

HyperSDFusion: Bridging Hierarchical Structures in Language and Geometry for Enhanced 3D Text2Shape Generation

Zhiying Leng^{1,2*} Tolga Birdal³ Xiaohui Liang^{2,4†} Federico Tombari¹

¹Technical University of Munich, Germany ³Imperial College London, U.K.

²Beihang University, China

⁴Zhongguancun Laboratory, China

{zhiyingleng, liang_xiaohui}@buaa.edu.cn, t.birdal@imperial.ac.uk, tombari@in.tum.de

Abstract

3D shape generation from text is a fundamental task in 3D representation learning. The text-shape pairs exhibit a hierarchical structure, where a general text like “chair” covers all 3D shapes of the chair, while more detailed prompts refer to more specific shapes. Furthermore, both text and 3D shapes are inherently hierarchical structures. However, existing Text2Shape methods, such as SDFusion, do not exploit that. In this work, we propose HyperSDFusion, a dual-branch diffusion model that generates 3D shapes from a given text. Since hyperbolic space is suitable for handling hierarchical data, we propose to learn the hierarchical representations of text and 3D shapes in hyperbolic space. First, we introduce a hyperbolic text-image encoder to learn the sequential and multi-modal hierarchical features of text in hyperbolic space. In addition, we design a hyperbolic text-graph convolution module to learn the hierarchical features of text in hyperbolic space. In order to fully utilize these text features, we introduce a dual-branch structure to embed text features in 3D feature space. At last, to endow the generated 3D shapes with a hierarchical structure, we devise a hyperbolic hierarchical loss. Our method is the first to explore the hyperbolic hierarchical representation for text-to-shape generation. Experimental results on the existing text-to-shape paired dataset, Text2Shape, achieved state-of-the-art results. We release our implementation under [HyperSDFusion.github.io](https://github.com/HyperSDFusion).

1. Introduction

Text-to-Shape synthesis [5, 7, 26, 31, 36] involves the task of generating high-quality and faithful shapes given a text prompt and holds significant promise for a wide range of applications including augmented/virtual reality and design, offering the potential for automated, diverse, and cost-

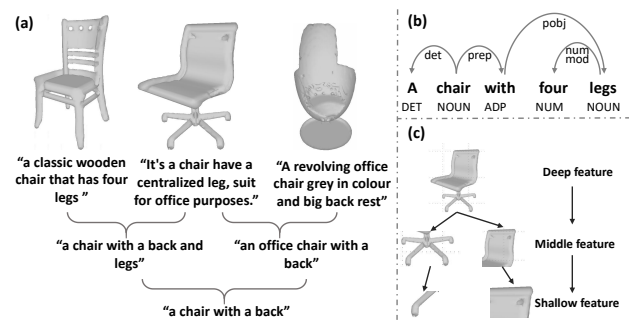


Figure 1. Hyperbolic text-shape representations. (a) The hierarchical structure between text and 3D shape. (b) The syntactic tree of text. (c) The hierarchical part-to-whole relationships of 3D shape.

effective 3D content. Unlike image-based media [1, 18, 60], natural language provides a more direct means of expression. However, effectively marrying the realms of 3D geometry and natural language is challenging, leading to no established standard for text-guided 3D shape generation.

We argue that *hierarchy*, preordering elements of a set in increasing complexity is fundamental to linking geometry and language as illustrated in Fig. 1. 3D shapes inherently exhibit *compositionality* [33, 44], possessing hierarchical part-to-whole relationships [32]. On the other hand, language exhibits a hierarchical tree-like syntactic structure [8, 12, 57], rooted in inter-word relationships. Recognizing these parallel hierarchical natures requires rethinking text-to-shape correspondence also within a similar hierarchical framework. For example, a general prompt like “a chair” can correspond to thousands of 3D shapes. In contrast, a more detailed description like “a wooden chair with armrests and four legs” narrows down the possible shapes to those with specific attributes. Fully embracing and leveraging such hierarchical nuances can significantly improve the fidelity and specificity of generated shapes, making strides in the field of text-to-shape synthesis.

Existing text-to-shape methods can be divided into two categories depending on the data type they handle: one for

*This work was done when the author was at TUM.

†corresponding author

paired text-shape data [5, 7, 14, 25, 28, 31, 47] and the other for unpaired data [21, 26, 36, 39, 41, 42, 50, 52]. Some methods [21, 26, 36, 39, 50, 52] for unpaired data generate intermediary images, and then transform these images into 3D shapes by 3D generative models, like NeRFs [26, 52]. Others [41, 42] leverage a shared text-image embedding space, generating 3D shapes by an image-3D generator. These methods do not directly learn the text-shape representation. In contrast, methods using text-shape paired data have the advantage of directly learning the text-shape representation by GANs [5], Variational Autoencoders [14, 31], or Diffusion Models [7, 25], generating 3D shapes from texts. To date, all these methods ignore the joint hierarchical structure of 3D shapes and natural language.

In this work, we focus on text-shape paired data. Inspired by prior works on leveraging hierarchies in images [2, 13, 20], point clouds [27, 32], or text-image pairs [9], we propose to embed the tree-like hierarchical structure of text and shape, jointly, into a more natural non-Euclidean, hyperbolic space. This incurs less distortion than in Euclidean space, primarily due to the exponential expansion of hyperbolic space ideally suited to representing trees [30]. Our approach, deemed **HyperSDFusion**, first utilizes a Signed Distance Field (SDF) based Autoencoder [7] to embed the SDF representation of 3D shapes into a compact latent space, learning a latent feature for each 3D shape. Then, to concurrently exploit the sequential (word order) and hierarchical structures of an input prompt, we propose a dual-branch latent diffusion model in hyperbolic space to generate a desired latent feature close to the ground truth latent feature from a noise. In one branch, we leverage a pre-trained text-image model [9], learning both sequential features of language and multi-modal hierarchical features of text-image, in hyperbolic space. In a parallel branch, we devise a hyperbolic text-graph convolution model that parses the input prompt into a syntax graph and learns the hierarchical features of language in hyperbolic space. Notably, to maintain the hierarchical structure of 3D shapes during the generation process, we introduce a hyperbolic hierarchical loss, which correlates the distance between 3D deep and shallow features with the distance to the origin of the Poincaré ball (hyperbolic space).

Finally, we conduct a series of experiments on the existing text-shape paired dataset, Text2Shape [5]. The experiments demonstrate that our method achieves high-quality generation results while preserving the hierarchical characteristics of text and shape. Our main contributions are:

- We are the first to learn a joint hierarchical representation of text and shape in the hyperbolic space, improving the quality of text-to-shape generation.
- We introduce a dual-branch diffusion to fully capture both sequential and hierarchical structures of texts in hyperbolic space.

- Our proposed hyperbolic hierarchical loss ensures that the generation process of the diffusion model maintains the hierarchical structure of 3D shapes.

2. Related works

Text-to-Shape Generation. In recent years, text-to-shape generation has garnered significant attention in 3D representation learning. In this task, annotating paired text-shape data is time-consuming and laborious. The Text2Shape dataset [5], a text-shape paired dataset, is widely used in text-to-shape generation. This dataset was proposed by Chen *et al.* as the first text-shape paired dataset, and Chen *et al.* implemented a GAN-based text-shape generation on it. Later, a series of methods [7, 14, 25, 28, 31, 47] were proposed for paired data, which can be divided into VAE-based methods, Autoencoder-based methods, Auto-regressive-based method, and diffusion model-based methods according to the advanced models used. In VAE-based methods [14, 31], Fu *et al.* [14] propose ShapeCrafter, a recursive text-shape generation method by recursively embedding text features. In Autoencoder-based methods [28, 47], Tian *et al.* [47] propose a structure-aware method to align the text feature space and the 3D shape feature space at the part level, by dividing the feature space. In Auto-regressive-based methods, Luo *et al.* [29] propose an improved Auto-regressive Model for 3D shape generation, by applying discrete representation learning in a latent vector instead of volumetric grids. In diffusion-based methods [7, 25], Cheng *et al.* [7] propose SDFusion that employs latent diffusion to generate an ideal latent feature, close to the ground truth latent feature embedded by a Quantised-VAE.

Recently, with the development of multi-modal learning, researchers have used existing multi-modal models to generate 3D by unpaired text-shape data. Some methods [26, 36, 39, 52] use images as intermediate generations, firstly leveraging pre-trained text-to-image models to produce images, then employ 3D reconstruction methods such as NeRFs to reconstruct 3D shapes from the images. Other methods [41, 42] leverage a shared text-image embedding space, generating 3D shapes by an image-3D generator. However, these methods cannot directly learn the representation between text and 3D. In this work, we focus on 3D generation by paired text-shape data, aiming to directly learn the representation between text and 3D.

Diffusion Models. Diffusion models as new powerful generative models have shown record-breaking performance in many applications, like 3D shape generation [17, 43, 48, 55, 58], image synthesis [11, 40, 51], human motion generation [19, 54, 56], video generation [3, 15], etc. Diffusion models can be divided into two categories based on whether they directly generate the final output: standard diffusion models [6, 56] and latent diffusion models [34, 40]. The

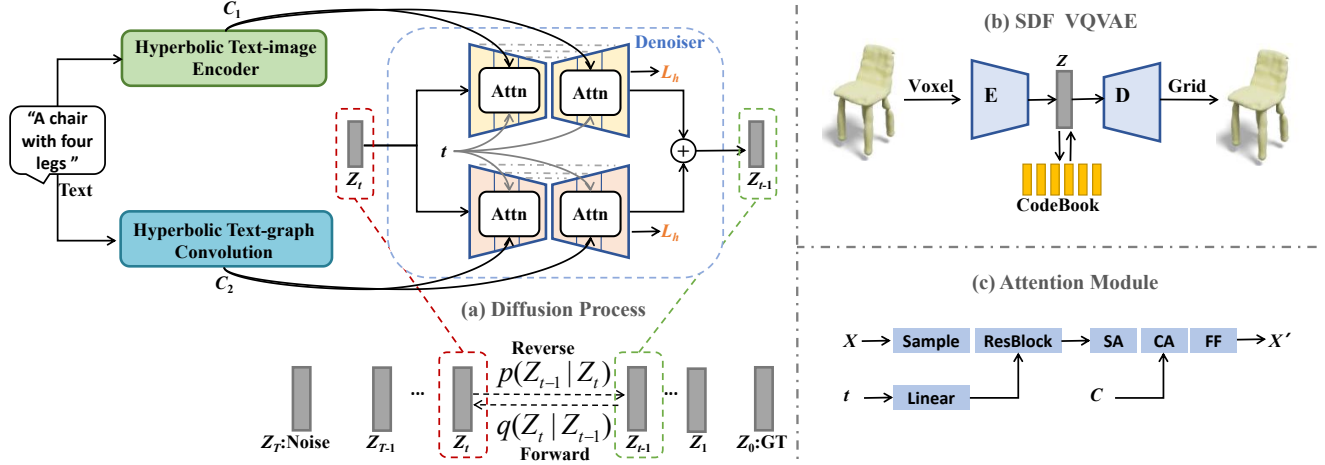


Figure 2. Overview of the proposed HyperSDFusion. (a) The forward and reverse processes of the proposed dual-branch diffusion model from Z_0 to Z_T . In particular, the detailed denoising process of the latent feature Z_t based on text conditions $\{C_1, C_2\}$ is showcased. (b) The architecture of a VQVAE for 3D shape represented by SDF. (c) The attention module in the denoiser of the diffusion model.

standard diffusion models directly generate the final output, such as images or 3D shapes. However, as the scale or resolution of the output increases, standard diffusion models consume significant GPU resources. The latent diffusion models utilize a learned latent space to generate a latent feature and then transform it into the final output, substantially reducing GPU consumption. In this work, we also employ the latent diffusion model to generate a latent feature, and then transform it into a 3D shape.

Hyperbolic Representation learning. In recent years, there has been a growing interest in deep representation learning in hyperbolic spaces [35, 53, 59]. It has been proved that hyperbolic space is more suitable than Euclidean space for processing data with a tree-like structure or power-law distribution due to its exponential growth property. There are a few works about hyperbolic representation learning in Computer Vision [2, 9, 13, 20, 23, 27, 32], such as learning hierarchical representation of images in hyperbolic space [2, 13, 20], analyzing and utilizing the hierarchical property of point cloud in hyperbolic space [27, 32], etc. Desai *et al.* [9] presented the hierarchical structure in text-image and proposed a hyperbolic contrastive learning model for text-image paired data. As far as we know, we are the first to learn the hierarchical representation of text and shape in hyperbolic space, which is beneficial for text-to-shape generation.

3. Preliminaries

Hyperbolic space. Hyperbolic space is a non-Euclidean space, also an n -dimensional Riemannian manifold of constant negative curvature. The hyperbolic space can be modeled by several isometric models [53]. The most popular model is the Poincaré ball model, which we adopt. The r -

dimensional Poincaré ball $B_c^n := \{x \in \mathbb{R}^n \mid \|x\|^2 < r^2\}$ with a negative curvature c , endowed with the canonical metric of the Euclidean space g_x^B admits the structure of a Riemannian manifold $\mathcal{B} := (B_c^n, g_x^B)$ with the geodesic distance $d(\cdot) : \mathcal{B} \times \mathcal{B} \rightarrow \mathbb{R}$ given by:

$$d(x, y) = \frac{1}{\sqrt{c}} \operatorname{arccosh} \left(1 + \frac{2\|x - y\|^2}{(1 - \|x\|^2)(1 - \|y\|^2)} \right). \quad (1)$$

Identifying its tangent space $\mathcal{T}_x \mathcal{B}$ with \mathbb{R}^n allows us to transform data between Euclidean and hyperbolic spaces through the exponential and logarithmic maps with origin at x , respectively denoted by Exp_x and Log_x . Explicit expressions are explained in Mettes *et al.* [30].

Hyperbolic Graph Convolution (HGC) HGC extends the standard graph convolution to hyperbolic spaces and is suitable for processing tree-like graph data. Generally, an HGC consists of feature initialization, feature updating and aggregation, and activation. Formally, a graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{x_i^E \mid i = 0, \dots, N\}$ is the node feature set in Euclidean space, and \mathcal{E} is the edge set. Firstly, an Exp function transfers node features to hyperbolic space, initializing node features on the hyperbolic manifold, x_i^B . Then a Möbius-layer [22] updates and aggregates node features, which is a generalization of the fully connected layer in hyperbolic space. Finally, a non-linear hyperbolic activation σ^B acts on these features by: (i) mapping them back to Euclidean space, (ii) transforming them by traditional (Euclidean) non-linear layers, and (iii) mapping them back to the hyperbolic space. This yields hyperbolic node features y_i^B . This procedure is defined as:

$$y_i^B = \sigma^B(\operatorname{Möbius}(\operatorname{Exp}(x_i^B))). \quad (2)$$

We refer the reader to Yang *et al.* [53] for further details.

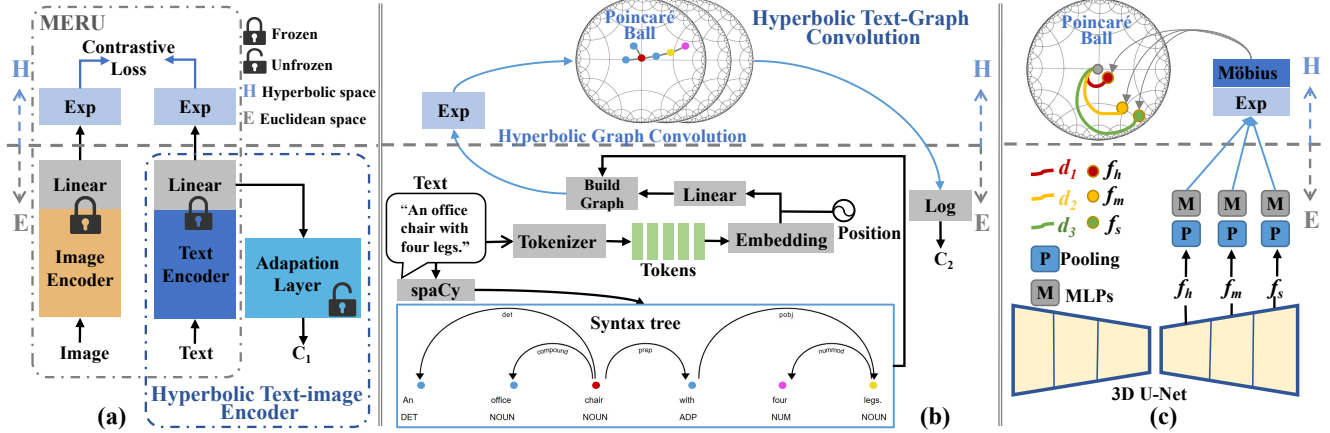


Figure 3. Illustration of our proposed modules. (a) Given a text, the hyperbolic text-image encoder learns both sequential and multi-modal hierarchical features of the text, C_1 . (b) The hyperbolic text-graph convolution module learns hierarchical syntactic features of the text, C_2 . (c) Hyperbolic hierarchical loss supervises the hierarchical structure of 3D shape features in hyperbolic space.

4. Method

In this section we present our method, called HyperSDFusion, for text-to-shape generation. Similar to the visual attribute hierarchies of Li *et al.* [24], we encode the semantic hierarchy in a hyperbolic space, where the root of the hierarchy (i.e., at the center of Poincaré ball) is a category, e.g., chair. The finer-grained separations are at lower levels, such as subcategories or details. Finally, the category-irrelevant features are at the lowest level, e.g., legs of a chair (*cf.* Fig. 1). We achieve this through hyperbolic text-shape feature learning under the supervision of our proposed hyperbolic hierarchical loss. We now describe our model and provide additional details in our supplementary material.

4.1. A Dual-branch Latent Diffusion Model

The architecture of our proposed HyperSDFusion for text-to-shape generation is shown in Fig. 2. HyperSDFusion is a dual-branch Latent diffusion model, including 3D shape compression, the forward and reverse process of the latent diffusion model based on text conditions.

3D Shape Compression As the scale of the 3D shape increases, the GPU consumption yields a significant increment. Embedding the 3D shape into the low-dimensional latent space tends to greatly reduce resource consumption. Hence, we encode 3D shapes into a compact latent space, representing each 3D shape as a latent feature.

As shown in Fig. 2(b), we firstly represent 3D shapes as a Truncated Signed Distance Field (TSDF) Γ of size $R \times R \times R \times 1$, where R is the resolution, and 1 is the dimension of distance. A Vector Quantised-Variational AutoEncoder (VQ-VAE) [49] is employed to learn the latent feature of Γ . The encoder E of VQ-VAE encodes Γ into a latent representation $Z = E(\Gamma)$, where $Z \in \mathbb{R}^{d \times d \times d \times m}$, d is the resolution of the latent feature and m is the dimen-

sion of features. Importantly, the encoder downsamples the TSDF by a factor $f = R/d$. Then, Z is discretized by the codebook VQ . Finally, the decoder D of VQ-VAE reconstructs the 3D shape from the discretized Z , represented as $\Gamma' = D(VQ(Z))$. The reconstruction loss between Γ and Γ' follows Van *et al.* [49].

Forward Process of The Latent Diffusion Model. With the learned latent feature Z of each shape as ground truth, the forward process is an iterative process that adds Gaussian noise to Z in a Markovian manner, as shown in Fig. 2(a). In detail, the ground truth latent feature of shape is the input at time 0, defined as Z_0 . Then Z_0 is noised by Gaussian noise ϵ time by time. After T times, Z_0 is noised to Z_T , which is close to standard Gaussian noise. The whole forward process is represented as:

$$q(Z_{1:T}|Z_0) = \prod_{t=1}^T q(Z_t|Z_{t-1}), \quad (3)$$

where q is the feature distribution at each time.

Reverse Process of The Latent Diffusion Model based on Text Conditions. Given a text prompt as the condition, the reverse process is to generate the latent feature of the 3D shape coincident with the text from random Gaussian noise. In detail, starting from random Gaussian noise Z_T at time T , we gradually sample it to remove noise by a reverse Markov chain. As shown in Fig. 2(a), the noisy latent feature at time t , Z_t , is denoised to Z_{t-1} by a denoiser conditioned on text features, C_1, C_2 . The denoiser is a dual-branch 3D U-Net structured attention model [40]. Each module in the attention model consists of an up/down-sample, residual blocks, self-attention, cross-attention, and a feed-forward layer, as shown in Fig. 2(c). After T times, the output Z_0 is close to the ground truth latent feature. The

reverse process is formulated as follows:

$$p(Z_{0:t}|Z_T) = p(Z_T) \prod_{t=1}^T p(Z_{t-1}|Z_t), \quad (4)$$

where p is the feature distribution in the reverse process.

Dual-branch Denoiser. To well utilize text information, we design a dual-branch denoiser. On the one hand, text is composed of ordered words, indicating a sequential structure. We introduce a hyperbolic text-image encoder to learn the multi-modal sequential feature of the text, defined as C_1 . On the other hand, natural language also exhibits a syntactic structure [8] that can be parsed into a syntax tree, as shown in Fig. 1(b). We propose a hyperbolic text-graph convolution module to learn the hierarchical structure of the text, defined as C_2 .

A common way to utilize the two conditions C_1, C_2 is to concatenate them as a single condition and inject them into a single-branch denoiser. However, this way may cause feature interference. Hence, we propose a dual-branch denoiser to utilize respectively, which consists of two parallel 3D U-Nets structures. As shown in Fig. 2(a), C_1 and C_2 are fed separately as conditions to the cross-attention of the 3D U-Nets. Finally, the outputs of the two branches are concatenated. The dual-branch denoiser can preserve the independence of the two conditions, preventing confusion or loss of information. In this way, one branch perceives the sequential structure from C_1 , while the other branch perceives the hierarchical structure from C_2 .

4.2. Text Feature Learning in Hyperbolic Space

As previously mentioned, text exhibits both sequential structure and syntactic hierarchical structure. We now expose our hyperbolic text-image encoder (HTIE) and hyperbolic text-graph convolution module (HTGC) used to better capture these structures.

Hyperbolic Text-image Encoder. Our hyperbolic text-image encoder learns text features, not only capturing the sequential structure of text but also embedding the multi-modal hierarchical structure of text-image. As shown in Fig. 3(a), HTIE consists of a transformer-based text encoder and an adaptation layer. The text encoder captures long-range dependencies between words in a text, learning the sequence features of text. Besides, the text encoder is pre-trained from MERU [9], a text-image contrastive learning model that learns the hierarchical features between text and image in a hyperbolic space. Hence, our text encoder also embeds multi-modal hierarchical features of text and image. Then, we introduce an adaptation layer to narrow the gap between the pre-trained text encoder and the text-shape dataset, implemented by multiple transformer layers. In this way, our HTIE learns a text condition, C_1 .

Hyperbolic Text-graph Convolution Module. A text can be parsed into a syntax tree according to its syntactic structure. As shown in Fig. 3(b), each word is endowed with Part-Of-Speech (POS), and the dependency between words is indicated by directed edges. Since hyperbolic space is more suitable for processing tree-like structure data than Euclidean space [53], we propose a hyperbolic text-graph convolution module to learn the syntactic structure of text, including syntax tree construction, text-graph initialization in Euclidean space, and learning in hyperbolic space.

Firstly, the input text is processed by spaCy [16], a natural language processing library yielding a syntax tree. Secondly, the node and edge features are initialized in the Euclidean space E . A given text prompt is converted into tokens and then mapped to a vector representation by an embedding layer. The vector representation is added with the position encoding, following a linear layer to transform features. At this point, the features of tokens in Euclidean space are initialized as node features V^E of the text-graph, and the branches of the syntax tree serve as edges \mathcal{E} in the text-graph. The text-graph is defined as $\mathcal{G}^E = \{V^E, \mathcal{E}\}$. Thirdly, the text-graph is projected to hyperbolic space, learning the hierarchical structure of the text-graph. In detail, the projected text-graph is represented as $\mathcal{G}^H = \text{Exp}(\mathcal{G}^E) = \{V^H, \mathcal{E}\}$, where Exp is as defined in Section 3, and H represents hyperbolic space. In hyperbolic space modeled by the Poincaré Ball model, stacked hyperbolic graph convolutions propagate and update node features on the text-graph, better capturing the hierarchical structure of the text-graph. At last, the updated node features of the text-graph are projected back into the Euclidean space by the Log , yielding the other text condition C_2 .

4.3. Hyperbolic Hierarchical Loss for 3D shape

In this work, the denoiser predicts the latent feature of 3D shapes from a random Gaussian noise conditioned on text, denoted as Z' . To ensure that Z' closely approximates the ground-truth latent feature, the mean squared error L_{mse} between Z and Z' is employed as the loss function. However, L_{mse} only supervises the similarity between features, ignoring the hierarchical structure of 3D shapes.

As mentioned before, 3D shapes inherently exhibit a hierarchical structure, namely the part-whole hierarchy. In the feature space of 3D shape, this hierarchy manifests as hierarchical relationships between deep and shallow features. In this work, these deep and shallow features are the multi-scale outputs of the 3D U-Net. As depicted in Fig. 3(c), the small-scale features f_h are the global feature of the 3D shape, the large-scale features f_s are the local feature, and the medium-scale features f_m line in between. To supervise these features to maintain the hierarchy, we propose a hyperbolic hierarchical loss to regularize the feature space.

These features are first passed through a pooling layer

and multi-layer perceptrons to unify their scale and feature dimension. Then, they are projected to the hyperbolic space by Exp , and transformed to a unified distribution by a shared Möbius layer. Due to the hierarchical property of deep and shallow features, features in the hyperbolic space should also maintain a tree-like hierarchical structure, as illustrated in Fig. 3(c). Here, we use the relative distance between features to constrain the feature distribution, i.e., $d_2 > d_1, d_3 > d_2$. d_i is the geodesic distance from the i -th feature to the center point. Therefore, the hyperbolic hierarchical loss is defined as:

$$L_h = \max(0, -d_2 + d_1) + \max(0, -d_3 + d_2). \quad (5)$$

The whole training loss is given by $L = L_{mse} + \alpha L_h$, where α is a balancing factor.

5. Experiments

To evaluate the performance of our method, we conducted a series of experiments on the paired text-shape dataset, Text2Shape [5]. The details of the experiments are described in the following subsections.

Implementation Details. Our method generates a 3D shape from a text, which is represented as a TSDF with the size of $64 \times 64 \times 64 \times 1$, where 64 is the resolution. Our method includes a 3D VQ-VAE and a dual-branch diffusion model. Firstly, we trained the 3D VQ-VAE on the ShapeNet dataset [4] with 13 categories of 3D shapes. The 3D VQ-VAE compresses the TSDF to a compact latent feature Z with the size of $16 \times 16 \times 16 \times 3$. Then, the 3D VQ-VAE is frozen, and the dual-branch diffusion model is trained on the Text2Shape dataset [5]. The optimizer for the training is AdamW with an initial learning rate of $1e-5$. During inference of the diffusion model, the sampler is the DDIM sampler [45], where the sample step T is set to 100.

Dataset. In this work, we choose Text2Shape dataset [5] to conduct experiments. The reason is that this work aims to learn hierarchical representations of text and 3D directly, which requires paired data. Text2Shape dataset [5] is a widely used dataset in paired text-to-shape generation. This dataset provides rich natural language annotations for tables and chairs in ShapeNet [4], describing their shape, color, texture, material, etc. This work only focuses on how to generate shapes from text.

Evaluation Metrics. In order to assess the generation quality, we adopt Intersection over Union (IoU), Chamfer Distance (CD), F-score, and Fréchet Inception Distance (FID) as evaluation metrics. IoU computes the intersection volume between generated and GT 3D shapes. CD is computed between generated and GT point clouds that are sampled from 3D shapes by farthest point sampling. F-score computes the harmonic mean between precision and recall, based on the distance between generated 3D shapes and

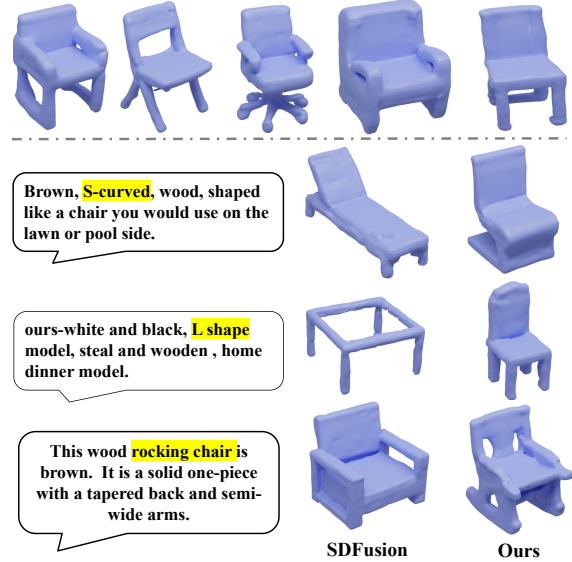


Figure 4. The showcase of text-to-shape generation results. Above the dotted line are some examples generated by our method, and below is the result compared to SDFusion [7].

Table 1. Evaluating text-to-shape generation on random 1000 samples of Text2Shape [5] test set. Bold indicates the best results.

Methods	IoU \uparrow	CD \downarrow	F-score \uparrow	FID \downarrow
Liu <i>et al.</i> [28]	12.21	1.41	13.34	1.00
SDFusion [7]	13.98	1.246	12.67	5.47
HyperSDFusion(Ours)	16.21	0.7433	15.16	0.70

GT 3D shapes. We set the distance threshold is set to 1%, as in Tatarchenko *et al.* [46]. FID calculates the Fréchet distance between the feature representations of images rendered from predicted and GT 3D shapes.

For evaluating the hierarchy of generated 3D shapes, we introduce two metrics, Hierarchical Mutual Difference (HMD) and Hyperbolic Distance (HD). HMD computes the Chamfer distance between 3D shapes generated from a general text and a detailed text, assessing the text-shape hierarchical structure. HD computes the geodesic distance in hyperbolic space among the deep and shallow features of 100 randomly selected generated 3D shapes, evaluating the hierarchical structure of 3D shapes.

5.1. Text-to-Shape Generation Results

Method for Comparison. There are few existing diffusion-based methods for paired text-to-shape generation, including SDFusion [7] and Diffusion-SDF [25]. SDFusion learns the mapping between a 3D latent space and text feature space, while Diffusion-SDF is for patch-wise latent spaces. Because our method aims to learn the hierarchical representation in a joint text-shape latent space, we compare our method with SDFusion [7] in the experiment. SDFusion learns text features in Euclidean space, and generates the

Table 2. Comparison results on texts with different lengths.

Method	words ≤ 8				8 < words ≤ 16				16 < words			
	IoU↑	F-sc.↑	CD↓	FID↓	IoU↑	F-sc.↑	CD↓	FID↓	IoU↑	F-sc.↑	CD↓	FID↓
SDFusion [7]	8.52	12.15	1.71	7.51	8.04	10.28	1.92	8.19	10.13	12.01	1.48	6.59
Ours	15.41	13.58	0.99	3.64	14.56	15.37	0.81	4.22	13.25	16.11	0.71	4.91

Table 3. Ablations on architectures and loss functions. ① refers to the single-branch diffusion model. ② represents the dual-branch diffusion model. In bold indicates the best results.

Group	Model	IoU↑	CD↓	F-score↑
0	Baseline	12.74	11.04	8.93
	Baseline+T5	12.26	11.29	8.14
1	Baseline+CLIP	13.99	9.41	8.77
	Baseline+HTIE	14.26	10.49	9.42
	Baseline+HTIE+HTGC+①	14.18	10.45	8.988
2	Baseline+HTIE+HTGC+②	15.73	8.99	9.50
	Baseline+HTIE+GCN+②	15.41	9.28	9.19
	Baseline+HTIE+HTGC+②+L _e	13.51	11.72	9.79
3	Baseline+HTIE+HTGC+②+L _h	16.44	9.053	9.917

3D shape without hierarchical supervision.

Results on Text-to-shape Generation. Tab. 1 compares our HyperSDFusion against the previous SOTA text-to-shape generation method, SDFusion [7]. Quantitatively, our method outperforms SDFusion by a significant margin (87% decrease in FID, 40% decrease in CD, 23% improvement in IoU, and 19.7% gain in F-score.). We achieved the best results on all these metrics for assessing generation quality. In particular, the large drop in FID indicates that shapes generated by our method are visually closer to the desired ones indicating high-quality text-to-shape generation. Some samples generated by our method are shown in Fig. 4. We observe that these shapes are complete and crisp in detail, when compared to SDFusion, which performs poorly on long texts. In contrast, our approach adequately captures text features and generates shapes faithful to the text, such as S-shapes, L-shapes, and rocking chairs.

Generation Performance on Texts with Different Lengths. For evaluating the generation performance on texts with different lengths, we conducted a comparison on Text2Shape++ dataset [14], which is built on Text2Shape dataset. In Text2Shape++, each text prompt is represented as a phrase sequence, and each phrase sequence corresponds to one or more shapes. We generate shapes with texts of length less than 8, more than 8 less than 16, and more than 16, respectively, to compare the performance of our method with SDFusion [7]. The comparison results are listed in Tab. 2. It can be observed that our method outperforms the existing methods on both long and short texts. We also showcase more generations of texts with different lengths in Fig. 6(a). Our method generates shapes consistent with both short and long texts. These results indicate that our method sufficiently learns and utilizes the text feature.




Figure 5. Visualizing text-shape hierarchical structure. Highlighted parts of the prompt represent the detailed information.

Table 4. The result of comparing the performance of capturing the text-shape hierarchical structure. More results are in the supplemental material.

Text	HMD↓	
	SDFusion [7]	Ours
"A chair."	-	-
"A soft chair."	1.29	0.57
"A brown color soft chair with 4 stands."	0.70	0.20
"A new-fashioned, grey soft chair, with two arms and four legs base."	0.60	0.43

5.2. Ablation Studies

We conducted ablation studies on a mini-set of Text2Shape dataset to demonstrate the effectiveness of our proposed method. We choose SDFusion [7] as the baseline because the text encoder of SDFusion is the standard sequential model in Euclidean space, Bert [10], and its denoiser is without the hierarchical supervision.

Hyperbolic Text-image Encoder. We compare our HTIE with the text encoder of the baseline, which is most used in existing methods [7, 25, 31]. We replace the text encoder of the baseline with our HTIE. As listed in Group 1 of Tab. 3, the baseline with our HTIE significantly improves the quantitative results, both IoU, CD, and F-score, indicating that our HTIE enhances the generation quality, owing to the multi-modal hierarchical features learned by the pre-trained MERU model.

In order to further verify the performance of our HTIE, we compare it with the other common text encoder, T5 [38], and the Euclidean text-image encoder, CLIP [37]. Results in Group 1 of Tab. 3 show that T5 performs worse than CLIP and our HTIE. This is because T5, pre-trained on text-to-text generation, learns embedding in favor of text generation over contextual understanding. Moreover, our HTIE performs better than CLIP, manifesting that hyperbolic space is more suitable for text-image multi-modal learning.

Dual-branch Diffusion Model. Our HTIE and hyperbolic text-graph convolution module (HTGC) learn different text features. We compare two ways of utilizing these two text features in diffusion models, single-branch and dual-branch. The comparison results are shown in Group 2 of Tab. 3. The performance of the single-branch slightly decreases on IoU and F-score. The reason is that concatenating features of the single-branch lead to feature interference.

ence. In contrast, our dual-branch diffusion model further improves the performance of the model with HTIE, illustrating that our dual-branch architecture more effectively leverages the text features captured by HTIE and HTGC. Learned text features are visualized in the supplementary.

In addition, we compare our HTGC with the standard Euclidean graph convolution (GCN), as listed in Group 2 of Tab. 3. It can be observed that our method yields better results, demonstrating hyperbolic space is more suitable for learning the syntactic structure of texts.

Hyperbolic Hierarchical Loss. We compared the performance of our model with and without hyperbolic hierarchical loss. As shown in Group 3 of Tab. 3, the model with the hyperbolic hierarchical loss yields improvements in IoU and F-score while remaining comparable in CD, compared without the loss. It suggests that supervising the hierarchical structure of features during the denoising process contributes to improving the generation quality.

As listed in Group 3 of Tab. 3, we also compare our hyperbolic hierarchical loss with an Euclidean version computed by Euclidean distance, L_e . The results show that L_h is better, which reflects the advantage of hyperbolic space in maintaining the hierarchical structure.

5.3. Analysis of Hierarchical Learning

Analysis for Text-shape Hierarchical Structure. Text-shape exhibits hierarchical structure from general texts and detailed texts. Our method embeds the hierarchical structure by hyperbolic learning, like our HTIE and HTGC. We use HMD to qualitatively evaluate the ability to capture the text-shape hierarchical structure. As shown in Tab. 4, we enumerate the HMD of SDFusion [7] and our method on texts, ranging from general text to detailed text. We can observe that our method achieves smaller HMD between 3D shapes generated from different levels. It indicated that the shape generated from general text has hierarchical relationships with the shape generated from detailed texts. The hierarchical relationships are showcased in Fig. 5. Compared with SDFusion, given a general text, “a chair”, our method learned a shape without much detail. Indeed, a general text has no specific information, just like the root of the text-shape tree. Then adding the word “soft” to the text, the chair generated by our method looks softer. Finally, by adding more detailed words, like “four legs”, “new-fashion”, “two arms”, our method generated accurate shapes with natural wrappings. We also visualized 3D shape features generated from general and detailed texts in hyperbolic space. Dots in Fig. 6(b) represent 3D latent shape features generated by texts. The hierarchical distribution of these dots is consistent with the text hierarchy.

Analysis of the Hierarchical Structure of 3D Shape. We employ a hyperbolic hierarchical loss to regularize the feature space in the denoiser. We use HD to qualitatively eval-

Method	d_1	d_2	d_3	Order
SDFusion[7]	406.29	431.76	255.83	$d_2 > d_1 > d_3$
Ours	3.04	3.44	6.06	$d_3 > d_2 > d_1$

Table 5. Comparing the performance in maintaining the hierarchical structure of 3D shape.

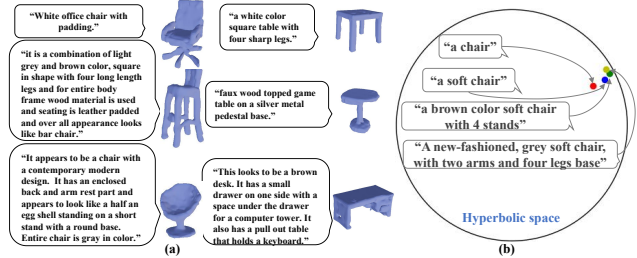


Figure 6. (a) More generation results, especially generated from long and complex text. (b) 3D Shape features generated from general and detailed texts in Poincaré Ball.

uate the ability to maintain the hierarchical structure of 3D shapes. As shown in Tab. 5, d_1 , d_2 , and d_3 are the distances of deep, middle, and shallow features to the origin in Euclidean space and hyperbolic space. If we compute the HD, the order provided by SDFusion is $d_2 > d_1 > d_3$, which does not follow the hierarchical structure of the point cloud. Instead, the order computed by our method is $d_3 > d_2 > d_1$, correctly following the hierarchy from deep to shallow features. It indicates the denoiser supervised by our hyperbolic hierarchical loss guarantees the tree-like hierarchical structure of 3D shape. We also provide the visualization of the feature distribution in the supplementary materials.

6. Conclusion

We propose a hyperbolic learning method for text-to-shape generation, namely HyperSDFusion. The key innovation lies in learning the inherent hierarchical structure of text and shape in hyperbolic space. In detail, we introduce a dual-branch diffusion model to fully utilize sequential and hierarchical features of text. The sequential features of the text are captured by the designed hyperbolic text-image encoder, simultaneously embedding multi-modal image-shape features. A hyperbolic text-graph convolution module is devised for learning hierarchical text features. Additionally, we propose a hyperbolic hierarchical loss to impart generated 3D shapes with hierarchical structure. Experimental results of our method on the Text2Shape dataset demonstrate the advantage of our method on text-to-shape generation. In the future, we will investigate more direct links to bridge language and 3D geometry.

Acknowledgments This work was supported by the National Nature Science Foundation of China under Grant 62272019, in part by China Scholarship Council. T. Birdal acknowledges support from the Engineering and Physical Sciences Research Council [grant EP/X011364/1].

References

- [1] Titas Anciukevičius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J Mitra, and Paul Guerrero. Rendering diffusion: Image diffusion for 3d reconstruction, inpainting and generation. In *CVPR*, pages 12608–12618, 2023. [1](#)
- [2] Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne Van Noord, and Pascal Mettes. Hyperbolic image segmentation. In *CVPR*, pages 4453–4462, 2022. [2](#), [3](#)
- [3] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2video: Video editing using image diffusion. In *ICCV*, pages 23206–23217, 2023. [2](#)
- [4] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [6](#)
- [5] Kevin Chen, Christopher B Choy, Manolis Savva, Angel X Chang, Thomas Funkhouser, and Silvio Savarese. Text2shape: Generating shapes from natural language by learning joint embeddings. In *ACCV*, 2018. [1](#), [2](#), [6](#)
- [6] Shoufa Chen, Peize Sun, Yibing Song, and Ping Luo. Diffusiondet: Diffusion model for object detection. In *ICCV*, pages 19830–19843, 2023. [2](#)
- [7] Yen-Chi Cheng, Hsin-Ying Lee, Sergey Tulyakov, Alexander G Schwing, and Liang-Yan Gui. Sdfusion: Multimodal 3d shape completion, reconstruction, and generation. In *CVPR*, pages 4456–4465, 2023. [1](#), [2](#), [6](#), [7](#), [8](#)
- [8] Noam Chomsky. *Syntactic structures*. Mouton de Gruyter, 2002. [1](#), [5](#)
- [9] Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, and Shanmukha Ramakrishna Vedantam. Hyperbolic image-text representations. In *International Conference on Machine Learning*, 2023. [2](#), [3](#), [5](#)
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [7](#)
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 34:8780–8794, 2021. [2](#)
- [12] Ondřej Dušek and Filip Jurcicek. Sequence-to-sequence generation for spoken dialogue via deep syntax trees and strings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, 2016. [1](#)
- [13] Aleksandr Ermolov, Leyla Mirvakhabova, Valentin Khrulkov, Nicu Sebe, and Ivan Oseledets. Hyperbolic vision transformers: Combining improvements in metric learning. In *CVPR*, pages 7409–7419, 2022. [2](#), [3](#)
- [14] Rao Fu, Xiao Zhan, Yiwen Chen, Daniel Ritchie, and Srinath Sridhar. Shapecrafter: A recursive text-conditioned 3d shape generation model. *NeurIPS*, 35:8882–8895, 2022. [2](#), [7](#)
- [15] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weillbach, and Frank Wood. Flexible diffusion modeling of long videos. *NeurIPS*, 35:27953–27965, 2022. [2](#)
- [16] Matthew Honnibal and Ines Montani. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. 2017. [5](#)
- [17] Ka-Hei Hui, Ruihui Li, Jingyu Hu, and Chi-Wing Fu. Neural wavelet-domain diffusion for 3d shape generation. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. [2](#)
- [18] Animesh Karnewar, Andrea Vedaldi, David Novotny, and Niloy J Mitra. Holodiffusion: Training a 3d diffusion model using 2d images. In *CVPR*, pages 18423–18433, 2023. [1](#)
- [19] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *ICCV*, pages 2151–2162, 2023. [2](#)
- [20] Valentin Khrulkov, Leyla Mirvakhabova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic image embeddings. In *CVPR*, pages 6418–6428, 2020. [2](#), [3](#)
- [21] Gwanghyun Kim, Ji Ha Jang, and Se Young Chun. Podia-3d: Domain adaptation of 3d generative model across large domain gap using pose-preserved text-to-image diffusion. In *ICCV*, pages 22603–22612, 2023. [2](#)
- [22] Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in pytorch. *arXiv preprint arXiv:2005.02819*, 2020. [3](#)
- [23] Zhiying Leng, Shun-Cheng Wu, Mahdi Saleh, Antonio Montanaro, Hao Yu, Yin Wang, Nassir Navab, Xiaohui Liang, and Federico Tombari. Dynamic hyperbolic attention network for fine hand-object reconstruction. In *ICCV*, pages 14894–14904, 2023. [3](#)
- [24] Lingxiao Li, Yi Zhang, and Shuhui Wang. The euclidean space is evil: Hyperbolic attribute editing for few-shot image generation. In *ICCV*, pages 22714–22724, 2023. [4](#)
- [25] Muheng Li, Yueqi Duan, Jie Zhou, and Jiwen Lu. Diffusion-sdf: Text-to-shape via voxelized diffusion. In *CVPR*, pages 12642–12651, 2023. [2](#), [6](#), [7](#)
- [26] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. [1](#), [2](#)
- [27] Fangzhou Lin, Yun Yue, Songlin Hou, Xuechu Yu, Yajun Xu, Kazunori D Yamada, and Ziming Zhang. Hyperbolic chamfer distance for point cloud completion. In *ICCV*, pages 14595–14606, 2023. [2](#), [3](#)
- [28] Zhengzhe Liu, Yi Wang, Xiaojuan Qi, and Chi-Wing Fu. Towards implicit text-guided 3d shape generation. In *CVPR*, pages 17896–17906, 2022. [2](#), [6](#)
- [29] Simian Luo, Xuelin Qian, Yanwei Fu, Yinda Zhang, Ying Tai, Zhenyu Zhang, Chengjie Wang, and Xiangyang Xue. Learning versatile 3d shape generation with improved autoregressive models. In *ICCV*, pages 14139–14149, 2023. [2](#)
- [30] Pascal Mettes, Mina Ghadimi Atigh, Martin Keller-Ressel, Jeffrey Gu, and Serena Yeung. Hyperbolic deep learning in computer vision: A survey. *arXiv preprint arXiv:2305.06611*, 2023. [2](#), [3](#)
- [31] Paritosh Mittal, Yen-Chi Cheng, Maneesh Singh, and Shubham Tulsiani. Autosdf: Shape priors for 3d completion, reconstruction and generation. In *CVPR*, 2022. [1](#), [2](#), [7](#)
- [32] Antonio Montanaro, Diego Valsesia, and Enrico Magli. Re-thinking the compositionality of point clouds through regularization in the hyperbolic space. *NeurIPS*, 35:33741–33753, 2022. [1](#), [2](#), [3](#)

- [33] Muhammad Ferjad Naeem, Evin Pınar Örnek, Yongqin Xian, Luc Van Gool, and Federico Tombari. 3d compositional zero-shot learning with decompositional consensus. In *ECCV*, pages 713–730. Springer, 2022. 1
- [34] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, pages 4195–4205, 2023. 2
- [35] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. Hyperbolic deep neural networks: A survey. *IEEE TPAMI*, 44(12):10023–10044, 2021. 3
- [36] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 1, 2
- [37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 7
- [38] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 7
- [39] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, Yuanzhen Li, and Varun Jampani. Dreambooth3d: Subject-driven text-to-3d generation. In *ICCV*, pages 2349–2359, 2023. 2
- [40] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 2, 4
- [41] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. Clip-forge: Towards zero-shot text-to-shape generation. In *CVPR*, pages 18603–18613, 2022. 2
- [42] Aditya Sanghi, Rao Fu, Vivian Liu, Karl DD Willis, Hooman Shayani, Amir H Khasahmadi, Srinath Sridhar, and Daniel Ritchie. Clip-sculptor: Zero-shot generation of high-fidelity and diverse shapes from natural language. In *CVPR*, pages 18339–18348, 2023. 2
- [43] Jaehyeok Shim, Changwoo Kang, and Kyungdon Joo. Diffusion-based signed distance fields for 3d shape generation. In *CVPR*, pages 20887–20897, 2023. 2
- [44] Habib Slim, Xiang Li, Yuchen Li, Mahmoud Ahmed, Mohamed Ayman, Ujjwal Upadhyay, Ahmed Abdelreheem, Arpit Prajapati, Suhail Pothigara, Peter Wonka, et al. 3dcompat++: An improved large-scale 3d vision dataset for compositional recognition. *arXiv preprint arXiv:2310.18511*, 2023. 1
- [45] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 6
- [46] Maxim Tatarchenko, Stephan R Richter, René Ranftl, Zhuwen Li, Vladlen Koltun, and Thomas Brox. What do single-view 3d reconstruction networks learn? In *CVPR*, pages 3405–3414, 2019. 6
- [47] Xi Tian, Yong-Liang Yang, and Qi Wu. Shapescollider: Structure-aware 3d shape generation from text. In *ICCV*, pages 2715–2724, 2023. 2
- [48] Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, Karsten Kreis, et al. Lion: Latent point diffusion models for 3d shape generation. *Advances in Neural Information Processing Systems*, 35:10021–10039, 2022. 2
- [49] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *NeurIPS*, 30, 2017. 4
- [50] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, pages 12619–12629, 2023. 2
- [51] Qiucheng Wu, Yujian Liu, Handong Zhao, Trung Bui, Zhe Lin, Yang Zhang, and Shiyu Chang. Harnessing the spatial-temporal attention of diffusion models for high-fidelity text-to-image synthesis. In *ICCV*, pages 7766–7776, 2023. 2
- [52] Jiale Xu, Xintao Wang, Weihao Cheng, Yan-Pei Cao, Ying Shan, Xiaohu Qie, and Shenghua Gao. Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In *CVPR*, pages 20908–20918, 2023. 2
- [53] Menglin Yang, Min Zhou, Zhihao Li, Jiahong Liu, Lujia Pan, Hui Xiong, and Irwin King. Hyperbolic graph neural networks: a review of methods and applications. *arXiv preprint arXiv:2202.13852*, 2022. 3, 5
- [54] Ye Yuan, Jiaming Song, Umar Iqbal, Arash Vahdat, and Jan Kautz. Physdiff: Physics-guided human motion diffusion model. In *ICCV*, pages 16010–16021, 2023. 2
- [55] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023. 2
- [56] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 2
- [57] Xinyuan Zhang, Yi Yang, Siyang Yuan, Dinghan Shen, and Lawrence Carin. Syntax-infused variational autoencoder for text generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2069–2078, 2019. 1
- [58] Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. Locally attentional sdf diffusion for controllable 3d shape generation. *ACM Transactions on Graphics (TOG)*, 42(4):1–13, 2023. 2
- [59] Min Zhou, Menglin Yang, Bo Xiong, Hui Xiong, and Irwin King. Hyperbolic graph neural networks: A tutorial on methods and applications. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 5843–5844, 2023. 3
- [60] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *CVPR*, pages 12588–12597, 2023. 1