

3D Semantic Subspace Traverser: Empowering 3D Generative Model with Shape Editing Capability

Ruowei Wang Yu Liu Pei Su Jianwei Zhang Qijun Zhao✉
Sichuan University

{ruoweiwang, liuyuvincen, supei}@stu.scu.edu.cn {zhangjianwei, qjzhao}@scu.edu.cn

Abstract

Shape generation is the practice of producing 3D shapes as various representations for 3D content creation. Previous studies on 3D shape generation have focused on shape quality and structure, without or less considering the importance of semantic information. Consequently, such generative models often fail to preserve the semantic consistency of shape structure or enable manipulation of the semantic attributes of shapes during generation. In this paper, we proposed a novel semantic generative model named 3D Semantic Subspace Traverser that utilizes semantic attributes for category-specific 3D shape generation and editing. Our method utilizes implicit functions as the 3D shape representation and combines a novel latent-space GAN with a linear subspace model to discover semantic dimensions in the local latent space of 3D shapes. Each dimension of the subspace corresponds to a particular semantic attribute, and we can edit the attributes of generated shapes by traversing the coefficients of those dimensions. Experimental results demonstrate that our method can produce plausible shapes with complex structures and enable the editing of semantic attributes. The code and trained models are available at https://github.com/TrepangCat/3D_Semantic_Subspace_Traverser

1. Introduction

Building a generative model to synthesize 3D shapes is a predominant topic in computer graphics and 3D vision since it is widely used in applications such as 3D content creation/reconstruction [9, 20, 44, 42], autonomous driving [3, 7, 16], AR/VR [21], and robotics [19, 57]. Typically, a 3D shape generative model takes a random vector as input and produces a 3D shape as output. To evaluate the effectiveness of generative models, researchers consider three aspects: quality, diversity, and controllability. The ideal 3D shape generative model should produce visually realistic and diverse 3D shapes while also enabling local structure

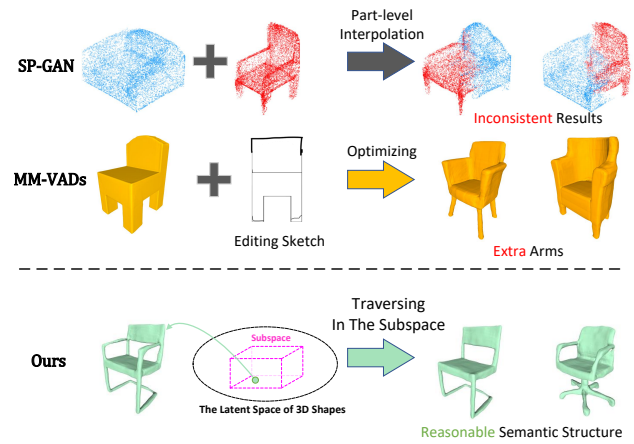


Figure 1. Different Editing Paradigm. Since SP-GAN is a point cloud generative model, we present its results as point clouds. As illustrated, the first row shows a typical method using part-level interpolation for shape editing. The second row is an example of the optimization-based method for shape editing. The last row shows how we traverse the local linear subspace to produce edited shapes with reasonable semantic structures.

manipulation to achieve controllability.

Previous 3D shape generative models like variational autoencoder (VAE) [18, 37, 39, 65, 61], generative adversarial network (GAN) [60, 1, 38, 8, 29, 34, 69], flow-based models [64, 35, 33], and diffusion models [48, 66, 48, 70] have been applied to various 3D representations, including point clouds [1], voxels [60], mesh [18] and so on. By adopting the rapidly developed implicit function-based shape representation [8, 4, 53], the generated shapes show promising quality and diversity [69]. However, the ability to control and manipulate the local structures of generated shapes during the generation process remains a challenging task. It is crucial to develop methods for effective shape control that allows for the modification of local structures during the generation process.

Previous methods that edit local shapes can be classified into two groups: interpolation-based methods [58, 61, 38, 51, 26, 39, 43, 40] and optimization-based methods

[23, 12, 13, 46, 9, 20]. The interpolation-based ones utilize part-level interpolation to edit shapes. However, since it does not consider the semantic relation between shape parts, it produces inconsistent results (see in the first row of Fig. 1). The optimization-based methods take the editing inputs, which are sketches, point moving operations, and so on, to optimize shapes. But its results focus on matching the editing inputs while ignoring other shape parts (see in the second row of Fig. 1). In light of this, we propose a new paradigm that utilizes the learned semantic information in the subspace (see in the last row of Fig. 1). During the generation process, we can edit shapes by traversing semantic dimensions in the subspace to produce diverse results. More details are illustrated as follows:

We present 3D Semantic Subspace Traverser, a novel semantic generative model for 3D shape generation and editing by leveraging the proposed local linear subspace models and the latent-space GAN. We adopt deep implicit function as the 3D representation and apply our GAN to the latent shape space, which is produced by a VAE. The GAN generates shape features in the form of a shape code grid, which can be decoded into 3D shapes by the VAE decoder. To empower our 3D generative model with shape editing capability, we embed the local linear subspace models into the GAN. Those local linear subspace models can discover semantic dimensions from feature maps in an unsupervised manner. By traversing along those semantic dimensions, our GAN can control the semantic attributes of the generated shapes, as shown in Fig. 1. The proposed 3D Semantic Subspace Traverser generates plausible 3D shapes and enables semantically editing of 3D shapes, as evidenced by quantitative and qualitative results.

To summarize, the contributions of this work are:

- We propose 3D Semantic Subspace Traverser, a novel semantic generative model that facilitates semantic shape editing for 3D shapes.
- We present a new latent-space GAN that leverages shape code grids to enable implicit-function-based 3D generation, producing 3D shapes with diverse topological structures.
- We introduce a 3D local linear subspace model to unsupervisedly and progressively mine interpretable and controllable dimensions from generator layers.
- Superior performance has been achieved in both 3D shape generation and editing.

2. Related Work

2.1. Generative Models for 3D Shape

GAN-based methods. 3D-GAN [60] first introduces the GAN network to the voxel-based 3D generation, since the

format of the voxel is suitable for the neural network. G2L-GAN [58] produces the voxels with the part labels. PA-GENet [36] generates the voxel parts of the shapes and assembles them together. PC-GAN [1] first uses the GAN network for the point cloud and proposes the latent-space GAN (l-GAN) which trains the GAN on the latent space of a pre-trained autoencoder. Tree-GAN [56] uses tree structures [10] for the GAN network with graph convolutions. Progressive-PCGAN [2] and PDGN [28] progressively generate the point cloud from a low resolution to a high resolution. Inspired by StyleGAN [31], MRGAN [14] and SP-GAN [38] use AdaiN [27] for part-aware generation and shape editing. While SP-GAN manually chooses parts for editing, MRGAN learns to split shapes into parts in an unsupervised manner. Because the voxel requires large memory footprints and the point cloud lacks topological structures, researchers turn to deep implicit functions [8, 53] to represent complex shapes with little memory cost. IM-GAN [8] designs a l-GAN to generate the occupancy values of the shapes. In IM-GAN, the GAN generates the shape code vectors which are further decoded to the occupancy values by a pre-trained decoder. Implicit-Grid [29] proposes a grid-based implicit function for shape generation, which can capture more local details and enables the spatial control of the generation process. ShapeGAN [34] proposes an SDF-based GAN and studies how to combine the signed distance function (SDF) with the voxel-based or the point-based discriminator. GET3D [15] generates meshes with textures and is trained with losses defined on rendered 2D images. SDF-styleGAN [69] introduces the advanced StyleGAN2 [32] to 3D shape generation, and proposes both global and local discriminators for the generated SDF values.

Other generative methods. PQ-NET [61] proposes a sequence-to-sequence autoencoder to generate 3D shapes via sequential part assembly. PolyGen [52] predicts the vertices and faces of the meshes sequentially by using a Transformer-based framework. GCA [68] introduces a new generative cellular automata to generate 3D shapes progressively. In addition, the flow-based generative model [64, 35, 33], and the diffusion model [48, 62] are also applied to the point cloud. 3DILG [67] takes the transformer as the backbone and applies irregular grids to the implicit function, which allows shape codes to have arbitrary positions instead of pre-defined ones.

Our method combines a shape code grid [41] with a novel l-GAN which consists of a variational autoencoder (VAE) and a GAN. Compared to IM-GAN and Implicit-Grid which also use the l-GAN, our continuous shape code grid naturally solves the discontinuity problem in the generated shapes by the trilinear interpolation. Different from the discrete grids in Implicit-Grid, ShapeFormer [63], and AutoSDF [50], our continuous one is better at represent-

Method	Representation	Part-level Interpolation	Semantic Editing
3D-GAN [2017]	voxel	✗	✗
G2L-GAN [2018]	voxel	✗	✗
PAGENet [2020]	voxel	✓	✗
GCA [2021]	voxel	✗	✗
PC-GAN [2018]	point cloud	✗	✗
Tree-GAN [2019]	point cloud	✗	✗
PointFlow [2019]	point cloud	✗	✗
DPF-Net [2020]	point cloud	✗	✗
MRGAN [2021]	point cloud	✓	✗
SP-GAN [2021]	point cloud	✓	✗
IM-GAN [2019]	implicit function	✗	✗
ShapeGAN [2020]	implicit function	✗	✗
Implicit-Grid [2021]	implicit function	✗	✗
SDF-styleGAN [2022]	implicit function	✗	✗
3D Semantic Subspace Traverser	implicit function	✓	✓

Table 1. Typical 3D generative models and their applications.

ing shapes with small grid resolution. Since the shape code grids produced by the VAE follow the Gaussian distribution, our GAN can easily learn to generate the shape code grids with complex structures. For shape editing, thanks to the spacial prior of the shape code grid and the local linear subspace model, we can do part-level interpolation and semantic editing. In Tab. 1, we summarize typical 3D generative models and their applications.

2.2. Shape Editing

The 3D part-aware generative methods, such as PQ-NET [61], PAGENet [36], MRGAN [14], and SP-GAN [38], can do shape editing via part-level shape interpolation. PQ-NET and PAGENet learn the part segmentations from the part labels. MRGAN learns to disentangle the shape parts without any part-level supervision. In SP-GAN, a sphere prior provides the part information. There are some 2D-to-3D shape editing methods, such as Sketch2Mesh [20] and MM-VADs [9]. Sketch2Mesh uses an encoder-decoder network to reconstruct shapes from sketches and refines 3D shapes through differentiable rendering. MM-VADs proposes a general multimodal generative model to edit and reconstruct 3D shapes from 2D sketches. Both of them can manipulate a 3D shape by editing its corresponding 2D sketch. DIF-NET [12] proposes a deformed implicit field representation, which builds the dense correspondences among shapes via a template. Thanks to the dense correspondences, DIF-NET can edit a shape by moving a 3D point to a new position. Similarly, DualSDF [23] proposes a new representation with two levels of granularity, i.e., the SDF and the primitive. DualSDF can edit the SDF of the shape by moving the primitives. SPAGHETTI [26] can directly edit neural implicit shapes and enable part-level control by applying transforms, blending, and inserting fragments of different shapes to local regions of the generated object.

ShapeCrafter [13] utilizes text descriptions to recursively edit shapes.

There are some non-generative part-aware 3D shape editing methods, such as Wei et al.[59] and DeepMetaHandles[45]. Wei et al.[59] relies on pre-defined semantic attributes, limiting editing possibilities. DeepMetaHandles[45] learns to move a set of control points, lacking the ability to add or remove parts.

For shape editing, our method does not need any part-level supervision or dense correspondences. We utilize the spacial prior of the shape code grid and the proposed local linear subspace model to discover the semantic dimensions of the local regions. By traversing those semantic dimensions, we can edit the semantic attributes of the shapes, such as adding or removing parts, changing the type of legs, and so on.

2.3. Interpretability Learning for GANs

GANSpace [24] uses principal component analysis (PCA) in the feature space to identify the important latent directions, which represent the interpretable variations. SeFa [55] tries to find the semantic directions by decomposing the weights of the pre-trained generator. EigenGAN [25] embeds the linear subspace models to the intermediate feature maps of the generator to discover the semantic dimensions. By traversing those dimensions, EigenGAN can edit the semantic attributes of the outputs.

The methods mentioned above can only learn to control the global semantic attributes. In contrast, we propose the local linear subspace model to learn both the local and global semantic attributes of a feature map. We use differentiable embedding to embed the local linear subspace model to a local region of the feature map. The size of the region we embed determines whether the subspace model learns local or global attributes.

3. Method

3.1. Overview

We design 3D Semantic Subspace Traverser, a novel semantic 3D generative model for 3D shape generation and editing, as illustrated in Fig. 2. In the following sections, we first introduce our local linear subspace model and how to embed it into the local regions of the feature map by differentiable embedding in Sec. 3.2. Second, we present the details of our 3D Semantic Subspace Traverser in Sec. 3.3. Then, we talk about how we encode shapes into the latent shape code grids in Sec. 3.4. Lastly, we present the loss function and implementation details in Sec. 3.5 and Sec. 3.6 respectively.

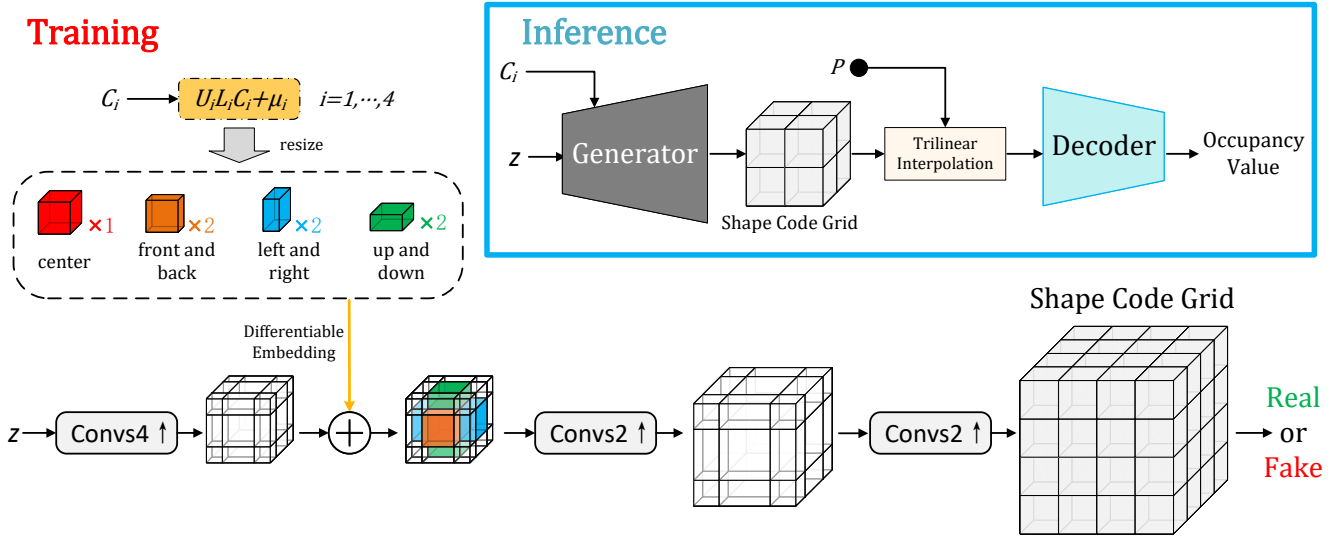


Figure 2. The overview of our 3D Semantic Subspace Traverser. z and C_i are random vectors. $\text{Convs}X$ represents the transposed convolutional block with stride X . To learn the semantic information, we embed four local linear subspace models (the red cube, the blue cubes, the green cubes, and the orange cubes) into the feature map by differentiable embedding. It is important to note that two orange cubes (as well as blue and green ones) are combined into a single subspace, while they are embedded in distinct regions. It is viable by using transform matrices containing two identity matrices. While the resolution of the feature map is r^3 , the edge length of the red cube is $\frac{1}{2}r$. We embed the red cube into the center of the feature map, and other cubes are embedded into the positions beside the red cube. We present a brief inference process in the top right corner to show how we generate and edit shapes.

3.2. Local Linear Subspace Model

Different from the linear subspace model in EigenGAN [25], our local linear subspace models can be embedded into local regions to learn semantics. In the forward propagation, we first take a random vector $C = [c_1, \dots, c_n]^T$ as the coordinates to sample a point $\phi \in R^{c \times d \times h \times w}$ from the local linear subspace:

$$\phi = ULC + \mu = \sum_{i=1}^n u_i l_i c_i + \mu \quad (1)$$

Here $U = [u_1, \dots, u_n]$, $u_i \in R^{c \times d \times h \times w}$ are the orthonormal bases. n is the number of bases. $L = \text{diag}(l_1, \dots, l_n)$ is a diagonal matrix in which l_i is the importance of the basis vector u_i . μ is the origin of the subspace.

Then, we use differentiable embedding to embed this point ϕ into the local region of the feature map $m \in R^{c \times D \times H \times W}$:

$$\hat{m} = m + f(\phi, X, Y, Z) \quad (2)$$

$$f(\phi, X, Y, Z) = \sum_d \sum_h \sum_w \phi_{cdhw} X_{dD} Y_{hH} Z_{wW} \quad (3)$$

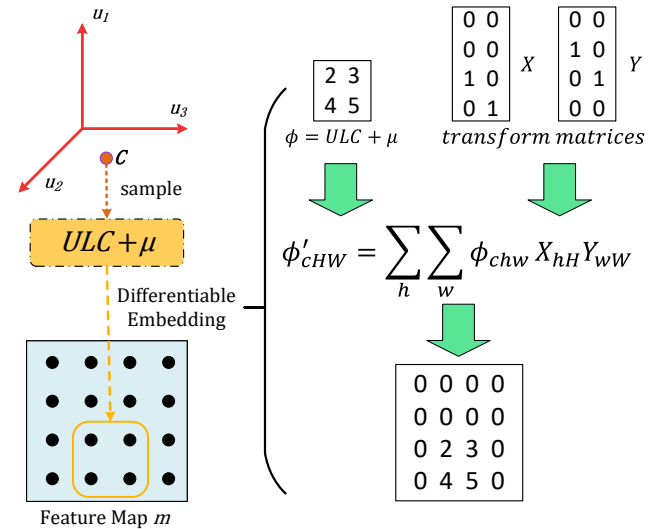


Figure 3. 2D Example of the embedding process. We first sample a point ϕ from the subspace. Then to make the embedding process differentiable, we use transform matrices instead of indexing to embed the point ϕ into the specified local region of the feature map m .

Since the sizes of m and ϕ are not the same, embedding by indexing is not differentiable. To solve this problem, we use a padding function f to expand the size of ϕ for differentiable embedding. As shown in Eq. (3), we express the tensor calculations by Einstein notation. $X \in R^{d \times D}$, $Y \in R^{h \times H}$, and $Z \in R^{w \times W}$ are transform matrices that consist of a square identity matrix and zeroes, like $\begin{bmatrix} 0 & I & 0 \end{bmatrix}^T$. Those transform matrices expand different dimensions of the ϕ respectively and control the position of the ϕ after padding. In other words, they decide which region of the feature map m we embed the ϕ into. By this means, we replace the indexing with the transform matrices, which makes the embedding differentiable. Fig. 3 shows a 2D example to illustrate the differentiable embedding. All the parameters of the local linear subspace model $S = (U, L, \mu)$ are learnable and are trained with the holistic 3D Semantic Subspace Traverser network.

After training, the local linear subspace model captures the semantic attributes of the local region in the feature map m . We can semantically manipulate the feature map by changing the coordinates C of the sample point in the subspace.

3.3. 3D Semantic Subspace Traverser

Fig. 2 shows the overview of our 3D Semantic Subspace Traverser. We embed four local linear subspace models into the feature map m between the generator’s layers to control the semantic attributes as:

$$\hat{m} = m + \sum_{i=1}^4 g(f(\phi_i, X_i, Y_i, Z_i)) \quad (4)$$

Here g is an identity function or a simple convolution. The embedding layout is illustrated in Fig. 2. We design a centrosymmetric layout according to the symmetry of training data. In this layout, the center subspace makes global changes to shapes, and other subspaces edit local shape parts. Such an embedding layout, in our opinion, is generally effective and can manipulate both local and global regions of shapes.

Our generator takes a random noise $z \sim N(0, 1)$ and a set of random vectors $\{C_i | C_i \sim N(0, 1), i = 1, \dots, 4\}$ as input. C_i is the coordinates for the i^{th} local linear subspace to sample a point. $G(z, C_i)$ is the generated shape code grid.

We design a 3D PatchGAN discriminator D to focus on the patches of the shape code grid. The discriminator follows the basic idea of the 2D PatchGAN discriminator from Pix2Pix [30]. It classifies some overlapping $N \times N \times N$ patches in the shape code grid as real or fake. By this means, we penalize the shape code grid at the scale of patches to improve the diversity of the generated shapes. More network details of the GAN are provided in the supplementary material.

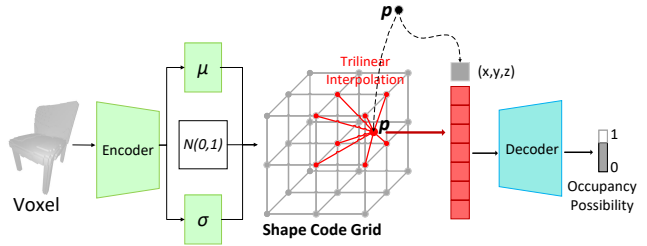


Figure 4. The VAE for shape encoding. In this figure, we take a 3^3 shape code grid as an example. Seven red squares are the shape features of the point \mathbf{p} and its six surrounding points.

At the inference time, the generated shape code grid is decoded into occupancy values by a pretrained implicit decoder, which we will talk about in Sec. 3.4. And we can extract the mesh from the occupancy values by marching cubes algorithm [47].

3.4. Shape Encoding

To train our 3D Semantic Subspace Traverser, we take a VAE V to encode 3D shapes into the latent shape code grids. The holistic framework of our VAE is shown in Fig. 4. We train the VAE in a voxel-superresolution manner. The encoder first encodes r^3 voxels into the mean μ and the standard deviation σ of a distribution. Then we sample a shape code grid $s \in R^{c \times g \times g \times g}$ from this distribution by reparameterization. c and g are the channel and resolution of the shape code grid respectively.

The decoder takes a point \mathbf{p} and a shape code grid s as the inputs to predict the occupancy values of the point \mathbf{p} . First, we extract the shape features $F(\mathbf{p})$ from the grid at location \mathbf{p} by using trilinear interpolation as shown in Fig. 4. Then, in order to get the local features around the location \mathbf{p} , we also extract the shape features $F(\mathbf{p}_i)$ of six surrounding points \mathbf{p}_i . Those surrounding points are in a distance d from \mathbf{p} along three Cartesian axes. We concatenate the shape features $F(\mathbf{p})$, the local features $F(\mathbf{p}_i)$, and the coordinates of \mathbf{p} together, and feed them to the implicit decoder to get the occupancy value of \mathbf{p} .

Once trained, the shape code grid s produced by the encoder is used to train our 3D Semantic Subspace Traverser. More network details of the VAE are provided in the supplementary material.

3.5. Model Learning

The loss functions of our 3D Semantic Subspace Traverser are the non-saturating loss \mathcal{L}_{GAN} [17], the R_1 regularization \mathcal{R}_1 [49], and the subspace regularization loss $\mathcal{L}_{subspace}$ [25]. To be specific, the loss of the generator is

$$\mathcal{L}_G = \mathbb{E}[\zeta(-D(G(z, C_i)))] + \alpha \mathcal{L}_{subspace} \quad (5)$$

Here $G(\cdot)$ and $D(\cdot)$ are outputs of the generator and discriminator respectively. $\mathcal{L}_{subspace} = \|U^T U - I\|_F^2$. α is the weight of $\mathcal{L}_{subspace}$ and $\alpha = 1$. $\zeta(x) = \log(1 + \exp(x))$.

The loss of the discriminator is

$$\mathcal{L}_D = \mathbb{E}[\zeta(D(G(z, C_i)))] + \mathbb{E}[\zeta(D(s))] + \beta \mathcal{R}_1(s, \Theta_D) \quad (6)$$

Here s is the shape code grid produced by the encoder of VAE. Θ_D are the parameters of the discriminator D . β is the weight of \mathcal{R}_1 and $\beta = 1$.

We train our VAE by the binary cross-entropy loss and the KL loss:

$$\mathcal{L}_{VAE} = \mathcal{L}_{BCE}(V(a, \mathbf{p}), o) + KL(N(\mu, \sigma) || N(0, 1)) \quad (7)$$

Here \mathcal{L}_{BCE} is the binary cross entropy loss. a is the input voxel. $V(\cdot)$ is the predicted occupancy values of VAE. o are the ground truth occupancy values. KL is the KL loss. μ and σ are the mean μ and the standard deviation produced by the encoder of VAE.

3.6. Implementation Details

In 3D Semantic Subspace Traverser, we set $n = 6$, which is the number of dimensions of all subspace models. In experience, when $n \leq 3$, the subspace models have difficulty in learning semantics. When $n > 3$, the subspace model is able to capture semantics. $n = 6$ is chosen as a trade-off between the learned semantics and the number of parameters. In the VAE, we set $r = 64$, $c = 8$, $g = 16$, and $d = 0.0536$.

4. Experiments

4.1. Dataset and Details

We use the ShapeNet Core (v1) dataset [5] to train our network. For comparison, we adopt the data split of [69] [22]: 70% of the data for training, 20% of the data for testing, and the rest for validating which is not used in our experiments. For data preprocessing, we follow the steps in [11]. To be specific, we first convert the meshes to watertight ones. Second, we normalize those meshes to make the lengths of their largest bounding box edges equal to one. Third, we sample some points and compute their occupancy values for each mesh. Finally, we voxelize meshes to produce 64^3 voxels. Our VAE is trained with five categories of the ShapeNet Core (v1) dataset, including the chair, airplane, table, car, and rifle. Different from the VAE, we train one specific GAN for each specific category. In other words, our GAN is category-specific.

In the following subsections, we first present the shape editing results in Sec. 4.2. Then, we show the quantitative evaluation of our method and other 3D generative methods in Sec. 4.3. Finally, we present the voxel super-resolution experiment in Sec. 4.4.



Figure 5. Editing Comparison with VAE+PCA. We use our method and VAE+PCA to edit the origin shape and produce four results.

4.2. Semantic Attributes for Shape Editing

We show various shape edit results in Fig. 6 and Fig. 7, where “Dim= i ” is the i^{th} dimension of all local linear subspace models. In five categories, different kinds of semantic attributes are learned by different dimensions. As shown in Fig. 6 and Fig. 7, moving along those dimensions, chairs’ semantic attributes of the generated shapes change. We can edit both local attributes, e.g., “seat” in Dim = 4, 5, 6, 9, “back” in Dim = 8, 10, 19, and “leg” in Dim = 12, 14; and global attributes, e.g., “type” in Dim = 0, 2, and “width” in Dim = 18. Specifically, when we traverse along the 4th and 6th dimensions, the same attributes are edited but with opposite results. The 12th and 14th dimensions show similar results. The 12th dimension translates all chairs to the ones with star-shaped legs, while the 14th dimension changes all legs to two types. Our method maintains the right semantic structures without inconsistency after editing shape parts. Since our training is unsupervised, there may be some levels of semantic entanglement. Despite this, each dimension typically maintains a dominant semantic attribute. More importantly, the resulting edits are reasonably coherent and surpass the outcomes produced by the VAE+PCA method in Fig. 5. We will discuss this comparison in more detail later.

Fig. 7 shows semantic attributes learned from different categories. For example, we transfer airliners to fighter planes by traversing the 1th dimension. Moreover, it is noticeable that certain types of data, such as chairs, exhibit more pronounced semantic structure changes compared to rifles or cars. It is because our method leverages subspace models, which excel at capturing these larger structural variations, leading to stronger visual results with the categories like the chair.

Comparison with VAE+PCA. We provide a meaningful baseline for comparison, which is to apply principal components analysis (PCA) to the VAE encoder output space to discover semantic dimensions. In Fig. 5, we show some editing results of our method and VAE+PCA. As illustrated, VAE+PCA typically generates unreasonable float artifacts, whereas our results preserve the proper semantic structures.

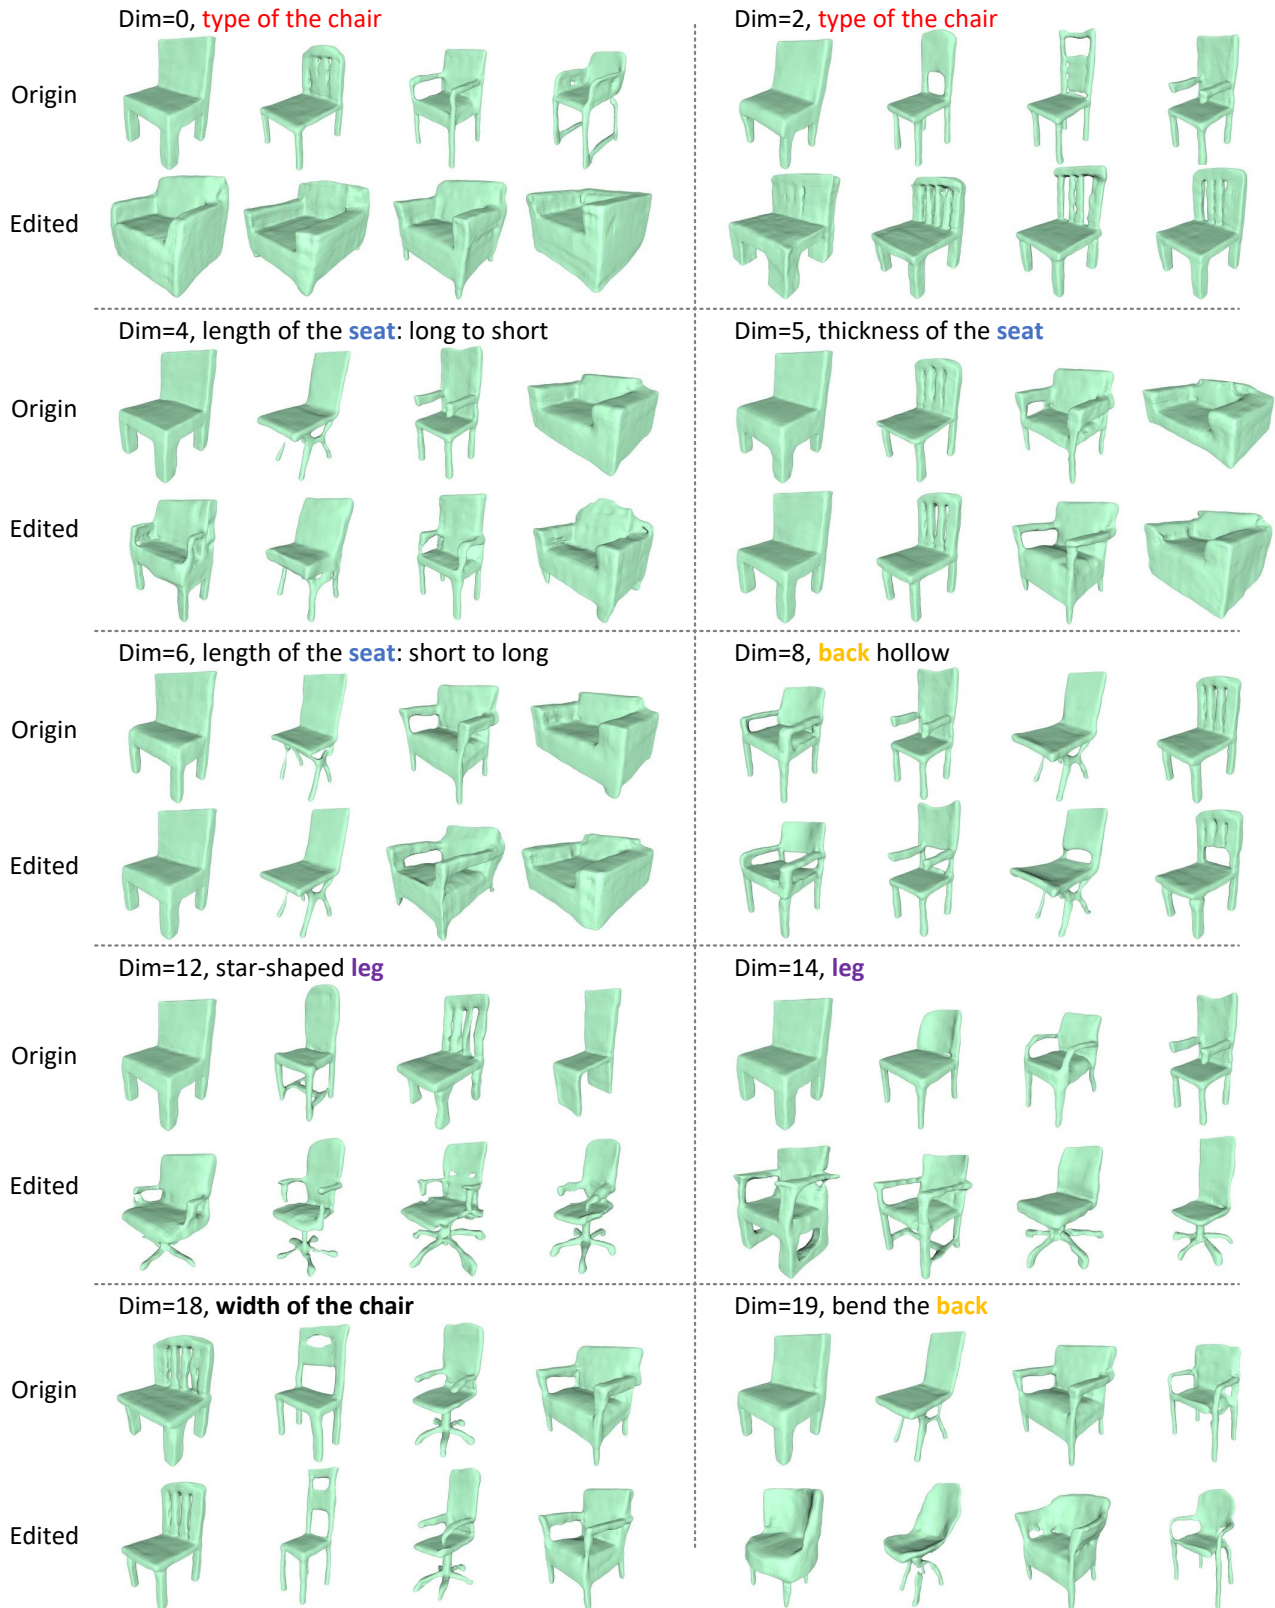


Figure 6. The learned semantic attributes of the local linear subspaces in the chair category. “Dim= i ” denotes the i^{th} dimension of all local linear subspace models.

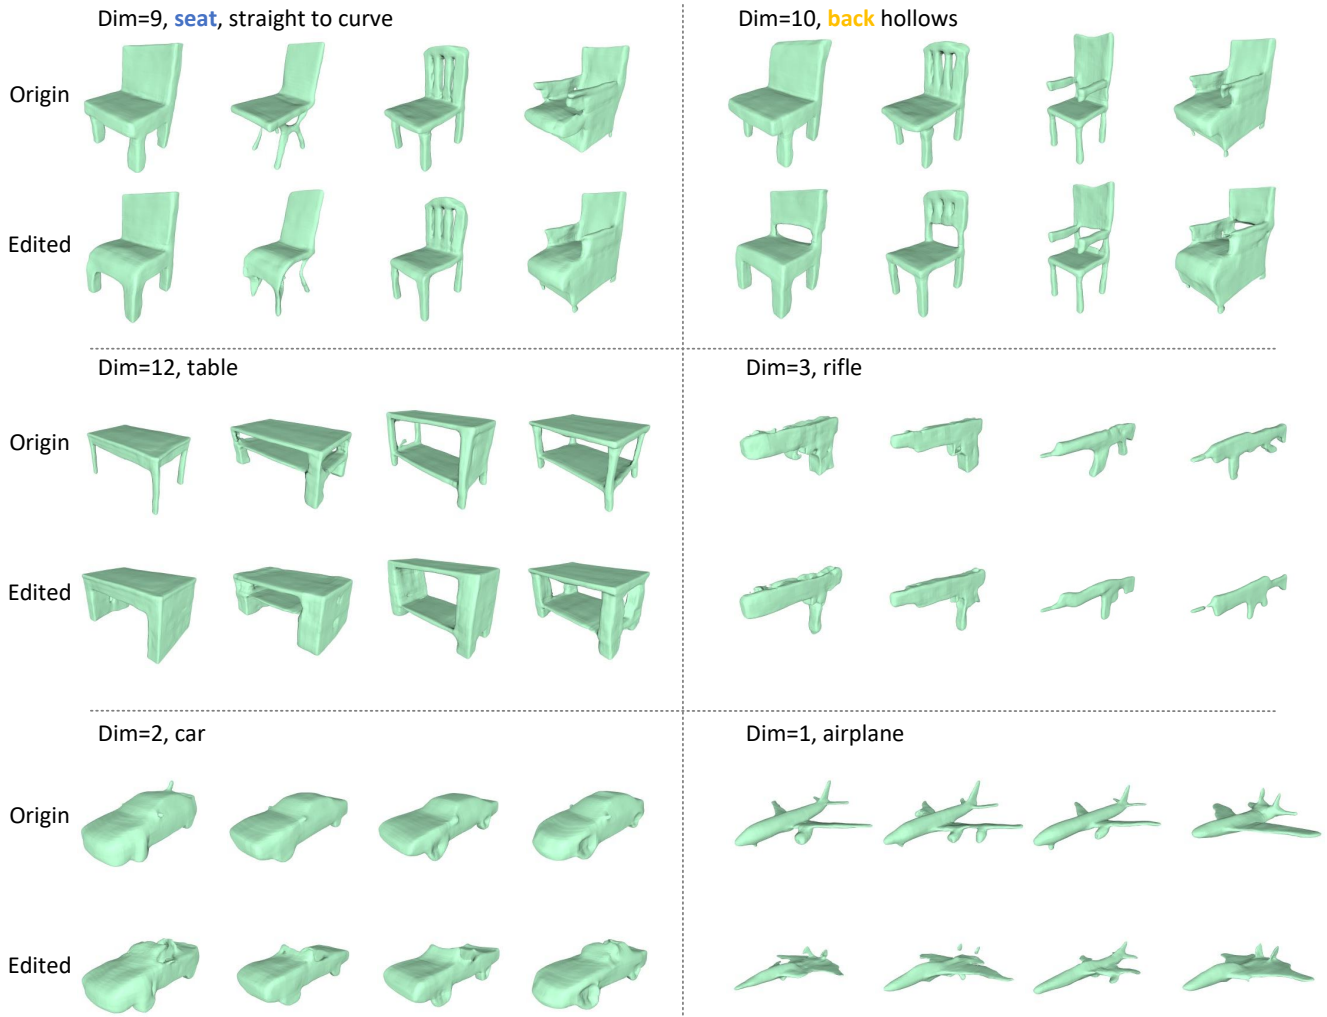


Figure 7. The learned semantic attributes of the local linear subspaces in various categories. “Dim= i ” denotes the i^{th} dimension of all local linear subspace models.

In conclusion, our approach captures semantic information more effectively than PCA.

4.3. Quantitative Evaluation of Shape Generation

To evaluate our shape generation quality, we compare our method to other representative methods that use implicit functions for 3D generation: IM-GAN [8], ShapeGAN [34], Implicit-Grid [29], and SDF-StyleGAN [69]. Among these methods, IM-GAN, Implicit-Grid, and our method use occupancy values to represent 3D shapes. ShapeGAN and SDF-StyleGAN use SDF values to represent 3D shapes.

For a fair comparison, we take ECD [29] and Shading-image-based FID (S-FID) [69] as our evaluation metrics, and also report COV [1] and MMD [1] values as the ref-

erence. Following the experiments setting in [69], ECD, MMD, and COV are all based on the light-field-distance (LFD) [6, 69] to measure the distance between two samples. LFD is based on silhouette images and measures the structure similarity between two shapes. When computing COV, MMD, and ECD, we take the testing dataset as the reference to compare with our generated shapes, which is five times the number of the reference. When computing S-FID, we use the training dataset as the reference, and the number of generated shapes is the same as the reference. We use the clean-fid algorithm [54] to compute the FID.

We show the quantitative results in Tab. 2. COV is listed for reference only as mentioned before. In all categories, our method performs best in MMD. And for three classes

Data	Method	COV(%) \uparrow	MMD \downarrow	ECD \downarrow	S-FID \downarrow
Chair	IMGAN	72.57	3326	1998	63.42
	Implicit-Grid	82.23	3447	1231	119.5
	ShapeGAN	65.19	3726	4171	126.7
	SDF-StyleGAN	75.07	3465	1394	36.48
	Ours	78.61	3160	810	99.64
Airplane	IMGAN	76.89	4557	2222	74.57
	Implicit-Grid	81.71	5504	4254	145.4
	ShapeGAN	60.94	5306	6769	162.4
	SDF-StyleGAN	74.17	4989	3438	65.77
	Ours	73.42	3879	1933	133.4
Car	IMGAN	54.13	2543	12675	141.2
	Implicit-Grid	75.13	2549	8670	209.3
	ShapeGAN	57.40	2625	14400	225.2
	SDF-StyleGAN	73.60	2517	6653	97.99
	Ours	64.67	1509	10355	219.5
Table	IMGAN	83.43	3012	907	51.70
	Implicit-Grid	85.66	3082	1089	87.69
	ShapeGAN	76.26	3236	1913	103.1
	SDF-StyleGAN	69.80	3119	1729	39.03
	Ours	79.78	2698	539	99.95
Rifle	IMGAN	71.16	5834	701	103.3
	Implicit-Grid	77.89	5921	357	125.4
	ShapeGAN	46.74	6450	3115	182.3
	SDF-StyleGAN	80.63	6091	510	64.86
	Ours	62.73	4251	1207	154.4

Table 2. The quantitative evaluation of different 3D generation methods.

with complex topological structures, i.e., chair, airplane, and table, we also achieve the best results in ECD. It shows that our method is good at generating shapes with complex topological structures. But in S-FID, we are not the best due to the unsmooth surfaces.

4.4. Voxel Super-Resolution

This section evaluates our shape compression ability by the voxel super-resolution experiment. We compare our VAE with IM-NET [8] and IF-NET [11] (see in Tab. 3). For a fair comparison, we use five classes in the data split of IF-NET, which are chairs, airplanes, tables, cars, and rifles. We train all models for 200 epochs. All methods take 32^3 voxels as inputs, and the output meshes are sampled at 128^3 . To measure the quality, we employ three recognized measures [11]: volumetric intersection over union (IoU) evaluates how well the defined volumes overlap, Chamfer-L2 evaluates the accuracy and completeness of the shape surface, and normal consistency evaluates the accuracy and completeness of the shape normals. Tab. 3 shows that our method achieves the second-best results, which is reasonable. In contrast to IM-NET, our method utilizes a grid to represent the shape code rather than a vector, which improves the quality of local regions’ reconstruction. And IF-NET outperforms us in terms of metrics since it utilizes multi-scale features. Our method makes a balance between

Method	IoU \uparrow	Chamfer- $L_2 \times 10^2 \downarrow$	Normal-Consis. \uparrow
IM-NET	0.754	0.160	0.877
IF-NET	0.907	0.002	0.960
Ours	0.804	0.107	0.925

Table 3. Quantitative Evaluation of voxel super-resolution.

quality and network complexity. Additionally, our main contribution is in the GAN part. Based on Tab. 3, we believe that the current VAE is good enough to support our method.

5. Conclusion and Discussion

Conclusion. In this paper, we introduce a novel semantic 3D generative model, 3D Semantic Subspace Traverser, which can edit semantic attributes of 3D shapes. We apply a 3D Semantic Subspace Traverser to the latent shape code and use our local linear subspace models to capture the semantic information. Once trained, each dimension of subspaces corresponds to a specific semantic attribute, and we can edit shapes by traversing along those dimensions.

Limitations & Discussion. The limitation of our method is the handcrafted embedding layout. It would be our future work to design a mechanism that can automatically choose the embedding layout.

And for the current handcraft embedding layout, we provide a generally effective design in Fig. 2 for users. It is also free to propose different embedding layouts for different tasks. There is some advice for designing the embedding layout. First, depending on where you wish to learn the semantics, you would do better to pick the suitable subspace size and embedding locations. Second, the number of dimensions in a subspace should be greater than three; otherwise, it is hard to learn the semantics.

Automatic dimension selection for editing is still an open issue in the literature. A possible solution is based on 3D segmentation: first, segment an original shape and its corresponding edited shape into semantic parts, and then compare corresponding parts to decide whether the semantic dimension is meaningful. Moreover, since our learned attributes are unpredictable, taking texts as guidance to lead dimensions to learn specific semantic attributes is a viable solution. These are also among our future work.

Acknowledgement. The research was sponsored by the National Natural Science Foundation of China (No. 62176170).

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3D point clouds. In *ICML*, 2018. 1, 2, 8
- [2] Mohammad Samiul Arshad and William J. Beksi. A progressive conditional generative adversarial network for generating dense and colored 3D point clouds. In *3DV*, 2020. 2
- [3] Aseem Behl, Omid Hosseini Jafari, Siva Karthik Mustikovela, Hassan Abu Alhaija, Carsten Rother, and Andreas Geiger. Bounding boxes, segmentations and object coordinates: How important is recognition for 3D scene flow estimation in autonomous driving scenarios? In *ICCV*, 2017. 1
- [4] Rohan Chabra, Jan E Lenssen, Eddy Ilg, Tanner Schmidt, Julian Straub, Steven Lovegrove, and Richard Newcombe. Deep local shapes: Learning local sdf priors for detailed 3D reconstruction. In *ECCV*, 2020. 1
- [5] Angel X. Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, Jianxiong Xiao, Li Yi, and Fisher Yu. ShapeNet: An Information-Rich 3D Model Repository. Technical Report arXiv:1512.03012 [cs.GR], Stanford University — Princeton University — Toyota Technological Institute at Chicago, 2015. 6
- [6] Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. On visual similarity based 3D model retrieval. In *CGF*, 2003. 8
- [7] Yiping Chen, Jingkang Wang, Jonathan Li, Cewu Lu, Zhipeng Luo, Han Xue, and Cheng Wang. LiDAR-Video Driving Dataset: Learning Driving Policies Effectively. In *CVPR*, 2018. 1
- [8] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 1, 2, 8, 9
- [9] Zezhou Cheng, Menglei Chai, Jian Ren, Hsin-Ying Lee, Kyle Olszewski, Zeng Huang, Subhransu Maji, and Sergey Tulyakov. Cross-modal 3D shape generation and manipulation. *arXiv preprint arXiv:2207.11795*, 2022. 1, 2, 3
- [10] Zhou Cheng, Chun Yuan, Jiancheng Li, and Haiqin Yang. Treenet: Learning sentence representations with unconstrained tree structure. In *IJCAI*, 2018. 2
- [11] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *CVPR*, 2020. 6, 9
- [12] Yu Deng, Jiaolong Yang, and Xin Tong. Deformed implicit field: Modeling 3D shapes with learned dense correspondence. In *CVPR*, 2021. 2, 3
- [13] Rao Fu, Xiao Zhan, Yiwen Chen, Daniel Ritchie, and Srinath Sridhar. Shapecrafter: A recursive text-conditioned 3D shape generation model. *NeurIPS*, 2022. 2, 3
- [14] Rinon Gal, Amit Bermano, Hao Zhang, and Daniel Cohen-Or. MRGAN: Multi-rooted 3D shape representation learning with unsupervised part disentanglement. In *ICCVW*, 2021. 2, 3
- [15] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3D: A generative model of high quality 3D textured shapes learned from images. In *NeurIPS*, 2022. 2
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 1
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 5
- [18] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3D surface generation. In *CVPR*, 2018. 1
- [19] Joris Guerry, Alexandre Boulch, Bertrand Le Saux, Julien Moras, Aurélien Plyer, and David Filliat. Snapnet-r: Consistent 3D multi-view semantic labeling for robotics. In *ICCVW*, 2017. 1
- [20] Benoit Guillard, Edoardo Remelli, Pierre Yvernay, and Pascal Fua. Sketch2Mesh: Reconstructing and Editing 3D Shapes from Sketches. In *ICCV*, 2021. 1, 2, 3
- [21] Lei Han, Tian Zheng, Yinheng Zhu, Lan Xu, and Lu Fang. Live semantic 3D perception for immersive augmented reality. *TVCG*, 2020. 1
- [22] Christian Häne, Shubham Tulsiani, and Jitendra Malik. Hierarchical surface prediction for 3D object reconstruction. In *3DV*, 2017. 6
- [23] Zekun Hao, Hadar Averbuch-Elor, Noah Snively, and Serge Belongie. Dualsdf: Semantic shape manipulation using a two-level representation. In *CVPR*, 2020. 2, 3
- [24] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *NeurIPS*, 2020. 3
- [25] Zhenliang He, Meina Kan, and Shiguang Shan. Eigengan: Layer-wise eigen-learning for gans. In *ICCV*, 2021. 3, 4, 5
- [26] Amir Hertz, Or Perel, Raja Giryes, Olga Sorkine-Hornung, and Daniel Cohen-Or. SPAGHETTI: Editing Implicit Shapes Through Part Aware Generation. *arXiv preprint arXiv:2201.13168*, 2022. 1, 3
- [27] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 2
- [28] Le Hui, Rui Xu, Jin Xie, Jianjun Qian, and Jian Yang. Progressive point cloud deconvolution generation network. In *ECCV*, 2020. 2
- [29] Moritz Ibing, Isaak Lim, and Leif Kobbelt. 3D shape generation with grid-based implicit functions. In *CVPR*, 2021. 1, 2, 8
- [30] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017. 5
- [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 2
- [32] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2

- [33] Hyeongju Kim, Hyeonseung Lee, Woo Hyun Kang, Joun Yeop Lee, and Nam Soo Kim. Softflow: Probabilistic framework for normalizing flow on manifolds. In *NeurIPS*, 2020. 1, 2
- [34] Marian Kleineberg, Matthias Fey, and Frank Weichert. Adversarial Generation of Continuous Implicit Shape Representations. In *EG*, 2020. 1, 2, 8
- [35] Roman Klokov, Edmond Boyer, and Jakob Verbeek. Discrete point flow networks for efficient point cloud generation. In *ECCV*, 2020. 1, 2
- [36] Jun Li, Chengjie Niu, and Kai Xu. Learning part generation and assembly for structure-aware shape synthesis. In *AAAI*, 2020. 2, 3
- [37] Jun Li, Kai Xu, Siddhartha Chaudhuri, Ersin Yumer, Hao Zhang, and Leonidas Guibas. GRASS: Generative recursive autoencoders for shape structures. *TOG*, 2017. 1
- [38] Ruihui Li, Xianzhi Li, Ka-Hei Hui, and Chi-Wing Fu. SP-GAN: Sphere-guided 3D shape generation and manipulation. *TOG*, 2021. 1, 2, 3
- [39] Shidi Li, Miaomiao Liu, and Christian Walder. EditVAE: Unsupervised parts-aware controllable 3D point cloud shape generation. In *AAAI*, 2022. 1
- [40] Feng Liu and Xiaoming Liu. Learning implicit functions for topology-varying dense 3D shape correspondence. In *NeurIPS*, 2020. 1
- [41] Feng Liu and Xiaoming Liu. Voxel-based 3D detection and reconstruction of multiple objects from a single image. In *NeurIPS*, 2021. 2
- [42] Feng Liu and Xiaoming Liu. 2D gans meet unsupervised single-view 3D reconstruction. In *ECCV*, 2022. 1
- [43] Feng Liu and Xiaoming Liu. Learning implicit functions for dense 3D shape correspondence of generic objects. *TPAMI*, 2023. 1
- [44] Feng Liu, Luan Tran, and Xiaoming Liu. Fully understanding generic objects: Modeling, segmentation, and reconstruction. In *CVPR*, 2021. 1
- [45] Minghua Liu, Minhyuk Sung, Radomir Mech, and Hao Su. Deepmetahandles: Learning deformation meta-handles of 3D meshes with biharmonic coordinates. In *CVPR*, 2021. 3
- [46] Zhengzhe Liu, Yi Wang, Xiaojuan Qi, and Chi-Wing Fu. Towards implicit text-guided 3D shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17896–17906, June 2022. 2
- [47] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3D surface construction algorithm. In *SIG-GRAPH*, 1987. 5
- [48] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3D point cloud generation. In *CVPR*, 2021. 1, 2
- [49] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *ICML*, 2018. 5
- [50] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3D completion, reconstruction and generation. In *CVPR*, 2022. 2
- [51] Kaichun Mo, Paul Guerrero, Li Yi, Hao Su, Peter Wonka, Niloy J. Mitra, and Leonidas J. Guibas. Structrnet: Hierarchical graph networks for 3D shape generation. *TOG*, 2019. 1
- [52] Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. Polygen: An autoregressive generative model of 3D meshes. In *ICML*, 2020. 2
- [53] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 1, 2
- [54] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022. 8
- [55] Yujun Shen and Bolei Zhou. Closed-form factorization of latent semantics in gans. In *CVPR*, 2021. 3
- [56] Dong Wook Shu, Sung Woo Park, and Junseok Kwon. 3D point cloud generative adversarial network based on tree structured graph convolutions. In *ICCV*, 2019. 2
- [57] Keisuke Tateno, Federico Tombari, Iro Laina, and Nassir Navab. Cnn-slam: Real-time dense monocular slam with learned depth prediction. In *CVPR*, 2017. 1
- [58] Hao Wang, Nadav Schor, Ruizhen Hu, Haibin Huang, Daniel Cohen-Or, and Hui Huang. Global-to-local generative model for 3D shapes. *TOG*, 2018. 1, 2
- [59] Fangyin Wei, Elena Sizikova, Avneesh Sud, Szymon Rusinkiewicz, and Thomas Funkhouser. Learning to infer semantic parameters for 3D shape editing. In *3DV*, 2020. 3
- [60] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In *NeurIPS*, 2016. 1, 2
- [61] Rundi Wu, Yixin Zhuang, Kai Xu, Hao Zhang, and Baoquan Chen. Pq-net: A generative part seq2seq network for 3D shapes. In *CVPR*, 2020. 1, 2, 3
- [62] Zijie Wu, Yaonan Wang, Mingtao Feng, He Xie, and Ajmal Mian. Sketch and text guided diffusion model for colored point cloud generation. *arXiv preprint arXiv:2308.02874*, 2023. 2
- [63] Xingguang Yan, Liqiang Lin, Niloy J Mitra, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Shapeformer: Transformer-based shape completion via sparse representation. In *CVPR*, 2022. 2
- [64] Guandao Yang, Xun Huang, Zekun Hao, Ming-Yu Liu, Serge Belongie, and Bharath Hariharan. Pointflow: 3D point cloud generation with continuous normalizing flows. In *ICCV*, 2019. 1, 2
- [65] Jie Yang, Kaichun Mo, Yu-Kun Lai, Leonidas J Guibas, and Lin Gao. DSG-Net: Learning disentangled structure and geometry for 3D shape generation. *TOG*, 2022. 1
- [66] Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten Kreis. LION: Latent point diffusion models for 3D shape generation. In *NeurIPS*, 2022. 1
- [67] Biao Zhang, Matthias Nießner, and Peter Wonka. 3DILG: Irregular Latent Grids for 3D Generative Modeling. In *NeurIPS*, 2022. 2

- [68] Dongsu Zhang, Changwoon Choi, Jeonghwan Kim, and Young Min Kim. Learning to generate 3D shapes with generative cellular automata. In *ICLR*, 2021. [2](#)
- [69] Xin-Yang Zheng, Yang Liu, Peng-Shuai Wang, and Xin Tong. SDF-StyleGAN: implicit sdf-based stylegan for 3D shape generation. *arXiv preprint arXiv:2206.12055*, 2022. [1](#), [2](#), [6](#), [8](#)
- [70] Linqi Zhou, Yilun Du, and Jiajun Wu. 3D shape generation and completion through point-voxel diffusion. In *ICCV*, 2021. [1](#)