We prove that (i) our PointMixup indeed finds the shortest path between two point clouds; (ii) the assignment does not change for any pairs of the mixed point clouds for any interpolation ratio; and (iii) our PointMixup is a linear interpolation, an important property since labels are also linearly interpolated. Figure 1 shows two pairs of point clouds, along with our interpolations. Third, we show the empirical benefits of our data augmentation across various tasks, including classification, few-shot learning, and semi-supervised learning. We furthermore show that our approach is agnostic to the network used for classification, while we also become more robust to noise and geometric transformations to the points.

2 Related Work

Deep learning for point clouds. Point clouds are unordered sets and hence early works focus on analyzing equivalent symmetric functions which ensures permutation invariance. [17, 15, 33]. The pioneering PointNet work by Qi *et al.* [15] presented the first deep network that operates directly on unordered point sets. It learns the global feature with shared multi-layer perceptions and a max pooling operation to ensure permutation invariance. PointNet++ [16] extends this idea further with hierarchical structure by relying on a heuristic method of farthest point sampling and grouping to build the hierarchy. Likewise, other recent methods follow to learn hierarchical local features either by grouping points in various manners [10, 12, 29, 30, 32, 36, 23]. Li *et al.* [12] propose to learn a transformation from the input points to simultaneously solve the weighting of input point features and permutation of points into a latent and potentially canonical order. Xu *et al.* [32] extends 2D convolution to 3D point clouds by parameterizing a family of convolution filters. Wang *et al.* [29] proposed to leverage neighborhood structures in both point and feature spaces.

In this work, we aim to improve point cloud classification for any point-based approach. To that end, we propose a new model-agnostic data augmentation. We propose a Mixup regularization for point clouds and show that it can build on various architectures to obtain better classification results by reducing the generalization error in classification tasks. A very recent work by Li *et al.* [11] also considers improving point cloud classification by augmentation. They rely on auto-augmentation and a complicated adversarial training procedure, whereas in this work we propose to augment point clouds by interpolation.

Interpolation-based regularization. Employing regularization approaches for training deep neural networks to improve their generalization performances have become standard practice in deep learning. Recent works consider a regularization by interpolating the example and label pairs, commonly known as Mixup [24, 8, 34]. Manifold Mixup [26] extends Mixup by interpolating the hidden representations at multiple layers. Recently, an effort has been made on applying Mixup to various tasks such as object detection [35] and segmentation [7]. Different from existing works, which are predominantly employed in the image domain, we propose a new optimal assignment Mixup paradigm for point clouds, in order to deal with their permutation-invariant nature.

Recently, Mixup [34] has also been investigated from a semi-supervised learning perspective [3, 27, 2]. Mixmatch [3] guesses low-entropy labels for unlabelled data-augmented examples and mixes labelled and unlabelled data using Mixup [34]. Interpolation Consistency Training [27] utilizes the consistency constraint between the interpolation of unlabelled points with the interpolation of the predictions at those points. In this work, we show that our PointMixup can be integrated in such frameworks to enable semi-supervised learning for point clouds.

3 Point cloud augmentation by interpolation

3.1 Problem setting

In our setting, we are given a training set $\{(S_m, c_m)\}_{m=1}^M$ consisting of M point clouds. $S_m = \{p_n^m\}_{n=1}^N \in \mathcal{S}$ is a point cloud consisting of N points, $p_n^m \in \mathbb{R}^3$ is the 3D point, \mathcal{S} is the set of such 3D point clouds with N elements. $c_m \in \{0, 1\}^C$ is the one-hot class label for a total of C classes. The goal is to train a function $h : \mathcal{S} \mapsto [0, 1]^C$ that learns to map a point cloud to a semantic label distribution. Throughout our work, we remain agnostic to the type of function h used for the mapping and we focus on data augmentation to generate new examples.

Data augmentation is an integral part of training deep neural networks, especially when the size of the training data is limited compared to the size of the model parameters. A popular data augmentation strategy is Mixup [34]. Mixup performs augmentation in the image domain by linearly interpolating pixels, as well as labels. Specifically, let $I_1 \in \mathbb{R}^{W \times H \times 3}$ and $I_2 \in \mathbb{R}^{W \times H \times 3}$ denote two images. Then a new image and its label are generated as:

$$I_{\text{mix}}(\lambda) = (1 - \lambda) \cdot I_1 + \lambda \cdot I_2, \quad (1)$$

$$c_{\text{mix}}(\lambda) = (1 - \lambda) \cdot c_1 + \lambda \cdot c_2, \quad (2)$$

where $\lambda \in [0, 1]$ denotes the mixup ratio. Usually λ is sampled from a beta distribution $\lambda \sim \text{Beta}(\gamma, \gamma)$. Such a direct interpolation is feasible for images as the data is aligned. In point clouds, however, linear interpolation is not straightforward. The reason is that point clouds are sets of points in which the point elements are orderless and permutation-invariant. We must, therefore, seek a definition of interpolation on unordered sets.

3.2 Interpolation between point clouds

Let $S_1 \in \mathcal{S}$ and $S_2 \in \mathcal{S}$ denote two training examples on which we seek to perform interpolation with ratio λ to generate new training examples. Given a pair of source examples S_1 and S_2 , an interpolation function, $f_{S_1 \rightarrow S_2} : [0, 1] \mapsto \mathcal{S}$ can be *any continuous function*, which forms a curve that joins S_1 and S_2 in a metric space (\mathcal{S}, d) with a proper distance function d . This means that it is up to us to define what makes an interpolation good. We define the concept of shortest-path interpolation in the context of point cloud:

Definition 1 (Shortest-path interpolation). *In a metric space (\mathcal{S}, d) , a shortest-path interpolation $f_{S_1 \rightarrow S_2}^* : [0, 1] \mapsto \mathcal{S}$ is an interpolation between the given pair of source examples $S_1 \in \mathcal{S}$ and $S_2 \in \mathcal{S}$, such that for any $\lambda \in [0, 1]$, $d(S_1, S^{(\lambda)}) + d(S^{(\lambda)}, S_2) = d(S_1, S_2)$ holds for $S^{(\lambda)} = f_{S_1 \rightarrow S_2}^*(\lambda)$ being the interpolant.*

We say that Definition 1 ensures the shortest path property because the triangle inequality holds for any properly defined distance $d : d(S_1, S^{(\lambda)}) + d(S^{(\lambda)}, S_2) \geq d(S_1, S_2)$. The intuition behind this definition is that the shortest path property ensures the uniqueness of the label distribution on the interpolated data.

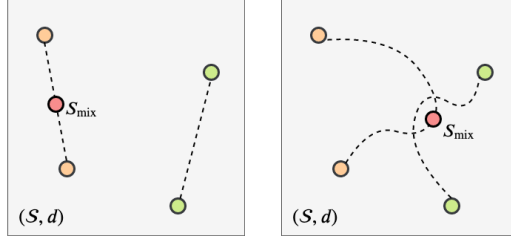


Fig. 2: **Intuition of shortest-path interpolation.** The examples lives on a metric space (\mathcal{S}, d) as dots in the figure. The dashed lines are the interpolation paths between different pairs of examples. When the shortest-path property is ensured (left), the interpolation paths from different pairs of source examples are likely to be not intersect in a complicated metric space. While in non-shortest path interpolation (right), the paths can intertwine with each other with a much higher probability, making it hard to tell which pair of source examples does the mixed data come from.

To put it otherwise, when computing interpolants from different sources, the interpolants generated by the shortest-path interpolation is more likely to be discriminative than the ones generated by a non-shortest-path interpolation.

To define an interpolation for point clouds, therefore, we must first select a reasonable distance metric. Then, we opt for the shorterst-path interpolation function based on the selected distance metric. For point clouds a proper distance metric is the Earth Mover’s Distance (EMD), as it captures well not only the geometry between two point clouds, but also local details as well as density distributions [5, 1, 13]. EMD measures the least amount of total displacement required for each of the points in the first point cloud, $x_i \in S_1$, to match a corresponding point in the second point cloud, $y_j \in S_2$. Formally, the EMD for point clouds solves the following assignment problem:

$$\phi^* = \arg \min_{\phi \in \Phi} \sum_i \|x_i - y_{\phi(i)}\|_2, \quad (3)$$

where $\Phi = \{\{1, \dots, N\} \mapsto \{1, \dots, N\}\}$ is the set of possible bijective assignments, which give one-to-one correspondences between points in the two point clouds. Given the optimal assignment ϕ^* , the EMD is then defined as the average effort to move S_1 points to S_2 :

$$d_{\text{EMD}} = \frac{1}{N} \sum_i \|x_i - y_{\phi^*(i)}\|_2. \quad (4)$$

3.3 PointMixup: Optimal assignment interpolation for point clouds

We propose an interpolation strategy, which can be used for augmentation that is analogous of Mixup [34] but for point clouds. We refer to this proposed PointMixup as *Optimal Assignment (OA) Interpolation*, as it relies on the optimal assignment on the basis of the EMD to define the interpolation between clouds.

Given the source pair of point clouds $S_1 = \{x_i\}_{i=1}^N$ and $S_2 = \{y_j\}_{j=1}^N$, the Optimal Assignment (OA) interpolation is a path function $f_{S_1 \rightarrow S_2}^* : [0, 1] \mapsto \mathcal{S}$. With $\lambda \in [0, 1]$,

$$f_{S_1 \rightarrow S_2}^*(\lambda) = \{u_i\}_{i=1}^N, \quad \text{where} \quad (5)$$

$$u_i = (1 - \lambda) \cdot x_i + \lambda \cdot y_{\phi^*(i)}, \quad (6)$$

in which ϕ^* is the optimal assignment from S_1 to S_2 defined by Eq. 3. Then the interpolant $S_{\mathbf{OA}}^{S_1 \rightarrow S_2, (\lambda)}$ (or $S_{\mathbf{OA}}^{(\lambda)}$ when there is no confusion) generated by the OA interpolation path function $f_{S_1 \rightarrow S_2}^*(\lambda)$ is the required augmented data for point cloud Mixup.

$$S_{\mathbf{OA}}^{(\lambda)} = \{(1 - \lambda) \cdot x_i + \lambda \cdot y_{\phi^*(i)}\}_{i=1}^N. \quad (7)$$

Under the view of $f_{S_1 \rightarrow S_2}^*$ being a path function in the metric space $(\mathcal{S}, d_{\text{EMD}})$, f is expected to be the shortest path joining S_1 and S_2 since the definition of the interpolation is induced from the EMD.

3.4 Analysis

Intuitively we expect that PointMixup is a shortest path linear interpolation. That is, the interpolation lies on the shortest path joining the source pairs, and the interpolation is linear with regard to λ in $(\mathcal{S}, d_{\text{EMD}})$, since the definition of the interpolation is derived from the EMD. However, it is non-trivial to show the optimal assignment interpolation abides to a shortest path linear interpolation, because the optimal assignment between the mixed point cloud and either of the source point cloud is unknown. It is, therefore, not obvious that we can ensure whether there exists a shorter path between the mixed examples and the source examples. To this end, we need to provide an in-depth analysis.

To ensure the uniqueness of the label distribution from the mixed data, we need to show that the *shortest path property* w.r.t. the EMD is fulfilled. Moreover, we need to show that the proposed interpolation is *linear* w.r.t the EMD, in order to ensure that the input interpolation has the same ratio as the label interpolation. Besides, we evaluate the *assignment invariance property* as a prerequisite knowledge for the proof for the linearity. This property implies that there exists no shorter path between interpolants with different λ , *i.e.*, the shortest path between the interpolants is a part of the shortest path between the source examples. Due to space limitation, we sketch the proof for each property. The complete proofs are available in the supplementary material.

We start with the shortest path property. Since the EMD for point cloud is a metric, the triangle inequality $d_{\text{EMD}}(A, B) + d_{\text{EMD}}(B, C) \geq d_{\text{EMD}}(A, C)$ holds (for which a formal proof can be found in [19]). Thus we formalize the shortest path property into the following proposition:

Property 1 (shortest path) *Given the source examples S_1 and S_2 , $\forall \lambda \in [0, 1]$, $d_{\text{EMD}}(S_1, S_{\mathbf{OA}}^{(\lambda)}) + d_{\text{EMD}}(S_{\mathbf{OA}}^{(\lambda)}, S_2) = d_{\text{EMD}}(S_1, S_2)$.*

Sketch of Proof From the definition of the EMD we can derive $d_{\text{EMD}}(S_1, S_{\text{OA}}^{(\lambda)}) + d_{\text{EMD}}(S_2, S_{\text{OA}}^{(\lambda)}) \leq d_{\text{EMD}}(S_1, S_2)$. Then from the triangle inequity of the EMD, only the equality remains. \square

We then introduce the assignment invariance property of the OA Mixup as an intermediate step for the proof of the linearity of OA Mixup. The property shows that the assignment does not change for any pairs of the mixed point clouds with different λ . Moreover, the assignment invariance property is important to imply that the shortest path between the any two mixed point clouds is part of the shortest path between the two source point clouds.

Property 2 (assignment invariance) $S_{\text{OA}}^{(\lambda_1)}$ and $S_{\text{OA}}^{(\lambda_2)}$ are two mixed point clouds from the same given source pair of examples S_1 and S_2 as well as the mix ratios λ_1 and λ_2 such that $0 \leq \lambda_1 < \lambda_2 \leq 1$. Let the points in $S_{\text{OA}}^{(\lambda_1)}$ and $S_{\text{OA}}^{(\lambda_2)}$ be $u_i = (1 - \lambda_1) \cdot x_i + \lambda_1 \cdot y_{\phi^*(i)}$ and $v_k = (1 - \lambda_2) \cdot x_k + \lambda_2 \cdot y_{\phi^*(k)}$, where ϕ^* is the optimal assignment from S_1 to S_2 . Then the identical assignment ϕ_I is the optimal assignment from $S_{\text{OA}}^{(\lambda_1)}$ to $S_{\text{OA}}^{(\lambda_2)}$.

Sketch of Proof We first prove that the identical mapping is the optimal assignment from S_1 to $S_{\text{OA}}^{(\lambda_1)}$ from the definition of the EMD. Then we prove that ϕ^* is the optimal assignment from $S_{\text{OA}}^{(\lambda_1)}$ to S_2 . Finally we prove that the identical mapping is the optimal assignment from $S_{\text{OA}}^{(\lambda_1)}$ to $S_{\text{OA}}^{(\lambda_2)}$ similarly as the proof for the first intermediate argument. \square

Given the property of assignment invariance, the linearity follows:

Property 3 (linearity) For any mix ratios λ_1 and λ_2 such that $0 \leq \lambda_1 < \lambda_2 \leq 1$, the mixed point clouds $S_{\text{OA}}^{(\lambda_1)}$ and $S_{\text{OA}}^{(\lambda_2)}$ satisfies that $d_{\text{EMD}}(S_{\text{OA}}^{(\lambda_1)}, S_{\text{OA}}^{(\lambda_2)}) = (\lambda_2 - \lambda_1) \cdot d_{\text{EMD}}(S_1, S_2)$.

Sketch of Proof The proof can be directly derived from the fact that the identical mapping is the optimal assignment between $S_{\text{OA}}^{(\lambda_1)}$ and $S_{\text{OA}}^{(\lambda_2)}$. \square

The linear property of our interpolation is important, as we jointly interpolate the point clouds and the labels. By ensuring that the point cloud interpolation is linear, we ensure that the input interpolation has the same ratio as the label interpolation.

On the basis of the properties, we find that PointMixup is a shortest path linear interpolation between point clouds in $(\mathcal{S}, d_{\text{EMD}})$.

3.5 Manifold PointMixup: Interpolate between latent point features

In standard PointMixup, only the inputs, *i.e.*, the XYZ point cloud coordinates are mixed. The input XYZs are low-level geometry information and sensitive to disturbances and transformations, which in turn limits the robustness of the

PointMixup. Inspired by Manifold Mixup [26], we can also use the proposed interpolation solution to mix the latent representations in the hidden layers of point cloud networks, which are trained to capture salient and high-level information that is less sensitive to transformations. PointMixup can be applied for the purpose of Manifold Mixup to mix both at the XYZs and different levels of latent point cloud features and maintain their respective advantages, which is expected to be a stronger regularizer for improved performance and robustness.

We describe how to mix the latent representations. Following [26], at each batch we randomly select a layer l to perform PointMixup from a set of layers L , which includes the input layer. In a point cloud network, the intermediate latent representation at layer l (before the global aggregation stage such as the max pooling aggregation in PointNet [15] and PointNet++ [16]) is $Z_{(l)} = \{(x_i, z_i^{(x)})\}_{i=1}^{N_z}$, in which x_i is 3D point coordinate and $z_i^{(x)}$ is the corresponding high-dimensional feature. For the mixed latent representation, given the latent representation of two source examples are $Z_{(l),1} = \{(x_i, z_i^{(x)})\}_{i=1}^{N_z}$ and $Z_{(l),2} = \{(y_i, z_i^{(y)})\}_{i=1}^{N_z}$, the optimal assignment ϕ^* is obtained by the 3D point coordinates x_i , and the mixed latent representation then becomes

$$\begin{aligned} Z_{(l),\mathbf{OA}}^{(\lambda)} &= \{(x_i^{\text{mix}}, z_i^{\text{mix}})\}, \quad \text{where} \\ x_i^{\text{mix}} &= (1 - \lambda) \cdot x_i + \lambda \cdot y_{\phi^*(i)}, \\ z_i^{\text{mix}} &= (1 - \lambda) \cdot z_i^{(x)} + \lambda \cdot z_{\phi^*(i)}^{(y)}. \end{aligned}$$

Specifically in PointNet++, three layers of representations are randomly selected to perform Manifold Mixup: the input, and the representations after the first and the second SA modules (See appendix of [16]).

4 Experiments

4.1 Setup

Datasets. We focus in our experiments on the **ModelNet40** dataset [31]. This dataset contains 12,311 CAD models from 40 man-made object categories, split into 9,843 for training and 2,468 for testing. We furthermore perform experiments on the **ScanObjectNN** dataset [25]. This dataset consists of real-world point cloud objects, rather than sampled virtual point clouds. The dataset consists of 2,902 objects and 15 categories. We report on two variants of the dataset, a standard variant OBJ_ONLY and one with heavy permutations from rigid transformations PB.T50_RS [25].

Following [12], we discriminate between settings where each dataset is pre-aligned and unaligned with horizontal rotation on training and test point cloud examples. For the unaligned settings, we randomly rotate the training point cloud along the up-axis. Then, before solving the optimal assignment, we perform a simple additional alignment step to fit and align the symmetry axes between the two point clouds to be mixed. Through this way, the point clouds are better




Fig. 3: **Baseline interpolation variants.** Top: point cloud interpolation through random assignment. Bottom: interpolation through sampling.

aligned and we obtain more reasonable point correspondences. Last, we also perform experiments using only 20% of the training data.

Network architectures. The main network architecture used throughout the paper is PointNet++ [16]. We also report results with PointNet [15] and DGCNN [29], to show that our approach is agnostic to the architecture that is employed. PointNet learns a permutation-invariant set function, which does not capture local structures induced by the metric space the points live in. PointNet++ is a hierarchical structure, which segments a point cloud into smaller clusters and applies PointNet locally. DGCNN performs hierarchical operations by selecting a local neighbor in the feature space instead of the point space, resulting in each point having different neighborhoods in different layers.

Experimental details. We uniformly sample 1,024 points on the mesh faces according to the face area and normalize them to be contained in a unit sphere, which is a standard setting [15, 16, 12]. In case of mixing clouds with different number of points, we can simply replicate random elements from the each point set to reach the same cardinality. During training, we augment the point clouds on-the-fly with random jitter for each point using Gaussian noise with zero mean and 0.02 standard deviation. We implement our approach in PyTorch [14]. For network optimization, we use the Adam optimizer with an initial learning rate of 10^{-3} . The model is trained for 300 epochs with a batch size of 16. We follow previous work [34, 26] and draw λ from a beta distribution $\lambda \sim \text{Beta}(\gamma, \gamma)$. We also perform Manifold Mixup [26] in our approach, through interpolation on the transformed and pooled points in intermediate network layers. In this work, we opt to use the efficient algorithm and adapt the open-source implementation from [13] to solve the optimal assignment approximation. Training for 300 epochs takes around 17 hours without augmentation and around 19 hours with PointMixup or Manifold PointMixup on a single NVIDIA GTX 1080 ti.

Baseline interpolations. For our comparisons to baseline point cloud augmentations, we compare to two variants. The first variant is random assignment

interpolation, where a random assignment $\phi^{\mathbf{RA}}$ is used, to connect points from both sets, yielding:

$$S_{\mathbf{RA}}^{(\lambda)} = \{(1 - \lambda) \cdot x_i + \lambda \cdot y_{\phi^{\mathbf{RA}}(i)}\}.$$

The second variant is point sampling interpolation, where random draws without replacement of points from each set are made according to the sampling frequency λ :

$$S_{\mathbf{PS}}^{(\lambda)} = S_1^{(1-\lambda)} \cup S_2^{(\lambda)},$$

where $S_2^{(\lambda)}$ denotes a randomly sampled subset of S_2 , with $\lfloor \lambda N \rfloor$ elements. ($\lfloor \cdot \rfloor$ is the floor function.) And similar for S_1 with $N - \lfloor \lambda N \rfloor$ elements, such that $S_{\mathbf{PS}}^{(\lambda)}$ contains exactly N points. The intuition of the point sampling variant is that for point clouds as unordered sets, one can move one point cloud to another through a set operation such that it removes several random elements from set S_1 and replace them with same amount of elements from S_2 .

4.2 Point cloud classification ablations

We perform four ablation studies to show the workings of our approach with respect to the interpolation ratio, comparison to baseline interpolations and other regularizations, as well robustness to noise.

Effect of interpolation ratio. The first ablation study focuses on the effect of the interpolation ratio in the data augmentation for point cloud classification. We perform this study on ModelNet40 using the PointNet++ architecture. The results are shown in Fig. 4 for the pre-aligned setting. We find that regardless of the interpolation ratio used, our approach provides a boost over the setting without augmentation by interpolation. PointMixup positively influences point cloud classification. The inclusion of manifold mixup adds a further boost to the scores.

Throughout further experiments, we use $\gamma = 0.4$ for input mixup and $\gamma = 1.5$ for manifold mixup in unaligned setting, and $\gamma = 1.0$ for input mixup and $\gamma = 2.0$ for manifold mixup in pre-aligned setting.

Comparison to baseline interpolations. In the second ablation study, we investigate the effectiveness of our PointMixup compared to the two interpolation baselines. We again use ModelNet40 and PointNet++. We perform the evaluation on both the pre-aligned and unaligned dataset variants, where for both we also report results with a reduced training set. The results are shown in Table 1. Across both the alignment variants and dataset sizes, our PointMixup obtains favorable results. This result highlights the effectiveness of our

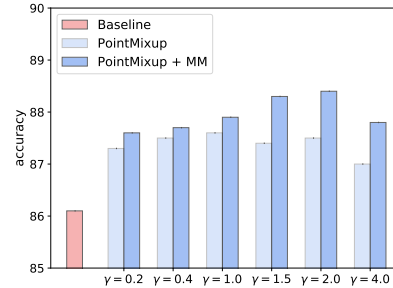


Fig. 4: **Effect of interpolation ratios.** MM denotes Manifold Mixup.

Table 1: **Comparison of PointMixup to baseline interpolations** on ModelNet40 using PointNet++. PointMixup compares favorably to excluding interpolation and to the baselines, highlighting the benefits of our shortest path interpolation solution.

	No mixup	Random assignment	Point sampling	PointMixup
Manifold mixup	×	×	✓	×
Full dataset				
Pre-aligned	91.9	91.6	91.9	92.2
Unaligned	90.7	90.8	91.1	90.9
Reduced dataset				
Pre-aligned	86.1	85.5	87.3	87.2
Unaligned	84.4	84.8	85.4	85.7

Table 2: **Evaluating our approach to other data augmentations** (left) **and its robustness to noise and transformations** (right). We find that our approach with manifold mixup (MM) outperforms augmentations such as label smoothing and other variations of mixup. For the robustness evaluation, we find that our approach with strong regularization power from manifold mixup provides more robustness to random noise and geometric transformations.

	PointMixup		Transforms	PointMixup		
	×	✓		×	w/o MM	w/ MM
Baseline with no mixing	86.1	–	Noise $\sigma^2 = 0.01$	91.3	91.9	92.3
Mixup	–	87.6	Noise $\sigma^2 = 0.05$	35.1	51.5	56.5
Manifold mixup	–	88.6	Noise $\sigma^2 = 0.1$	4.03	4.27	7.42
Mix input, not labels	–	86.6	Z-rotation [-30,30]	74.3	70.9	77.8
Mix input from same class	–	86.4	X-rotation [-30,30]	73.2	70.8	76.8
Mixup latent (layer 1)	–	86.9	Y-rotation [-30,30]	87.6	87.9	88.7
Mixup latent (layer 2)	–	86.8	Scale (0.6)	85.8	84.5	86.3
Label smoothing (0.1)	87.2	–	Scale (2.0)	59.2	67.7	72.9
Label smoothing (0.2)	87.3	–	DropPoint (0.2)	84.9	78.1	90.9

approach, which abides to the shortest path linear interpolation definition, while the baselines do not.

PointMixup with other regularizers. Third, we evaluate how well PointMixup works by comparing to multiple existing data regularizers and mixup variants, again on ModelNet40 and PointNet++. We investigate the following augmentations: (i) Mixup [34], (ii) Manifold Mixup [26], (iii) mix input only without target mixup, (iv) mix latent representation at a fixed layer (manifold mixup does so at random layers), and (v) label smoothing [22]. Training is performed on the reduced dataset to better highlight their differences. We show the results in Table 2 on the left. Our approach with manifold mixup obtains the highest scores. The label smoothing regularizer is outperformed, while we also obtain better scores than the mixup variants. We conclude that PointMixup is forms an effective data augmentation for point clouds.

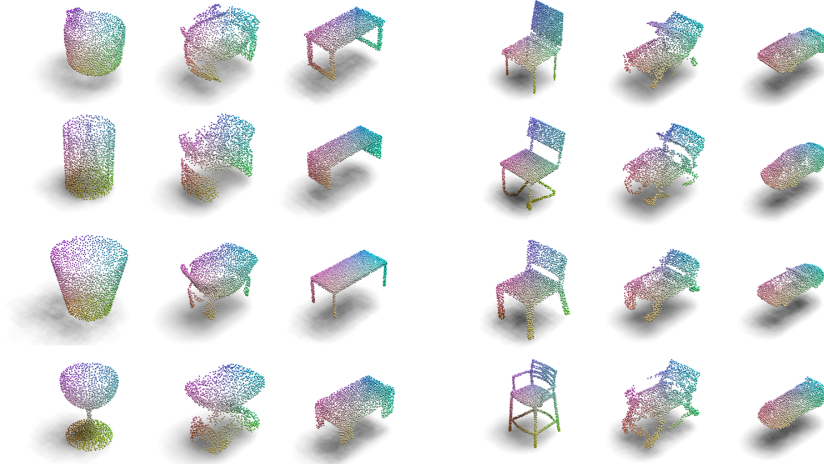


Fig. 5: **Qualitative examples of PointMixup.** We provide eight visualizations of our interpolation. The four examples on the left show interpolations for different configurations of cups and tables. The four examples on the right show interpolations for different chairs and cars.

Robustness to noise. By adding additional augmented training examples, we enrich the dataset. This enrichment comes with additional robustness with respect to noise in the point clouds. We evaluate the robustness by adding random noise perturbations on point location, scale, translation and different rotations. Note that for evaluation of robustness against up-axis rotation, we use the models which are trained with the pre-aligned setting, in order to test also the performance against rotation along the up-axis as a novel transform. The results are in Table 2 on the right. Overall, our approach including manifold mixup provides more stability across all perturbations. For example, with additional noise ($\sigma = 0.05$), we obtain an accuracy of 56.5, compared to 35.1 for the baseline. We similar trends for scaling (with a factor of two), with an accuracy of 72.9 versus 59.2. We conclude that PointMixup makes point cloud networks such as PointNet++ more stable to noise and rigid transformations.

Qualitative analysis. In Figure 5, we show eight examples of PointMix for point cloud interpolation; four interpolations of cups and tables, four interpolations of chairs and cars. Through our shortest path interpolation, we end up at new training examples that exhibit characteristics of both classes, making for sensible point clouds and mixed labels, which in turn indicate why PointMixup is beneficial for point cloud classification.

4.3 Evaluation on other networks and datasets

With PointMixup, new point clouds are generated by interpolating existing point clouds. As such, we are agnostic to the type of network or dataset. To highlight

Table 3: **PointMixup on other networks** (left) **and another dataset** (right). We find our approach is beneficial regardless the network or dataset.

	PointNet		DGCNN			ScanObjectNN			
	×	w/o MM	×	w/o MM	w/ MM	×	w/o MM	w/ MM	
Full	89.2	89.9	92.7	92.9	93.1	Standard	86.6	87.6	88.5
Reduced	81.3	83.4	88.2	88.8	89.0	Perturbed	79.3	80.2	80.6

this ability, we perform additional experiments on extra networks and an additional point cloud dataset.

PointMixup on other network architectures. We show the effect of PointMixup to two other networks, namely PointNet [15] and DGCNN [29]. The experiments are performed on ModelNet40. For PointNet, we perform the evaluation on the unaligned setting and for DGCNN with pre-aligned setting to remain consistent with the alignment choices made in the respective papers. The results are shown in Table 3 on the left. We find improvements when including PointMixup for both network architectures.

PointMixup on real-world point clouds. We also investigate PointMixup on point clouds from real-world object scans, using ScanObjectNN [25], which collects object from 3D scenes in SceneNN [9] and ScanNet [4]. Here, we rely on PointNet++ as network. The results in Table 3 on the right show that we can adequately deal with real-world point cloud scans, hence we are not restricted to point clouds from virtual scans. This result is in line with experiments on point cloud perturbations.

4.4 Beyond standard classification

The fewer training examples available, the stronger the need for additional examples through augmentation. Hence, we train PointNet++ on ModelNet40 in both a few-shot and semi-supervised setting.

Semi-supervised learning. Semi-supervised learning learns from a dataset where only a small portion of data is labeled. Here, we show how PointMixup directly enables semi-supervised learning for point clouds. We start from Interpolation Consistency Training [27], a state-of-the-art semi-supervised approach, which utilizes Mixup between unlabeled points. Here, we use our Mixup for point clouds within their semi-supervised approach. We evaluate on ModelNet40 using 400, 600, and 800 labeled point clouds. The result of semi-supervised learning are illustrated in Table 4 on the left. Compared to the supervised baseline, which only uses the available labelled examples, our mixup enables the use of additional unlabelled training examples, resulting in a clear boost in scores. With 800 labelled examples, the accuracy increases from 73.5% to 82.0%, highlighting the effectiveness of PointMixup in a semi-supervised setting.

Few-shot learning. Few-shot classification aims to learn a classifier to recognize unseen classes during training with limited examples. We follow [28, 18, 20, 6, 21] to regard few-shot learning a typical meta-learning method, which

Table 4: **Evaluating PointMixup in the context of semi-supervised (left) and few-shot learning (right).** When examples are scarce, as is the case for both settings, using our approach provides a boost to the scores.

	Semi-supervised classification			Few-shot classification	
	Supervised	[27]+PointMixup		[20]	+ PointMixup
400 examples	69.4	76.7	5-way 1-shot	72.3	77.2
600 examples	72.6	80.8	5-way 3-shot	80.2	82.2
800 examples	73.5	82.0	5-way 5-shot	84.2	85.9

learns how to learn from limited labeled data through training from a collection of tasks, *i.e.*, episodes. In an N -way K -shot setting, in each task, N classes are selected and K examples for each class are given as a support set, and the query set consists of the examples to be predicted. We perform few-shot classification on ModelNet40, from which we select 20 classes for training, 10 for validation, and 10 for testing. We utilize PointMixup within ProtoNet [20] by constructing mixed examples from the support set and update the model with the mixed examples before making predictions on the query set. We refer to the supplementary material for the details of our method and the settings. The results in Table 4 on the right show that incorporating our data augmentation provides a boost in scores, especially in the one-shot setting, where the accuracy increases from 72.3% to 77.2%.

5 Conclusion

This work proposes PointMixup for data augmentation on point clouds. Given the lack of data augmentation by interpolation on point clouds, we start by defining it as a shortest path linear interpolation. We show how to obtain PointMixup between two point clouds by means of an optimal assignment interpolation between their point sets. As such, we arrive at a Mixup for point clouds, or latent point cloud representations in the sense of Manifold Mixup, that can handle permutation invariant nature. We first prove that PointMixup abides to our shortest path linear interpolation definition. Then, we show through various experiments that PointMixup matters for point cloud classification. We show that our approach outperforms baseline interpolations and regularizers. Moreover, we highlight increased robustness to noise and geometric transformations, as well as its general applicability to point-based networks and datasets. Lastly, we show the potential of our approach in both semi-supervised and few-shot settings. The generic nature of PointMixup allows for a comprehensive embedding in point cloud classification.

Acknowledgment This research was supported in part by the SAVI/MediFor and the NWO VENI What & Where projects.

References

1. Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L.: Learning representations and generative models for 3d point clouds. In: ICML (2018)
2. Berthelot, D., Carlini, N., Cubuk, E.D., Kurakin, A., Sohn, K., Zhang, H., Raffel, C.: Remixmatch: Semi-supervised learning with distribution alignment and augmentation anchoring. In: CoRR (2019)
3. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. In: NeurIPS (2019)
4. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: CVPR (2017)
5. Fan, H., Su, H., Guibas, L.J.: A point set generation network for 3d object reconstruction from a single image. In: CVPR (2017)
6. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: ICML (2017)
7. French, G., Aila, T., Laine, S., Mackiewicz, M., Finlayson, G.: Consistency regularization and cutmix for semi-supervised semantic segmentation. In: CoRR (2019)
8. Guo, H., Mao, Y., Zhang, R.: Mixup as locally linear out-of-manifold regularization. In: AAAI (2019)
9. Hua, B.S., Pham, Q.H., Nguyen, D.T., Tran, M.K., Yu, L.F., Yeung, S.K.: Scenenn: A scene meshes dataset with annotations. In: 3DV (2016)
10. Li, J., Chen, B.M., Hee Lee, G.: So-net: Self-organizing network for point cloud analysis. In: CVPR. pp. 9397–9406 (2018)
11. Li, R., Li, X., Heng, P.A., Fu, C.W.: Pointaugment: an auto-augmentation framework for point cloud classification. In: CVPR. pp. 6378–6387 (2020)
12. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. In: NeurIPS (2018)
13. Liu, M., Sheng, L., Yang, S., Shao, J., Hu, S.M.: Morphing and sampling network for dense point cloud completion. In: CoRR (2019)
14. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: NeurIPS (2019)
15. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: CVPR (2017)
16. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017)
17. Ravanbakhsh, S., Schneider, J., Poczos, B.: Deep learning with sets and point clouds. arXiv preprint arXiv:1611.04500 (2016)
18. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: ICLR (2016)
19. Rubner, Y., Tomasi, C., Guibas, L.J.: The earth mover’s distance as a metric for image retrieval. IJCV (2000)
20. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: NeurIPS (2017)
21. Sung, F., Yang, Y., Zhang, L., Xiang, T., Torr, P.H., Hospedales, T.M.: Learning to compare: Relation network for few-shot learning. In: CVPR (2018)
22. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR (2016)

23. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: ICCV. pp. 6411–6420 (2019)
24. Tokozume, Y., Ushiku, Y., Harada, T.: Between-class learning for image classification. In: CVPR (2018)
25. Uy, M.A., Pham, Q.H., Hua, B.S., Nguyen, T., Yeung, S.K.: Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In: ICCV (2019)
26. Verma, V., Lamb, A., Beckham, C., Najafi, A., Mitliagkas, I., Courville, A., Lopez-Paz, D., Bengio, Y.: Manifold mixup: Better representations by interpolating hidden states. In: ICML (2019)
27. Verma, V., Lamb, A., Kannala, J., Bengio, Y., Lopez-Paz, D.: Interpolation consistency training for semi-supervised learning. IJCAI (2019)
28. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: NeurIPS (2016)
29. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. TOG (2019)
30. Wu, W., Qi, Z., Fuxin, L.: Pointconv: Deep convolutional networks on 3d point clouds. In: CVPR (2019)
31. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: CVPR (2015)
32. Xu, Y., Fan, T., Xu, M., Zeng, L., Qiao, Y.: Spidercnn: Deep learning on point sets with parameterized convolutional filters. In: ECCV (2018)
33. Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R.R., Smola, A.J.: Deep sets. In: NeurIPS. pp. 3391–3401 (2017)
34. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. In: ICLR (2017)
35. Zhang, Z., He, T., Zhang, H., Zhang, Z., Xie, J., Li, M.: Bag of freebies for training object detection neural networks. In: CoRR (2019)
36. Zhang, Z., Hua, B.S., Yeung, S.K.: Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics. In: ICCV (2019)