





Point Cloud Generation with Continuous Conditioning

Larissa T. Triess^{1,2}  Andre Bühler^{1,3} David Peter¹ 
Fabian B. Flohr^{1,3}  J. Marius Zöllner^{2,4} 

¹Mercedes-Benz AG ²Karlsruhe Institute of Technology
³University of Stuttgart ⁴Research Center for Information Technology

Abstract

Generative models can be used to synthesize 3D objects of high quality and diversity. However, there is typically no control over the properties of the generated object. This paper proposes a novel generative adversarial network (GAN) setup that generates 3D point cloud shapes conditioned on a continuous parameter. In an exemplary application, we use this to guide the generative process to create a 3D object with a custom-fit shape. We formulate this generation process in a multi-task setting by using the concept of auxiliary classifier GANs. Further, we propose to sample the generator label input for training from a kernel density estimation (KDE) of the dataset. Our ablations show that this leads to significant performance increase in regions with few samples. Extensive quantitative and qualitative experiments show that we gain explicit control over the object dimensions while maintaining good generation quality and diversity.

1 INTRODUCTION

In recent years many approaches evolved to analyze point clouds, such as classification or segmentation (Guo et al., 2021; Hackel et al., 2016; Li et al., 2018; Qi et al., 2017a; Thomas et al., 2019; Wang et al., 2019; Wu et al., 2019). Point clouds are a popular representation for 3D data acquired with depth sensing and laser scanning and are used in many applications, such as automated driving or human-robot interaction.

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

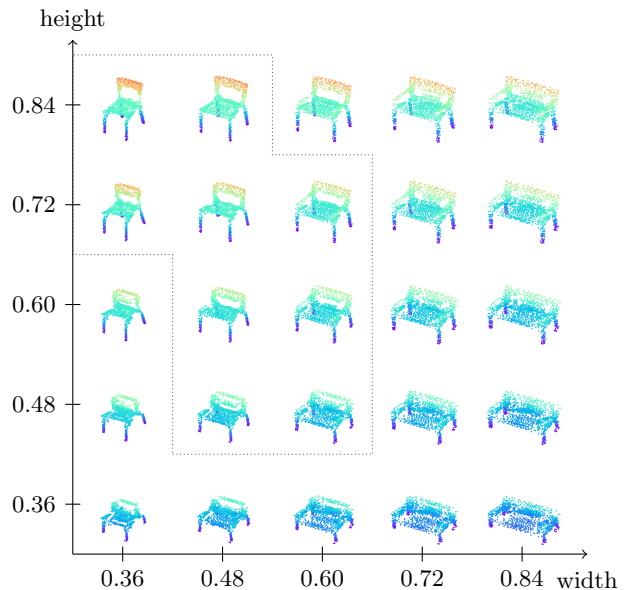


Figure 1: **Generated objects conditioned on different dimensions:** Our method can generate a diverse set of object shapes and be conditioned on object dimensions (values on the axes). The figure shows objects generated from the same latent vector \mathbf{z} , but with different continuous conditioning parameters \mathbf{y} . The generated objects are realistic and semantically meaningful. Our method generalizes to out-of-distribution dimensions (outside of dotted shape).

Here, often complex 3D scenes have to be analyzed by detecting objects, segmenting the scene, or estimating motion (Zhou and Tuzel, 2018; Lang et al., 2019; Liu et al., 2019; Milioto et al., 2019; Sirohi et al., 2021). The scenes are composed of a large number of different objects with hierarchical dependencies. Therefore, most 3D object detection approaches on real-world data use augmentation techniques to improve detection performance. This is often achieved by placing additional objects into the scene when training the networks (Yan et al., 2018; Lang et al., 2019; Baur et al.,

2019). This can greatly improve performance, especially for underrepresented classes, but is limited to the size and diversity of the provided object database. Therefore, it is favorable to have an unlimited number of objects at ones disposal.

Generative models, such as GANs (Goodfellow et al., 2014) or variational autoencoders (VAEs) (Kingma and Welling, 2014), are often used to generate completely new samples with high quality and diversity. These approaches are initially introduced for image generation, but lately a number of approaches for 3D generation have emerged (Achlioptas et al., 2018; Valsesia et al., 2019; Yang et al., 2019; Shu et al., 2019; Sun et al., 2020; Luo and Hu, 2021). However, none of these methods is capable to actively influence specific properties of the generated object, such as height and width (see Fig. 1). Therefore, we propose a method to use descriptions in form of continuous conditional parameters within a GAN to generate objects with desired properties. These properties can for example describe the aspect-ratios of the object, such that the GAN generates a custom-fit shape while maintaining generation quality and diversity. Using continuous conditions introduces new challenges, since it is a mathematically different problem than solving categorical conditioning problems, such as classification (Ding et al., 2021). The challenges arise from an infinitely large set of parameters and conditioning regions where no training data samples exist.

We formulate the continuous conditioning as a multi-task problem within the discriminator. We add an additional output head, similar to AcGAN (Odena et al., 2017), that estimates the continuous parameter. The overall loss function is then formulated as the weighted sum of the adversarial loss and the regression loss. Additionally, we propose an alternative label sampling that is used at training time. In a specifically designed experiment setup, we show that this sampling strategy improves the generation quality of the generator and opens the potential to generate novel out-of-distribution shapes. To the best of our knowledge, this work is first to use continuous conditional labels for 3D generation.

2 RELATED WORK

2.1 3D Generative Models

In recent years, a variety of methods emerged from GANs and VAEs to synthesize realistic data. Initially proposed to generate realistic 2D images, these concepts have widely been adapted to fulfill a variety of tasks on different modalities (Song et al., 2018; Chen et al., 2019; Karras et al., 2019). Wu et al. (2016) are

first to propose an unsupervised deep generative approach for 3D data generation from probabilistic input. Their voxel-based GAN allows to directly adapt concepts from 2D, but is limited in resolution due to computational inefficiency. Therefore, subsequent works focus on surface (Mescheder et al., 2019; Chen and Zhang, 2019; Michalkiewicz et al., 2019; Park et al., 2019) or point cloud (Qi et al., 2017a,b; Fan et al., 2017) representations instead. Achlioptas et al. (2018) propose architectures for both VAEs and GANs to generate point cloud objects. The use of relative simple models based on fully connected layers combined with PointNet layers limits the ability to produce more realistic objects. The method by Valsesia et al. (2019) exploits local topology by using a computationally heavy k -nearest neighbor technique to produce geometrically accurate point clouds. Shu et al. (2019) aim to improve the expressiveness of the generative models by introducing convolution-like layers and up-sampling techniques to the generator. Wang et al. (2020) deal with the choice of suitable discriminator architectures for 3D generation and show that models perform better when focusing on overall object shapes as well as sampling quality.

2.2 Conditional Generation

In many cases additional conditioning parameters are desired to generate specific object categories or styles. The most prominent example, conditional GAN (Mirza and Osindero, 2014), uses explicit conditioning and forms the basis of many other approaches (Zhu et al., 2017; Taigman et al., 2017; Hoffman et al., 2018). AcGAN (Odena et al., 2017) refines this concept of class conditioning by using an additional auxiliary classifier in the discriminator to ensure class specific content. To condition the output on arbitrary combinations of discrete attributes, some works use labeled images as input to the GAN (He et al., 2019; Perarnau et al., 2016). StyleGAN and others investigate how to enhance desirable properties in GAN latent spaces to influence the characteristics of generated images selectively (Karras et al., 2019, 2020; Härkönen et al., 2020).

2.3 Continuous Conditioning

Many of the aforementioned concepts use discrete label conditioning. However, attributes like rotation in angles or age in years are by definition continuous. Using continuous conditions is a mathematically different problem than solving categorical conditioning problems, such as classification (Ding et al., 2021). First, there may be few or no real samples for some regression labels and second, conventional label input methods, i.e. one-hot encoding, is not possible for an

infinite number of regression labels. CcGAN (Ding et al., 2021) is first to introduce a continuous conditional GAN for image generation. They solve the aforementioned problems by introducing a new GAN loss and a novel way to input the labels based on label projection (Miyato and Koyama, 2018). Shoshan et al. (2021) propose to use attribute specific pre-trained classifiers to enhance desired properties on the generative behavior. A subsequent training of mapping networks allows to generate noise vectors which produce explicit continuous attributes. Their method produces convincing results but requires an extensive amount of labeling and well pre-trained classifiers for each attribute.

2.4 3D Conditional Generation

The aforementioned conditioning strategies have all been proposed for image synthesis. There are some works that use text-based conditioning to generate 3D scenes (Chang et al., 2015a, 2014; Chen et al., 2018) with focus on database composition. Other approaches use symbolic part-tree descriptions to generate 3D objects with predefined compositions (Mo et al., 2020) or use occupancy networks for image based generation and coloring of 3D objects (Mescheder et al., 2019). Although being related, our method focuses on conditioning point cloud generation using continuous physical parameters.

2.5 Contributions

Our main contributions are threefold:

- We propose a GAN setup which formulates the generation process conditioned on continuous parameters in a multi-task setting by adapting the concept of auxiliary classifier GANs (Odena et al., 2017).
- We propose to sample the generator conditioning input for training from a KDE of the parameter distribution. Our ablations show that this leads to a significant performance increase in regions with few samples.
- We provide a number of qualitative evaluations that show that we gain explicit control over the object dimensions while maintaining generation quality and diversity.

3 BACKBONE

This section gives a short introduction to our backbone network. We build upon TreeGAN (Shu et al., 2019), a state-of-the-art GAN architecture that can

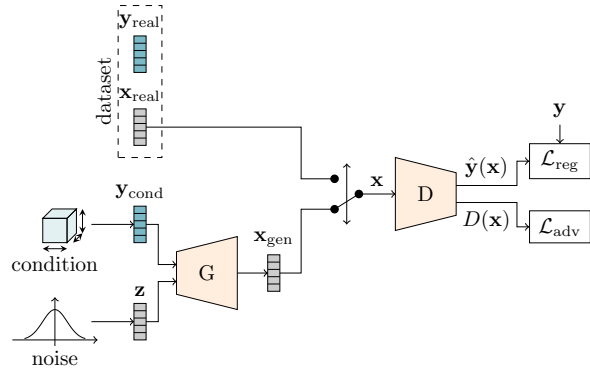


Figure 2: **Architecture:** The generator G generates a point cloud \mathbf{x}_{gen} from a random vector \mathbf{z} and a continuous parameter \mathbf{y}_{cond} . It is then alternately – with a real point cloud \mathbf{x}_{real} – fed to the discriminator, which predicts the probability of the sample stemming from the real distribution $D(\mathbf{x})$ and an estimate of the parameter $\hat{\mathbf{y}}(\mathbf{x})$. With these two outputs, the adversarial \mathcal{L}_{adv} and regression \mathcal{L}_{reg} losses are computed.

generate point cloud objects with high quality and diversity. The generator consists of stacked tree graph convolution layers (TreeGCN). It receives a random noise vector $\mathbf{z} \in \mathbb{R}^{96}$ as input and outputs a point cloud $\mathbf{x}_{\text{gen}} = G(\mathbf{z}) \in \mathbb{R}^{2048 \times 3}$. The loss function of the generator G is defined as

$$\mathcal{L}_{G,\text{adv}} = -\mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(\mathbf{x}_{\text{gen}})], \quad (1)$$

where D denotes the discriminator. The latent code distribution \mathcal{Z} is sampled from a Normal distribution $\mathbf{z} \in \mathcal{N}(\mathbf{0}, I)$.

The discriminator follows a PointNet architecture (Qi et al., 2017a). It either receives a real \mathbf{x}_{real} or a generated \mathbf{x}_{gen} point cloud $\in \mathbb{R}^{2048 \times 3}$ as input and outputs a single scalar $D(\mathbf{x})$. The output estimates whether the sample originates from the distribution of the real or generated samples. The loss function of the discriminator D is defined as

$$\begin{aligned} \mathcal{L}_{D,\text{adv}} = & \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(\mathbf{x}_{\text{gen}})] - \mathbb{E}_{\mathbf{x} \sim \mathcal{R}} [D(\mathbf{x}_{\text{real}})] \\ & + \lambda_{\text{gp}} \cdot \mathbb{E}_{\hat{\mathbf{x}}} \left[(\|\Delta_{\hat{\mathbf{x}}} D(\hat{\mathbf{x}})\|_2 - 1)^2 \right], \end{aligned} \quad (2)$$

where $\hat{\mathbf{x}}$ is sampled from an interpolation between real and fake point clouds and λ_{gp} is the weighting parameter for the gradient penalty term (Gulrajani et al., 2017).

4 METHOD

Fig. 2 presents the GAN setup of our proposed continuous conditioning architecture. The concept for the continuous conditioning is introduced in Section 4.1

while Section 4.2 explains the required parameter sampling Section 4.3 describes the model architecture and the training losses are introduced in Section 4.4.

4.1 Continuous Parameters

We aim to control the outer dimensions of the generated object with the continuous parameters while maintaining diversity regarding the type and shape of the object. While the object size and shape shall be disentangled, our method is still allowed to modify the shape slightly to better match the defined bounding box in terms of creating realistic and valid shapes, i.e. a small table can have a single leg, while a long table typically does not. We later refer to this property as *semantically meaningful*. The object is defined as a vector of points $\mathbf{x} \in \mathbb{R}^{N \times 3}$ with the number of points $N = 2048$ and the dimensions $[x, y, z]$. The parameter $\mathbf{y} \in \mathbb{R}^3$ defines the extent of the object in each dimension $\mathbf{y} = (\Delta x, \Delta y, \Delta z) \in [0, 1]^3$. Therefore, the parameters \mathbf{y}_{real} for the training data \mathbf{x}_{real} can easily be computed from the data itself with $\mathbf{y}_{\text{real}} = \|\max(\mathbf{x}_{\text{real}}) - \min(\mathbf{x}_{\text{real}})\|$.

4.2 Label Sampling for Training

At training time, we have to sample parameters \mathbf{y}_{cond} as a second input to the generator G . The easiest method is to sample randomly in $[0, 1]$. However, this can lead to a description of unsuitable object dimensions. To circumvent this issue, we can sample the dimensions from the training dataset. This ensures that only actually possible conditioning parameters are sampled. However, this limits the generator training to a fixed number of input conditioning. Therefore, we propose to sample the dimensions from the distribution of the dimensions within the training dataset. We compute a KDE over the whole dataset prior to training. This again assures that only suitable dimensions are drawn, but it does not limit them to be present in the dataset. In all cases, the aim of the generator is to generate point clouds that have dimensions which are close to the conditioned dimensions \mathbf{y}_{cond} .

4.3 Model

This section only describes the changes made to our backbone architecture for which we use TreeGAN, as described in Section 3. Fig. 3 shows how the label conditioning is incorporated into the generator and discriminator models. For the generator (Fig. 3a), we extend the architecture to receive an additional input, the conditioning vector \mathbf{y}_{cond} . This vector and the noise vector $\mathbf{z} \in \mathbb{R}^{96}$ are both passed through a linear layer. The concatenated results are then fed to the otherwise unmodified TreeGCN (Shu et al., 2019)




Figure 3: **Model details:** (a) shows the label input of the generator. Both inputs are fed through a fully-connected (FC) layer, concatenated and then fed to the tree graph convolution network (TreeGCN). (b) shows the discriminator. After a common feature extractor (PointNet), the model splits in two identical parts, where $D(\mathbf{x})$ is the adversarial feedback and $\hat{\mathbf{y}}(\mathbf{x})$ is the prediction of the continuous object description.

which outputs the point cloud $\mathbf{x}_{\text{gen}} = G(\mathbf{y}_{\text{cond}}, \mathbf{z})$. The discriminator input \mathbf{x} is either a point cloud from the dataset \mathbf{x}_{real} or a generated point cloud from the generator \mathbf{x}_{gen} (Fig. 3b). The network has two outputs, one for the standard adversarial feedback $D(\mathbf{x})$ and one to estimate the continuous parameter conditioning of the presented data $\hat{\mathbf{y}} = \hat{\mathbf{y}}(\mathbf{x})$. This concept is adapted from Auxiliary Classifier GAN (ACGAN) (Odena et al., 2017; Atienza, 2019). The idea is to leverage potential synergies between the two tasks within the shared discriminator layers.

4.4 Losses

The training objective is composed of two parts, the adversarial loss \mathcal{L}_{adv} and the regression loss \mathcal{L}_{reg} for the continuous parameter. Just like TreeGAN, we utilize the Wasserstein objective function with gradient penalty for the adversarial loss (Arjovsky et al., 2017; Gulrajani et al., 2017). For the parameter regression, we use L2-norm

$$\mathcal{L}_{\text{reg}}(\mathbf{y}, \hat{\mathbf{y}}) = \|\mathbf{y} - \hat{\mathbf{y}}\|_2 \quad (3)$$

with the predicted label $\hat{\mathbf{y}}$ and its target label \mathbf{y} .

The adversarial part of the generator loss $\mathcal{L}_{G, \text{adv}}$ is defined in Eq. (1) with a slight change to the definition of the generated point cloud \mathbf{x}_{gen} . It is now also dependent on the input conditioning and is therefore defined as $G(\mathbf{y}_{\text{cond}}, \mathbf{z})$ (instead of the unconditioned version $G(\mathbf{z})$). For the generator regression, we compute the \mathcal{L}_{reg} between the discriminator prediction for the condition parameter of the generated sample $\hat{\mathbf{y}}_{\text{gen}}$ and the actually requested parameter \mathbf{y}_{cond} . This results in the overall generator loss

$$\mathcal{L}_G = \lambda_{\text{adv}} \cdot \mathcal{L}_{G, \text{adv}} + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \hat{\mathbf{y}}_{\text{gen}}) \quad (4)$$

with the loss weighting factors λ_{adv} and λ_{reg} .

Eq. (2) defines the adversarial loss of the discriminator. The regression loss is defined as the error between the parameter of the real data \mathbf{y}_{real} and its corresponding prediction $\hat{\mathbf{y}}_{\text{real}}$. Analogously to the generator, this leads to the discriminator loss

$$\mathcal{L}_D = \lambda_{\text{adv}} \cdot \mathcal{L}_{D,\text{adv}} + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{real}}, \hat{\mathbf{y}}_{\text{real}}) \quad (5)$$

with the same loss weights λ_{adv} and λ_{reg} as in Eq. (4).

5 EXPERIMENTS

This section gives an overview on the experiments with the results reported in Section 6. Further details are provided in the supplementary material.

5.1 Dataset and Metrics

For our experiments, we use ShapeNetPart (Yi et al., 2016), a dataset with part annotations of more than 30,000 3D shapes in 16 object categories from ShapeNetCore (Chang et al., 2015b). To compare our results in terms of generation quality, we use a pre-trained version of the original *Fréchet Point Cloud Distance* (FPD) (Shu et al., 2019). The adherence of the conditioning properties are evaluated by calculating the mean squared error (MSE) for each dimension extent (Δx , Δy , Δz) of the generated object \mathbf{x}_{gen} versus the desired input parametrization \mathbf{y}_{cond} . These are the two most important metrics, but we also report coverage (COV), minimum matching distance (MMD), and Jensen-Shannon Divergence (JSD) to make our work comparable to existing methods (Achlioptas et al., 2018; Shu et al., 2019). The evaluation also focuses on qualitative results to show the advantages of our method.

5.2 Implementation Details

We build our model upon the existing implementation of TreeGAN (Shu et al., 2019) that we further refer to as our backbone model. The training parameters are also identical. We only apply the changes introduced in Section 4 to keep our method directly comparable to the backbone and do not incorporate any further mechanisms to enhance generation quality.

The loss weights λ_{adv} and λ_{reg} from Eq. (4) and Eq. (5) are variable and learned together with the rest of the model parameters, as proposed by Kendall et al. (2018).

We train all networks for 3000 epochs and select the checkpoint with the lowest combined metric score. The combined metric is defined as the product of FPD and MSE. This ensures fidelity as well as correctness and considers differing value ranges.

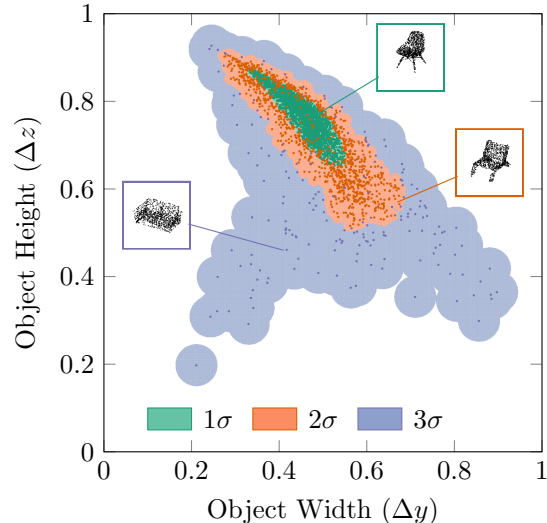


Figure 4: **Region-classified dimension distribution:** The plot shows all samples (marks) of the chair class for the entire training dataset in terms of their dimension extent in height z and width y . The three colors represent the resulting regions from k -nearest neighbor classifier with $k=20$ based on a KDE. The regions correspond to $1\sigma \approx 68\%$ (green), $2\sigma \approx 27\%$ (orange), and $3\sigma = 5\%$ (blue) of the entire data distribution.

5.3 Baselines

We compare our method to two baselines. These baselines represent possible design choices to achieve the desired aim of generating point clouds of specific dimensions when no explicit method already exists.

B1 – Backbone with re-sampling: A simple method to generate shapes of desired dimensions with an existing model, such as the backbone, is to generate point clouds from multiple sampled z vectors and then choose the one that has the dimensions that are closest to the ones requested. Here it is expected that the fidelity of the generated samples is similar to the ones from the backbone network itself. However, the sampling process is time consuming and the dimensions might not exactly fit, especially if these dimensions are not well reflected in the training dataset. For our experiments, we sample 10 versions per object and select the one with the smallest MSE.

B2 – Backbone with scaling: Another variant is to add a subsequent scaling block that squashes the generated point cloud into the desired dimensions. In contrast to B1, it is expected that the regression error on the dimensions is zero. However, the generation fidelity can be severely impacted.

5.4 Distribution Sampling

A key aspect of our method is that it is possible to actively sample from different regions of the conditioning vector distribution. Therefore, we specifically investigate the generation capabilities within different regions of the data distribution. Fig. 4 shows the distribution of the object dimensions of the training dataset for the chair class. The distribution is divided into three sections, indicated by the coloring. The sections are classified into regions where $1\sigma \approx 68\%$, $2\sigma \approx 27\%$, and $3\sigma \approx 5\%$, of the data lies, computed with a KDE. In additional experiments, we generate 1000 samples for each region and then compare FPD and MSE to show the effectiveness of our method.

6 RESULTS

6.1 Quantitative Results

Table 5 contains quantitative comparisons in terms of typical metrics used to evaluate the quality of generated point cloud objects. The rows for the backbone are included for reference but cannot be used for comparison, as it does not have the conditioning ability. B1 and B2 each symbolize the two corner cases. On the one hand, B1 obtains the lowest FPD but the highest MSE. This is the result of exploiting the good generation quality of the backbone and combine it with a sampling mechanism to obtain object of desired dimensions. However, it becomes clear that it is hard to obtain the desired shape configuration in an acceptable inference time, while the inference time depends on how many samples have to be drawn for one inference step. On the other hand, there is B2 which obtains an MSE of zero by construction. It simply scales the object to the desired size. However, this comes at the cost of realism, as evident from the increase in FPD. Our proposed approach lies in between B1 and B2 regarding FPD and obtains a very low MSE of only 0.28%. The other metrics show relatively equal performance for all methods.

The numbers indicate that our method is capable to ensure desired dimensions while maintaining high quality. However, we want to stress that the full potential of our method is better visible in the following qualitative results.

It is important to note that these baselines are only applicable since it is possible to easily compute the dimension parameter from the data itself. For many other applications this is not the case and our method offers the suitable solution (see also Section 7).

Additionally to the presented results, we also evaluated the label incorporation of traditional

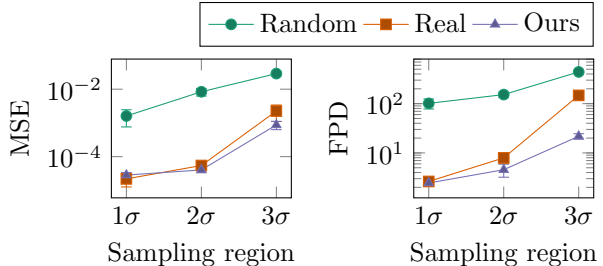


Figure 5: **Distribution sampling performance:** The figure shows performances for three data sampling regions for our model trained with three different label sampling strategies. The regions correspond to the labels on the x-axis, as defined in Fig. 4. For each region, 1000 samples are generated. The left plot shows the MSE of the dimension regression. On the right, the FPD of the generated samples is presented.

cGAN (Mirza and Osindero, 2014) and CcGAN (Ding et al., 2021) combined with the backbone network. However, these methods did not achieve satisfying results in our setup. In particular, cGAN completely ignored the conditioning and therefore received MSE values over 50% (at FPD ≈ 1.9 for “chair”). While for CcGAN, training was unstable and therefore resulted in metric scores that are several magnitudes higher than usual. For more information and additional results, we refer to the supplemental material.

6.2 Label and Region Sampling Ablations

Fig. 5 shows the results of our distribution sampling experiments on the chair class (explained in Section 5.4). We report results for three differently trained versions of our model. They differ in the way the conditioning parameter \mathbf{y}_{cond} is sampled at training time (see Section 4.2). The green dots correspond to random sampling in $[0, 1]$, while orange means that only parameters existing in the training dataset are being used, i.e. the marks of Fig. 4. Our proposed version (purple) uses all possible dimensions sampled from a KDE of the training distribution, i.e. the entire region in Fig. 4.

It can be observed that naturally MSE and FPD increase when moving away from the distribution center of gravity (increasing σ). For random sampling, we see significantly worse performance in all three sampling categories compared to the other two methods. For 1σ and 2σ the two other sampling methods perform almost equally well. However, we see a significant performance improvement of our proposed sampling strategy in the most sparsely populated region where only 5% of all training data lies. This shows the advantage of the distribution-based training over a data-based

Table 1: **Quantitative Comparison:** We report results for the classes “Chair” and “Airplane” in terms of the metrics used by Shu et al. (2019). Additionally, we report the regression error (MSE) for our introduced task. We report the results for the backbone network (TreeGAN) for reference. All evaluations are conducted on a hold-out validation split. Shu et al. (2019) use the entire dataset for training, therefore values vary slightly. Best results are marked in *bold*, second best in *italic*.

Shape	Model	FPD (\downarrow)	MSE [%] (\downarrow)	MMD (\downarrow)		COV (\uparrow)		JSD (\downarrow)
				CD	EMD	CD	EMD	
Chair	Backbone	0.9525	–	0.0020	0.1027	0.4875	0.2500	0.1082
	Baseline 1 (B1)	1.3674	26.07	0.0023	0.1013	0.4750	0.2625	0.1123
	Baseline 2 (B2)	1.9259	0.00	0.0021	0.1003	0.4875	0.2625	0.1068
	Ours	<i>1.5290</i>	<i>0.28</i>	0.0022	0.1059	0.4625	0.3125	0.1434
Airplane	Backbone	1.2947	–	0.0002	0.0805	0.4375	0.1375	0.1887
	Baseline 1 (B1)	<i>1.0209</i>	15.08	0.0003	0.0812	0.4500	0.1125	0.1819
	Baseline 2 (B2)	1.6613	0.00	0.0003	0.0783	0.5250	0.1375	0.1834
	Ours	0.8691	<i>0.30</i>	0.0003	0.0724	0.5000	0.1250	0.1291




Figure 6: **Manipulation of chair height:** Each row of objects is generated from a different latent vector. The object height is increased from left to right. The upper two rows show chairs generated with our method, while the bottom row shows the backbone network with scaling to the requested height (B2).

training, especially for increasing σ .

6.3 Continuous Parameter Interpolation

A decisive advantage of our method is that we can actively and directly influence the object dimensions. Fig. 6 and Fig. 7 show several examples of generated objects where the continuous parameter for the object size is changed. For different latent vectors, our networks generates different shapes of good quality and diversity. If now object sizes are changed, quality declines for the outermost samples (on the edge of the distribution). However, our method still produces semantically meaningful objects in contrast to the baseline B2. For example, the overall shape of the office chair in the upper row of Fig. 6 stays approximately the same while the backrest gets a slight tilt to the back to resemble an easy chair for lower heights, while B2 preserves the shape completely and simply compresses the object. We want to stress that these figures




Figure 7: **Manipulation of table width:** The object width is increased from left to right while keeping length and height constant. The upper three rows show tables generated with our method, the bottom row is generated with baseline B2. While B2 simply stretches the table, our method generates *semantically meaningful* object shapes, e.g. in second row, the wide tables have more than one leg.

are for demonstration purposes, the main application is rather to define the object size and then sample shapes (compare Fig. 8) and not to sample a shape and then modify the size, as shown here.

6.4 Out-of-Distribution Generation

Fig. 1 shows a larger span of size conditioning, where only the samples enclosed in the dotted shape are conditioned on parameters sampled from within the distribution. The dataset does not contain any chairs with sizes that are larger or smaller than the gray dots indicate (refer to Fig. 4). The generated shapes for extreme low or high height or width do not necessarily resemble realistic objects, but the generator still maintains the approximate object and shape configuration.




Figure 8: **Generation Diversity:** The upper row shows examples generated from diverse latent vectors at a fixed dimension \mathbf{y} . The lower row shows the $k=5$ nearest neighbors from the training dataset in terms of being closest to \mathbf{y} . This shows that our method generates a variety of shapes for the same dimension and does not perform a simple lookup from the dataset.

All generated objects show smooth transitions between the configurations, even when out-of-distribution and do not collapse in shape.

6.5 Model Properties

Fig. 8 shows five generated examples and five examples from the dataset that have the same dimensions. The generated shapes are diverse and considerably distinct from the dataset samples. This shows two things: First, our method does not simply generate the same object for a given dimension when changing the latent vector (*diversity*). Second, it also does not learn a simple lookup by reproducing the samples from the dataset that lie close to the desired dimensions (*novelty*).

As in related work, we demonstrate smooth representations within our model by interpolating between two latent vectors. Fig. 9 shows that we obtain smooth transitions from left to right for different object shapes. This property was already included in the backbone network, but this experiment shows that it was not compromised by the introduction of our conditioning method.

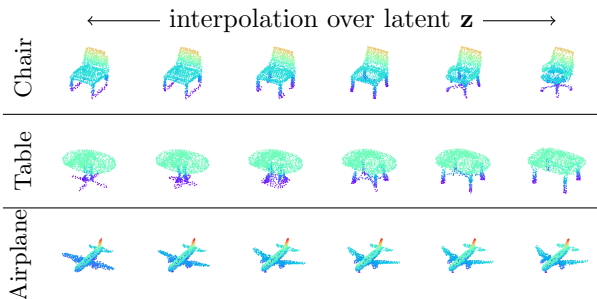


Figure 9: **Latent interpolation:** Interpolation between two random latent vectors \mathbf{z} with constant \mathbf{y} .




Figure 10: **Ratio of object parts:** The left column shows two part-annotated examples from the dataset with airplane body in red, wings in blue, and the remaining parts in yellow. In the upper row, our method increases the ratio of wing points within the object, while the lower row increases the percentage of the body part from left to right. Our method adapts the overall shape to be *semantically meaningful* and match the continuous parameter for a fixed latent vector.

7 DISCUSSION

Influencing object dimensions is only one application out of many for explicit continuous conditioning. Fig. 10 demonstrates how our method can be used to influence the percentage of certain object parts relative to the entire shape. To this end, the former dimension conditioning is now defined as a ratio of the body parts $\mathbf{y} \in \mathbb{R}^2$ with $\mathbf{y} = (\text{wings} : \text{body}, \text{others} : \text{body}) \in [0, 1]^2$ for the class “Airplane” where $\sum \mathbf{y} = 1.0$. The labels for the real data are obtained by using the point-wise part segmentation information from the dataset. We count the number of points for each part and divide it with the number of points for the part “body” to obtain a ratio of points for each part. Modifying the relation of object parts is more entangled with the overall appearance, therefore the shape changes considerably. In contrast to the object dimensions it is not possible to trivially retrieve this parameter from the generated shape. For the object dimensions a simple calculation suffices, while for the part manipulation a reliable part segmentation model would be needed, which requires annotations for training and can introduce additional errors. Therefore, a naive regression loss for the conditioning vector or using a simple conditioning discriminator as in cGAN is not applicable. This makes the discriminator regression vital.

Next steps for this work include the adaptation to real-world data. Objects scanned with a laser are often only seen from a certain viewpoint in contrast to the objects in this work. The viewpoint can be modeled by using additional continuous parameters that define the angle and distance from which the object is observed. Regarding our proposed application for augmentation of autonomous driving LiDAR scenes, we can then define the size and position of a 3D bounding box for which the GAN then generates a custom-fit object. A challenge is to adapt the backbone architecture to

generate variable numbers of points, as common for real-world data.

8 CONCLUSION

This paper presented a novel GAN setup for 3D shape generation that uses continuous conditional parameters to actively influence the dimensions of generated shapes. In extensive experiments we showed that our method can generate custom-fit objects that adhere the desired configuration while maintaining good generation quality and diversity. We showed that our distribution label sampling is superior to sampling existing dataset parameters. Further, we demonstrated generalization to out-of-distribution generation and gave a preview on potential other applications. Future work includes the adaptation to real-world data.

References

- Achlioptas, P., Diamanti, O., Mitliagkas, I., and Guibas, L. (2018). Learning Representations and Generative Models for 3D Point Clouds. In *Proc. of the International Conf. on Machine learning (ICML)*.
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein GAN. *arXiv.org*.
- Atienza, R. (2019). A Conditional Generative Adversarial Network for Rendering Point Clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Baur, S. A., Moosmann, F., Wirges, S., and Rist, C. B. (2019). Real-time 3D LiDAR Flow for Autonomous Vehicles. In *Proc. IEEE Intelligent Vehicles Symposium (IV)*, pages 1288–1295.
- Chang, A., Monroe, W., Savva, M., Potts, C., and Manning, C. D. (2015a). Text to 3D Scene Generation with Rich Lexical Grounding. In *Proc. of the International Joint Conf. on Natural Language Processing (IJCNLP)*.
- Chang, A., Savva, M., and Manning, C. D. (2014). Learning Spatial Knowledge for Text to 3D Scene Generation. In *Proc. of the Conf. on Empirical Methods in Natural Language Processing (EMNLP)*.
- Chang, A. X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., and Yu, F. (2015b). ShapeNet: An Information-Rich 3D Model Repository. Technical report, Stanford University — Princeton University — Toyota Technological Institute at Chicago.
- Chen, K., Choy, C. B., Savva, M., Chang, A. X., Funkhouser, T., and Savarese, S. (2018). Text2Shape: Generating Shapes from Natural Language by Learning Joint Embeddings. In *Proc. of the Asian Conf. on Computer Vision (ACCV)*.
- Chen, L., Wu, L., Hu, Z., and Wang, M. (2019). Quality-Aware Unpaired Image-to-Image Translation. In *IEEE Trans. on Multimedia (TMM)*.
- Chen, Z. and Zhang, H. (2019). Learning Implicit Fields for Generative Shape Modeling. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Ding, X., Wang, Y., Xu, Z., Welch, W. J., and Wang, Z. J. (2021). CcGAN: Continuous Conditional Generative Adversarial Networks for Image Generation. In *Proc. of the International Conf. on Learning Representations (ICLR)*.
- Fan, H., Su, H., and Guibas, L. (2017). A Point Set Generation Network for 3D Object Reconstruction from a Single Image. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NIPS)*.
- Guirajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. (2017). Improved Training of Wasserstein GANs. In *Advances in Neural Information Processing Systems (NIPS)*.
- Guo, M.-H., Cai, J.-X., Liu, Z.-N., Mu, T.-J., Martin, R. R., and Hu, S.-M. (2021). PCT: Point cloud transformer. In *Computational Visual Media*.
- Hackel, T., Wegner, J. D., and Schindler, K. (2016). Contour Detection in Unstructured 3D Point Clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1610–1618.
- Härkönen, E., Hertzmann, A., Lehtinen, J., and Paris, S. (2020). GANSpace: Discovering Interpretable GAN Controls. In *Advances in Neural Information Processing Systems (NIPS)*.
- He, Z., Zuo, W., Kan, M., Shan, S., and Chen, X. (2019). AttGAN: Facial Attribute Editing by Only Changing What You Want. In *IEEE Trans. on Image Processing (TIP)*.
- Hoffman, J., Tzeng, E., Park, T., Zhu, J.-Y., Isola, P., Saenko, K., Efros, A. A., and Darrell, T. (2018). CyCADA: Cycle Consistent Adversarial Domain Adaptation. In *Proc. of the International Conf. on Machine learning (ICML)*.
- Karras, T., Laine, S., and Aila, T. (2019). A Style-Based Generator Architecture for Generative Adversarial Networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.

- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. (2020). Analyzing and Improving the Image Quality of StyleGAN. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Kendall, A., Gal, Y., and Cipolla, R. (2018). Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Kingma, D. P. and Welling, M. (2014). Auto-Encoding Variational Bayes. In *Proc. of the International Conf. on Learning Representations (ICLR)*.
- Lang, A. H., Vora, S., Caesar, H., Zhou, L., Yang, J., and Beijbom, O. (2019). PointPillars: Fast Encoders for Object Detection from Point Clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 12697–12705.
- Li, Y., Bu, R., Sun, M., Wu, W., Di, X., and Chen, B. (2018). PointCNN: Convolution On X-Transformed Points. In *Advances in Neural Information Processing Systems (NIPS)*.
- Liu, X., Qi, C. R., and Guibas, L. J. (2019). FlowNet3D: Learning Scene Flow in 3D Point Clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Luo, S. and Hu, W. (2021). Diffusion Probabilistic Models for 3D Point Cloud Generation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., and Geiger, A. (2019). Occupancy Networks: Learning 3D Reconstruction in Function Space. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Michalkiewicz, M., Pontes, J. K., Jack, D., Baktashmotlagh, M., and Eriksson, A. (2019). Deep Level Sets: Implicit Surface Representations for 3D Shape Inference. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*.
- Milioto, A., Vizzo, I., Behley, J., and Stachniss, C. (2019). RangeNet++: Fast and Accurate LiDAR Semantic Segmentation. In *Proc. IEEE International Conf. on Intelligent Robots and Systems (IROS)*.
- Mirza, M. and Osindero, S. (2014). Conditional Generative Adversarial Nets. *arXiv.org*.
- Miyato, T. and Koyama, M. (2018). cGANs with Projection Discriminator. In *Proc. of the International Conf. on Learning Representations (ICLR)*.
- Mo, K., Wang, H., Yan, X., and Guibas, L. (2020). PT2PC: Learning to Generate 3D Point Cloud Shapes from Part Tree Conditions. In *Proc. of the European Conf. on Computer Vision (ECCV)*.
- Odena, A., Olah, C., and Shlens, J. (2017). Conditional Image Synthesis with Auxiliary Classifier GANs. In *Proceedings of the International Conference on Machine Learning*.
- Park, J. J., Florence, P., Straub, J., Newcombe, R., and Lovegrove, S. (2019). DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Perarnau, G., van de Weijer, J., Raducanu, B., and Álvarez, J. M. (2016). Invertible Conditional GANs for image editing. In *Advances in Neural Information Processing Systems (NIPS) Workshops*.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. (2017a). PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Qi, C. R., Yi, L., Su, H., and Guibas, L. J. (2017b). PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *Advances in Neural Information Processing Systems (NIPS)*.
- Shoshan, A., Bhonker, N., Kviatkovsky, I., and Medioni, G. (2021). GAN-Control: Explicitly Controllable GANs. *arXiv.org*.
- Shu, D. W., Park, S. W., and Kwon, J. (2019). 3D Point Cloud Generative Adversarial Network Based on Tree Structured Graph Convolutions. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*.
- Sirohi, K., Mohan, R., Büscher, D., Burgard, W., and Valada, A. (2021). EfficientLPS: Efficient LiDAR Panoptic Segmentation. *arXiv.org*.
- Song, Y., Yang, C., Lin, Z., Liu, X., Huang, Q., Li, H., and Kuo, C.-C. J. (2018). Contextual-Based Image Inpainting: Infer, Match, and Translate. In *Proc. of the European Conf. on Computer Vision (ECCV)*.
- Sun, Y., Wang, Y., Liu, Z., Siegel, J. E., and Sarma, S. E. (2020). PointGrow: Autoregressively Learned Point Cloud Generation with Self-Attention. In *Proc. of the IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- Taigman, Y., Polyak, A., and Wolf, L. (2017). Unsupervised Cross-Domain Image Generation. In *Proc. of the International Conf. on Learning Representations (ICLR)*.
- Thomas, H., Qi, C. R., Deschaud, J.-E., Marcotegui, B., Goulette, F., and Guibas, L. J. (2019). KP-Conv: Flexible and Deformable Convolution for Point Clouds. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*.

- Valsesia, D., Fracastoro, G., and Magli, E. (2019). Learning Localized Generative Models for 3D Point Clouds via Graph Convolution. In *Proc. of the International Conf. on Learning Representations (ICLR)*.
- Wang, H., Jiang, Z., Yi, L., Mo, K., Su, H., and Guibas, L. J. (2020). Rethinking Sampling in 3D Point Cloud Generative Adversarial Networks. *arXiv.org*.
- Wang, Y., Sun, Y., Liu, Z., Sarma, S. E., Bronstein, M. M., and Solomon, J. M. (2019). Dynamic Graph CNN for Learning on Point Clouds. In *ACM Trans. on Graphics*.
- Wu, J., Zhang, C., Xue, T., Freeman, B., and Tenenbaum, J. (2016). Learning a Probabilistic Latent Space of Object Shapes via 3D Generative-Adversarial Modeling. In *Advances in Neural Information Processing Systems (NIPS)*.
- Wu, W., Qi, Z., and Fuxin, L. (2019). PointConv: Deep Convolutional Networks on 3D Point Clouds. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Yan, Y., Mao, Y., and Li, B. (2018). SECOND: Sparsely Embedded Convolutional Detection. In *Sensors*.
- Yang, G., Huang, X., Hao, Z., Liu, M.-Y., Belongie, S., and Hariharan, B. (2019). PointFlow: 3D Point Cloud Generation with Continuous Normalizing Flows. *arXiv.org*.
- Yi, L., Kim, V. G., Ceylan, D., Shen, I.-C., Yan, M., Su, H., Lu, C., Huang, Q., Sheffer, A., and Guibas, L. (2016). A Scalable Active Framework for Region Annotation in 3D Shape Collections. In *ACM Trans. on Graphics*.
- Zhou, Y. and Tuzel, O. (2018). VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. (2017). Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*.

Supplementary Material: Point Cloud Generation with Continuous Conditioning

A OVERVIEW

This supplementary material covers details of the deep neural network (DNN) architectures, hyperparameters, evaluation, and additional results. Section B gives all the training details of our proposed approach. For the sake of completeness, Section C discusses further experiments that are indicated in the main paper. Section D presents further qualitative and quantitative results to complement the results section of the main paper.

B IMPLEMENTATION DETAILS

B.1 Architecture

Table 2 lists all layers, inputs, and operations of our DNN architecture for the generator model. We use the code¹ from the original *PyTorch* implementation of TreeGAN (Shu et al., 2019). Except for the input layers, our configuration is equal to the one of TreeGAN. Table 3 lists all layers, inputs, and operations of the discriminator DNN architecture. Here, the PointNet (Qi et al., 2017a) contained in the TreeGAN code was used as a basis. The split of the network for the auxiliary classifier mode is located directly after the PointNet feature extractor layers. It is followed by two identical sequences of linear operations, where the adversarial head outputs a vector of size 1, while the regression head outputs a vector of size d of the continuous conditioning parameter.

B.2 Training

Two Adam optimizers are used for optimization, one for the parameters of the generator and one for the discriminator. For both, the learning rate is set to $1e^{-4}$. Additionally, the two weighting factors λ_{adv} and λ_{reg} for the losses are optimized. In the main paper, we formulate the losses for the generator and discriminator as

$$\mathcal{L} = \lambda_{\text{adv}} \cdot \mathcal{L}_{\text{adv}} + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}}$$

with the adversarial loss \mathcal{L}_{adv} and the regression loss \mathcal{L}_{reg} . In order to avoid simply learning weighting factors of zero and to ensure stable training convergence at the same time, the loss is implemented as

$$\mathcal{L} = \mathcal{L}_{\text{adv}} \cdot e^{v_{\text{adv}}} + v_{\text{adv}} + \mathcal{L}_{\text{reg}} \cdot e^{v_{\text{reg}}} + v_{\text{reg}}$$

with v_{adv} and v_{reg} being the trainable variables. Both variables are initialized to $v = 0$ at the beginning of the training, such that both loss parts are equally weighted.

C ADDITIONAL ANALYSIS

C.1 Loss Variations

Both the generator and the discriminator loss consist of an adversarial part and a regression part. The generator regression loss computes the error between the requested parameter \mathbf{y}_{cond} and the corresponding discriminator prediction, while the discriminator regression is defined as the error between the parameter of the real data \mathbf{y}_{real} and its corresponding prediction $\hat{\mathbf{y}}_{\text{real}}$, such that

$$\begin{aligned} \mathcal{L}_{G,\text{reg}} &= \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \hat{\mathbf{y}}_{\text{gen}}) \quad \text{and} \\ \mathcal{L}_{D,\text{reg}} &= \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{real}}, \hat{\mathbf{y}}_{\text{real}}). \end{aligned}$$

¹TreeGAN code: <https://github.com/seowok/TreeGAN>

Table 2: **Generator Architecture:** Detailed network architecture and input format definition. The ID of each row is used to reference the output of the row. \uparrow indicates that the layer directly above is an input. d is the number of dimensions of the conditioning parameter. In case of the dimensions extent $d = 3$, while for the influence of object part percentage $d = 1$.

ID	Inputs	Operation	Output Shape
1	\mathbf{z}	Sample latent vector from $\mathbf{z} \sim \mathcal{Z} = \mathcal{N}(\mathbf{0}, I)$	[96]
2	\mathbf{y}_{cond}	Sample continuous parameter from $\mathbf{y}_{\text{cond}} \sim \text{KDE}(\mathbf{y}_{\text{real}})$	[d]
Label Handling			
3	1	Linear Layer	[64]
4	2	Linear Layer	[32]
5	3, 4	Concatenate	[1×96]
Tree Graph Convolution (TreeGC) Network			
6	\uparrow	Tree Graph Convolution + LeakyReLU	[1×256]
7	\uparrow	Branching	[2×256]
8	\uparrow	Tree Graph Convolution + LeakyReLU	[2×256]
9	\uparrow	Branching	[4×256]
10	\uparrow	Tree Graph Convolution + LeakyReLU	[4×256]
11	\uparrow	Branching	[8×256]
12	\uparrow	Tree Graph Convolution + LeakyReLU	[8×128]
13	\uparrow	Branching	[16×128]
14	\uparrow	Tree Graph Convolution + LeakyReLU	[16×128]
15	\uparrow	Branching	[32×128]
16	\uparrow	Tree Graph Convolution + LeakyReLU	[32×128]
17	\uparrow	Branching	[2048×128]
18	\uparrow	Tree Graph Convolution	[2048×3]

Table 3: **Discriminator Architecture:** Detailed network architecture and input format definition. The ID of each row is used to reference the output of the row. \uparrow indicates that the layer directly above is an input. d is the number of dimensions of the conditioning parameter. In case of the dimensions extend $d = 3$, while for the influence of object part percentage $d = 1$.

ID	Inputs	Operation	Output Shape	Description
1	point cloud \mathbf{x}	x, y, z	$[2048 \times 3]$	Input point cloud $\mathbf{x} = \mathbf{x}_{\text{real}}$ for real data and $\mathbf{x} = \mathbf{x}_{\text{gen}}$ for generated data.
PointNet Feature Extractor				
2	\uparrow	Conv1D + LeakyReLU	$[2048 \times 64]$	Kernel size 1×1 , stride 1
3	\uparrow	Conv1D + LeakyReLU	$[2048 \times 128]$	Kernel size 1×1 , stride 1
4	\uparrow	Conv1D + LeakyReLU	$[2048 \times 256]$	Kernel size 1×1 , stride 1
5	\uparrow	Conv1D + LeakyReLU	$[2048 \times 512]$	Kernel size 1×1 , stride 1
6	\uparrow	Conv1D + LeakyReLU	$[2048 \times 1024]$	Kernel size 1×1 , stride 1
7	\uparrow	MaxPool	$[1024]$	Global features
Adversarial Output Head				
8	\uparrow	Linear Layer	$[1024]$	
9	\uparrow	Linear Layer	$[512]$	
10	\uparrow	Linear Layer	$[512]$	
11	\uparrow	Linear Layer	$[1]$	Output vector $D(\mathbf{x})$
Regression Output Head				
12	7	Linear Layer	$[1024]$	
13	\uparrow	Linear Layer	$[512]$	
14	\uparrow	Linear Layer	$[512]$	
15	\uparrow	Linear Layer	$[d]$	Output vector $\hat{\mathbf{y}}(\mathbf{x})$

It is notable that $\mathcal{L}_{D,\text{reg}}$ is only computed for the real samples and not for the generated samples and that $\mathcal{L}_{G,\text{reg}}$ uses the discriminator prediction $\hat{\mathbf{y}}_{\text{gen}}$ instead of the configuration of the actually generated object \mathbf{y}_{gen} . These design choices can be attributed to the fact that in many cases the actual parameter \mathbf{y}_{gen} of the generated point cloud \mathbf{x}_{gen} is unknown, i.e. cannot be trivially retrieved from \mathbf{x}_{gen} . Since our application of influencing object dimensions offers the possibility to simply compute \mathbf{y}_{gen} from \mathbf{x}_{gen} with $\mathbf{y}_{\text{gen}} = \|\max(\mathbf{x}_{\text{gen}}) - \min(\mathbf{x}_{\text{gen}})\|$, we additionally investigated loss variations that exploit this property.

First, we define the generator and discriminator regression losses as

$$\begin{aligned}\mathcal{L}_{G,\text{reg}} &= \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \mathbf{y}_{\text{gen}}) \quad \text{and} \\ \mathcal{L}_{D,\text{reg}} &= \frac{1}{2} \left[\mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{real}}, \hat{\mathbf{y}}_{\text{real}}) + \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{gen}}, \hat{\mathbf{y}}_{\text{gen}}) \right],\end{aligned}$$

respectively. We found that this leads to slightly better training convergence, but no significant outperformance in the final model in terms of FPD or MSE. Therefore, we conclude that the proposed losses of the main paper are a good mechanism to train the model when \mathbf{y}_{gen} is unknown, as for most applications.

Second, we investigated another loss configuration. As for the discriminator, our proposed method skips the generated part of the loss entirely. However, it is also possible to formulate the losses as follows

$$\begin{aligned}\mathcal{L}_{G,\text{reg}} &= \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \hat{\mathbf{y}}_{\text{gen}}) \quad \text{and} \\ \mathcal{L}_{D,\text{reg}} &= \frac{1}{2} \left[\mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{real}}, \hat{\mathbf{y}}_{\text{real}}) + \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \hat{\mathbf{y}}_{\text{gen}}) \right],\end{aligned}$$

where \mathbf{y}_{gen} from above is replaced with $\hat{\mathbf{y}}_{\text{real}}$ in the generator and \mathbf{y}_{cond} in the discriminator. However, we found that this leads to unstable training and results in a significantly worse final model. We attribute this to false feedback for the discriminator, especially in the beginning of the training. At that time, the generator is not yet well enough trained to output shapes that are close to the requested parameters ($\mathbf{y}_{\text{gen}} \neq \hat{\mathbf{y}}_{\text{gen}}$). Therefore, we found that it is best to not include the generated path for the discriminator regression at all.

C.2 Other Conditioning Concepts

Additionally to our proposed method we also investigate other configurations for the continuous conditioning of point cloud generation. These method either did not yield promising results or are limited in applicability, therefore they are not included in the main paper. However, for the sake of completeness and reproducibility, we include all relevant implementation details here and state the overall results of our experiments.

C.2.1 cGAN with Continuous Parameters

We adapted the concepts of cGAN Fig. 11 and CcGAN (Ding et al., 2021) to work with TreeGAN (Shu et al., 2019) as the backbone network. In contrast to our proposed approach, we call this as an implicit conditioning scheme, since there is no extrinsic excitation that forces the model to use the conditioning input explicitly. The discriminator receives both the point cloud and the conditioning parameter as an input. The loss function of the generator G is defined as

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(\mathbf{y}_{\text{cond}}, G(\mathbf{y}_{\text{cond}}, \mathbf{z}))]$$

where \mathcal{Z} represents the latent code distribution which follows a Normal distribution, such that $z \in \mathcal{N}(0, 1)$. The loss function of the discriminator is defined as

$$\begin{aligned}\mathcal{L}_D &= \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(\mathbf{y}_{\text{cond}}, G(\mathbf{y}_{\text{cond}}, \mathbf{z}))] \\ &\quad - \mathbb{E}_{\mathbf{x} \sim \mathcal{R}} [D(\mathbf{y}_{\text{real}}, \mathbf{x})] + \mathcal{L}_{\text{gp}}\end{aligned}$$

with the gradient penalty \mathcal{L}_{gp} as defined in the main paper.

For the first variant, we use the traditional label incorporation introduced by cGAN (Mirza and Osindero, 2014). We refer to this as the Vanilla cGAN. The generator details are depicted in Fig. 12a and the discriminator details are shown in Fig. 12b. The generator label incorporation simply concatenates the latent and the label vector prior to feeding it to the graph convolution network (GCN). In the discriminator, the point cloud is processed by PointNet which outputs a feature vector to which the label vector is concatenated before being processed by a final set of linear layers.

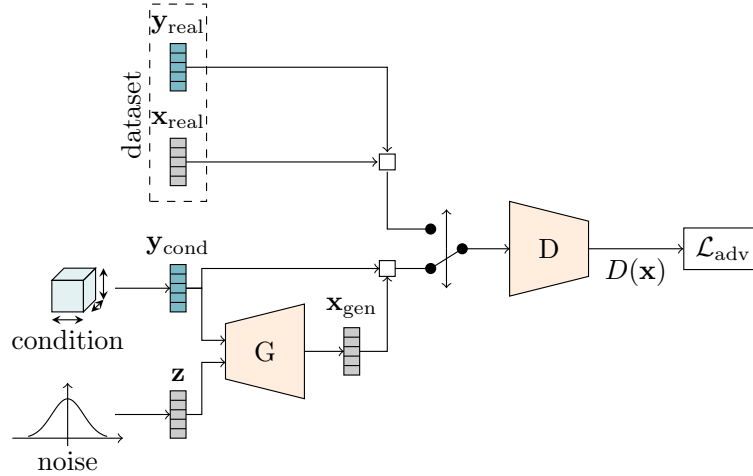


Figure 11: **cGAN with continuous parameters:** The generator G generates a point cloud \mathbf{x}_{gen} from a random vector \mathbf{z} and a continuous parameter \mathbf{y} . The discriminator either receives a set of point cloud and parameter from the real $\{\mathbf{y}_{\text{real}}, \mathbf{x}_{\text{real}}\}$ or the generated distribution $\{\mathbf{y}_{\text{cond}}, \mathbf{x}_{\text{gen}}\}$. It then outputs an estimate whether the set is real or generated with which the adversarial loss \mathcal{L}_{adv} is computed.




Figure 12: **Parameter handling details:** Generator and discriminator input handling for Vanilla and Projection cGAN.




Figure 13: **cGAN with additional regression**: The generator G generates a point cloud \mathbf{x}_{gen} from a random vector z and a regression label y . The discriminator either receives a real or a generated point cloud and predicts the probability of the sample originating from the real distribution. Additionally, the dimensions \mathbf{y}_{gen} of the generated point cloud are extracted which are then used to compute the regression error \mathcal{L}_{reg} .

For the second variant, we use the label input configuration proposed by CcGAN (Ding et al., 2021) to handle continuous parameters. Their approach is inspired by the method of label projection (Miyato and Koyama, 2018), which is why we refer to this variant as Projection cGAN. We used HVDL+NLI and conducted a coarse hyper-parameter search in steps of magnitudes (e.g. $10^{-1,2,3,\dots}$ to cover a wide range of options. Details for the generator and discriminator are shown in Fig. 12c and Fig. 12d, respectively. In contrast to Vanilla cGAN, both the latent and the label vectors are passed through a linear layer first, after which both are added together element-wise. For the discriminator, the label vector is fed through a linear layer after which the inner product with the features from PointNet is calculated. The features are fed through the final set of linear layers after which the result is added to the result of the inner product.

As mentioned in the main paper, both variants did not achieve satisfying results. The Vanilla variant ignored the conditioning entirely which led to very high MSE values (about four magnitudes higher than our proposed method). The performance in terms of FPD is close to the backbone. For the Projection variant, we observed very unstable training courses that often led to a collapse of the training with the model performing significantly worse in all metrics compared to all other models.

C.2.2 cGAN with Regression

The application of influencing the dimensions of an object has the major advantage that it is possible to directly compute the dimensions from the generated object to check whether the generator worked properly. This can also be used as a training signal. Fig. 13 shows an alternative approach for our proposed architecture, where a standard unconditioned discriminator is combined with an additional regression component. The generation of the desired dimensions is explicitly enforced with the regression loss, therefore we refer to this variant as the Regression cGAN. The generator loss is defined as

$$\mathcal{L}_G = -\mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(G(\mathbf{y}_{\text{cond}}, \mathbf{z}))] + \lambda_{\text{reg}} \cdot \mathcal{L}_{\text{reg}}(\mathbf{y}_{\text{cond}}, \mathbf{y}_{\text{gen}})$$

where \mathbf{y}_{gen} are the actual dimensions calculated from the generated point cloud, and λ_{reg} is the weighting factor for the regression loss. The generator loss is solely responsible to enforce the adherence of the dimension conditioning since the discriminator is not conditioned on the label input. The discriminator only judges from which distribution a sample originates from. Its loss function is defined as

$$\mathcal{L}_D = \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [D(G(\mathbf{y}_{\text{cond}}, \mathbf{z}))] - \mathbb{E}_{\mathbf{x} \sim \mathcal{R}} [D(\mathbf{x}_{\text{real}})] + \mathcal{L}_{\text{gp}}$$

The training behavior of the Regression cGAN is fundamentally different to the one of our proposed method. We observe that FPD is optimized first, hitting a minimum value at a quite early point during training where MSE

Table 4: **Quantitative Comparison:** We report results for the classes “Chair” and “Airplane”. All evaluations are conducted on a hold-out validation split. For both FPD and MSE smaller values indicate a better performance. MSE is given in %.

Model	Chair		Airplane	
	FPD	MSE	FPD	MSE
Reg. cGAN	1.4420	1.85	0.8802	20.19
Ours	1.5290	0.28	0.8691	0.30




Figure 14: **Dataset Class Distribution:** Shown is the class distribution of the ShapeNetPart dataset (Yi et al., 2016). We use the five largest classes for our experiments: “Table”, “Chair”, “Airplane”, “Car”, and “Lamp”.

is still quite high. The results for this checkpoint are listed in Table 4. From this point onward MSE is further minimized at the expense of FPD performance. In contrast to our proposed method, Regression cGAN aims at minimizing MSE down to zero while accepting FPD values that are magnitudes higher than for our proposed method. This means no realistic object shapes are being generated.

Considering these results and the fact that Regression cGAN can only be used for specific applications where \mathbf{y}_{gen} can easily be retrieved from the generated data, we found that our proposed method is superior to Regression cGAN. Our method offers a much easier and more stable handling of training while resulting in a very good performing model that is applicable to many scenarios.

D ADDITIONAL RESULTS

Table 5 displays the quantitative results of the five largest classes of the ShapeNetPart dataset and Table 6 gives details on the region-based performance. The class distribution of the dataset is shown in Fig. 14. For our experiments we did not consider classes with less than 1,000 object shapes. The quantitative results, especially when comparing our method to the chosen baselines, also depend on the distribution of the parameter we influence. For comparison, Fig. 15 shows the distribution of the five largest classes in terms of their extent in object height and width. While some classes, like “Table” have a wide distribution, others are more densely packed, like “Car”. Especially for the “Lamp” class, the stretching baseline (B2) achieves bad performance in terms of FPD, which can be attributed to its unique distribution.

Fig. 16 gives an overview on the quality and the diversity of the generated samples for the five different shapes.

Table 5: **Quantitative Comparison:** We report results for the five largest classes of the ShapeNetPart dataset in terms of the metrics used by Shu et al. (2019) and Achlioptas et al. (2018). Additionally, we report the regression error (MSE) for our introduced task. We freshly trained TreeGAN (Shu et al., 2019) to serve as our backbone network and report the results for reference. The baselines B1 and B2 are described in Sec. 5.3 of the main paper. Baseline 1 is the backbone network where ten versions per object are samples and the one with the best matching dimensions is chosen. Baseline 2 uses the backbone and then scales the object to the desired dimensions. All evaluations are conducted on a hold-out validation split. Shu et al. (2019) use the entire dataset for training, therefore values might vary slightly.

Shape	Model	FPD (\downarrow)	MSE [%] (\downarrow)	MMD (\downarrow)		COV (\uparrow)		JSD (\downarrow)
				CD	EMD	CD	EMD	
Table	Backbone	4.5245	–	0.0023	0.0863	0.4750	0.3750	0.1671
	Baseline 1	3.3009	52.88	0.0026	0.0912	0.4875	0.3500	0.1423
	Baseline 2	4.2851	0.00	0.0028	0.0920	0.5375	0.3125	0.1661
	Ours	3.0692	0.16	0.0019	0.1073	0.4875	0.3000	0.1313
Chair	Backbone	0.9525	–	0.0020	0.1027	0.4875	0.2500	0.1082
	Baseline 1	1.3674	26.07	0.0023	0.1013	0.4750	0.2625	0.1123
	Baseline 2	1.9259	0.00	0.0021	0.1003	0.4875	0.2625	0.1068
	Ours	1.5290	0.28	0.0022	0.1059	0.4625	0.3125	0.1434
Airplane	Backbone	1.2947	–	0.0002	0.0805	0.4375	0.1375	0.1887
	Baseline 1	1.0209	15.08	0.0003	0.0812	0.4500	0.1125	0.1819
	Baseline 2	1.6613	0.00	0.0003	0.0783	0.5250	0.1375	0.1834
	Ours	0.8691	0.30	0.0003	0.0724	0.5000	0.1250	0.1291
Car	Backbone	1.0816	–	0.0009	0.0656	0.4250	0.2375	0.0692
	Baseline 1	2.6293	5.15	0.0009	0.0651	0.4500	0.2000	0.0743
	Baseline 2	2.2045	0.00	0.0009	0.0634	0.4375	0.2750	0.0670
	Ours	1.7129	0.69	0.0010	0.0708	0.4125	0.1250	0.0714
Lamp	Backbone	2.9954	–	0.0037	0.1630	0.4500	0.2750	0.2642
	Baseline 1	3.4737	79.65	0.0029	0.1614	0.4500	0.2125	0.2569
	Baseline 2	36.5025	0.00	0.0041	0.1606	0.4500	0.2250	0.2653
	Ours	7.5012	0.77	0.0038	0.1917	0.4375	0.1750	0.2576
Total	Backbone	2.1697	–	0.0018	0.0996	0.4550	0.2550	0.1595
	Baseline 1	2.3584	35.77	0.0018	0.1000	0.4625	0.2275	0.1535
	Baseline 2	9.3159	0.00	0.0020	0.0989	0.4875	0.2425	0.1577
	Ours	2.9363	0.44	0.0018	0.1096	0.4200	0.2075	0.1466

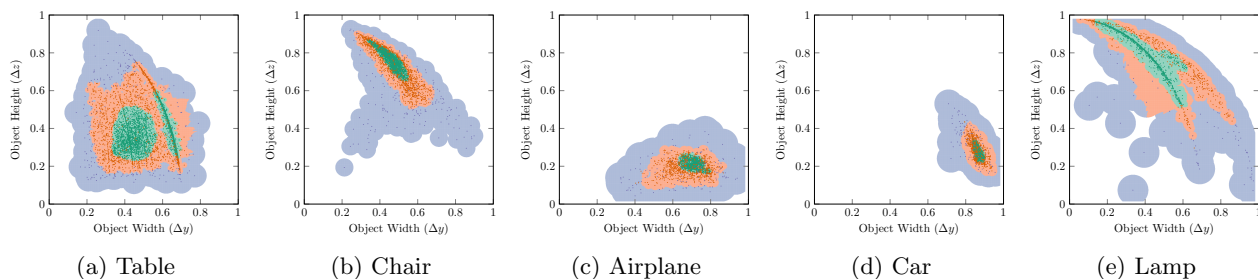


Figure 15: **Region-classified dimension distribution:** Shown is the sample distribution of five different object classes according to their extent in height z and width y . The length x of the object is not considered in this visualization. Each mark corresponds to one shape in the dataset. The three colors represent the resulting regions from k -nearest neighbor classifier with $k=20$ based on a KDE. The regions correspond to $1\sigma \approx 68\%$ (green), $2\sigma \approx 27\%$ (orange), and $3\sigma = 5\%$ (blue) of the entire data distribution. (Best viewed in color)

Table 6: **Region-based performance:** The table shows performances for three sampling regions ($\sigma_1, \sigma_2, \sigma_3$) of the data distribution for our architecture. The model is trained with two different sampling strategies each: our proposed version where labels are sampled from the KDE of the distributions (areas of Fig. 4), and the default version where labels are sampled from the list of labels contained in the dataset (marks of Fig. 4). For each σ -region, 1000 samples are generated. We report the MSE of the dimension regression and the FPD of the generated samples. Especially for less densely populated regions, i.e. σ_3 , our sampling strategy achieves better results.

Shape	Label Sampling	FPD (\downarrow)			MSE [%] (\downarrow)		
		σ_1	σ_2	σ_3	σ_1	σ_2	σ_3
Table	from distribution [ours]	4.7881	11.1556	83.1944	0.14	0.27	17.49
	from dataset samples	5.5063	23.1117	153.7859	0.25	1.15	52.01
Chair	from distribution [ours]	3.1137	2.4310	18.3729	0.22	0.26	6.51
	from dataset samples	2.5915	3.8257	138.0703	0.13	0.41	25.30
Airplane	from distribution [ours]	2.8487	1.6660	22.4949	0.20	0.36	7.62
	from dataset samples	2.5138	2.0912	42.4436	0.17	0.39	5.17
Car	from distribution [ours]	4.7105	11.8133	32.1120	0.61	0.89	3.74
	from dataset samples	3.7981	8.1593	60.2905	0.77	1.13	4.13
Lamp	from distribution [ours]	5.7873	20.2216	111.7557	0.44	1.84	11.50
	from dataset samples	4.4660	26.0117	290.8956	1.36	8.16	69.97




Figure 16: **Example Showcase:** The figure shows shapes generated with our proposed method from randomly sampled latent and conditioning parameters. The conditioning parameters are sampled from the KDE of the real data distribution, therefore only shapes for realistic dimensions are shown here.