



SketchDream: Sketch-based Text-to-3D Generation and Editing

FENG-LIN LIU, Institute of Computing Technology, CAS, China and University of Chinese Academy of Sciences, China

HONGBO FU, SCM, City University of Hong Kong, China and EMIA, HKUST, China

YU-KUN LAI, School of Computer Science and Informatics, Cardiff University, UK

LIN GAO*, Institute of Computing Technology, CAS, China and University of Chinese Academy of Sciences, China



Fig. 1. Our SketchDream system supports both generation and editing of high-quality 3D contents from 2D sketches. As shown in (a), given hand-drawn sketches and text prompts (on top of each example), our method generates high-quality rendering results of 3D contents from scratch. Existing text-to-3D generation approaches like MVDream [Shi et al. 2023b] generate photo-realistic results but cannot control component layouts and details, such as the door and window in the top example and the pose in the bottom example. In (b), we show sketch-based editing results of NeRFs reconstructed from real models. The newly generated components naturally interact with the original objects, with the unedited regions well preserved.

Existing text-based 3D generation methods generate attractive results but lack detailed geometry control. Sketches, known for their conciseness and expressiveness, have contributed to intuitive 3D modeling but are confined to producing texture-less mesh models within predefined categories. Integrating sketch and text simultaneously for 3D generation promises enhanced control over geometry and appearance but faces challenges from 2D-to-3D translation ambiguity and multi-modal condition integration. Moreover, further editing of 3D models in arbitrary views will give users more freedom to customize their models. However, it is difficult to achieve high generation quality, preserve unedited regions, and manage proper interactions between shape components. To solve the above issues, we propose a text-driven 3D content generation and editing method, SketchDream, which supports NeRF generation from given hand-drawn sketches and achieves free-view sketch-based local editing. To tackle the 2D-to-3D ambiguity challenge, we introduce

a sketch-based multi-view image generation diffusion model, which leverages depth guidance to establish spatial correspondence. A 3D ControlNet with a 3D attention module is utilized to control multi-view images and ensure their 3D consistency. To support local editing, we further propose a coarse-to-fine editing approach: the coarse phase analyzes component interactions and provides 3D masks to label edited regions, while the fine stage generates realistic results with refined details by local enhancement. Extensive experiments validate that our method generates higher-quality results compared with a combination of 2D ControlNet and image-to-3D generation techniques and achieves detailed control compared with existing diffusion-based 3D editing approaches.

CCS Concepts: • **Human-centered computing** → **Graphical user interfaces**; • **Computer systems organization** → **Neural networks**; • **Computing methodologies** → **Rendering**; Volumetric models.

Additional Key Words and Phrases: sketch-based interaction, diffusion models, neural radiance fields, 3D generation

ACM Reference Format:

Feng-Lin Liu, Hongbo Fu, Yu-Kun Lai, and Lin Gao. 2024. SketchDream: Sketch-based Text-to-3D Generation and Editing. *ACM Trans. Graph.* 43, 4, Article 44 (July 2024), 13 pages. <https://doi.org/10.1145/3658120>

1 INTRODUCTION

Creating high-quality 3D content is a popular topic with wide applicability in VR/AR, the movie industry, architecture, robotics simulation, etc. However, traditional 3D content production depends on sophisticated software and laborious procedures, making it challenging for amateur users to design their 3D models. To solve this

*Corresponding author: Lin Gao (gaolin@ict.ac.cn).

Authors' Contact Information: Feng-Lin Liu, Institute of Computing Technology, CAS, China and University of Chinese Academy of Sciences, China; Hongbo Fu, SCM, City University of Hong Kong, China and EMIA, HKUST, China, fuplus@gmail.com; Yu-Kun Lai, School of Computer Science and Informatics, Cardiff University, UK, LaiY4@cardiff.ac.uk; Lin Gao, Institute of Computing Technology, CAS, China and University of Chinese Academy of Sciences, China, gaolin@ict.ac.cn.



This work is licensed under a Creative Commons Attribution International 4.0 License. © 2024 Copyright held by the owner/author(s). ACM 1557-7368/2024/7-ART44 <https://doi.org/10.1145/3658120>

problem, sketching, a user-friendly and expressive interaction tool, has been utilized in 2D [Chen et al. 2009; Isola et al. 2017; Zhang et al. 2023] and 3D content generation [Gao et al. 2022, 2023b; Xiang et al. 2020; Zhang et al. 2021a; Zheng et al. 2023]. However, due to the restriction of network capacity and the ambiguity of 2D sketches, existing works can only generate limited single or multiple categories of objects. Additionally, since sketches only contain geometry information, how to effectively control the appearance has been under-explored in existing sketch-based 3D generation works.

Compared with sketches, text prompts can describe object category and appearance more easily. Thanks to the development of diffusion models, text-to-image generation [Midjournal. 2022; Rombach et al. 2022] has become successful in recent years. Based on these pre-trained 2D models, DreamFusion [Poole et al. 2023] designs Score Distillation Sampling (SDS) to optimize Neural Radiance Fields (NeRFs) and presents a text-to-3D generation framework. Follow-up works [Chen et al. 2023a; Lin et al. 2023; Qiu et al. 2023; Shi et al. 2023b; Wang et al. 2023c] further modify the supervision losses and optimization process to generate more realistic and higher-quality 3D models. Despite impressive results of text-based 3D generation, text prompts cannot precisely depict objects' shape, texture patterns, and layouts, and are thus less suitable for fine-grained control. Moreover, it is difficult for the above works to edit the local details of generated or real 3D models because of their global control of textual description. Although some works [Cheng et al. 2023; Sella et al. 2023; Zhuang et al. 2023] further achieve text-based local editing by labeling editing regions, it is still hard for them to easily control the shape and position of edited components.

To improve the controllability of 3D generation, we introduce sketch into the text-to-3D generation framework: sketch guides the shape and pattern, while text prompt controls the material and appearance. This goal is challenging to achieve for the following reasons. First, since a sketch only contains single-view information, the absence of multi-view supervision makes it difficult to generate complete 3D models. Moreover, directly adding single-view sketch constraints in other viewpoints degrades the generation quality and has lower faithfulness. For example, a straightforward approach [Zeng et al. 2023] is to first generate 2D images by ControlNet [Zhang et al. 2023] with a sketch-based condition and then utilize image-to-3D approaches [Liu et al. 2023b; Qian et al. 2023; Sun et al. 2023] to generate 3D contents. However, as shown in Fig. 8, this approach tends to generate distorted geometry in novel views different from input sketches and fuzzy texture details in the back view due to the information absence.

Further editing of the generated contents in other views provides users with more detailed control of results. Additionally, editing existing real models instead of generation from scratch is also a common situation. However, sketch-based 3D content editing is challenging because besides achieving high-quality content creation, the edited content should reasonably interact with the original content, with unedited regions well preserved. To solve the above issue, SKED [Mikaeili et al. 2023] requires the input of multi-view sketches, which are treated as binary masks to label edited regions. Although it achieves effective text- and sketch-based editing at the object level, maintaining unedited regions well and supporting part-level editing are still challenging.

To address the above issues, we propose SketchDream, a method for sketch-based text-to-3D generation and editing of photo-realistic contents. For sketch-based generation, since sketches are sparse and lack 3D information, it is ambiguous and difficult to directly generate 3D contents from them. To solve this problem, we use the depth information, which can bridge the 2D inputs and 3D models. Specifically, given an input 2D sketch and a text prompt, we utilize a 2D diffusion model [Zhang et al. 2023] to generate a corresponding depth map, which is then utilized to warp the input sketch. To propagate the sketch into 3D space and avoid the Janus problem (i.e., 3D models with multiple frontal faces), we build a 3D ControlNet based on MVDream [Shi et al. 2023b] to generate images in four camera views uniformly distributed in azimuth. To build the correspondence between the sketch and multi-view images, we modify the MVDream backbone to generate an additional image in the view of the input sketch. To ensure the consistency between different views and support sketch control, we design a 3D-attention control module, which takes the input 2D sketch and warped sketch in the nearby view with depth guidance as inputs to effectively control the 3D diffusion generation. Similar to MVDream [Shi et al. 2023b], we use SDS-based optimization to generate high-quality 3D contents.

Our framework naturally supports sketch-based 3D content editing. Given the generated or real 3D models, users can input an edited sketch with a 2D editing mask to modify local components. To generate high-quality editing results, we design a two-stage coarse-to-fine editing framework. In the coarse stage, the 2D mask is lifted into 3D space to construct a coarse 3D columnar mask. Then, the coarse stage generates an editing result that generally conforms to the edited sketch and the input text, but might have unoptimized quality, inadequate faithfulness, and mistaken interaction with the original object. We optimize the editing result quality in the fine stage. In particular, we extract a mesh model from the coarse-stage editing result and then label the edited region by the coarse 3D columnar mask. This new 3D mask precisely represents the component interaction and helps preserve the unedited regions. Moreover, we propose a local rendering strategy that calculates the sketch-based SDS supervision locally to enhance the sketch faithfulness and generation quality further.

Extensive experiments validate that our method generates higher-quality results than possible sketch-based text-to-3D baselines and existing sketch-based 3D editing approaches. Our main contributions can be summarized as follows.

- We propose the first sketch-based text-to-3D generation and editing method, which generates high-quality 3D objects under generalized categories and supports detailed editing of reconstructed or generated NeRFs.
- We propose a sketch-based multi-view image generation diffusion model, which utilizes a depth-guided warping strategy to create spatial correspondence, and a 3D-attention control module to ensure 3D consistency.
- To support local modification further, we develop a coarse-to-fine editing framework. The coarse stage generates initial results to label edited regions better, while the fine stage

generates high-quality editing results with a local rendering strategy.

2 RELATED WORK

Sketch-based 3D Generation. Sketch-based 3D generation has been extensively researched. Early works utilize retrieval strategies [Chen et al. 2003; Funkhouser et al. 2003] or carefully designed mapping approaches [Igarashi et al. 2006; Zeleznik et al. 2006] to determine 3D shapes by sketches. With the development of deep learning, recent works treat this problem as sketch-conditioned 3D reconstruction and predict volumetric grids, point clouds, mesh models, or even CAD commands [Li et al. 2020, 2022]. Although volumetric prediction methods [Delanoy et al. 2018; Wang et al. 2018] generate 3D models faithful to sketches, the grid resolution restricts the model quality. To enhance the details, another category of works [Gao et al. 2022; Lun et al. 2017; Wang et al. 2022b] map sketches into point clouds. To further construct 3D shape meshes, other works [Li et al. 2018; Xiang et al. 2020; Zhong et al. 2020b] predict the depth and normal maps, which are utilized to deform pre-defined templates or directly construct 3D shapes. Data augmentation [Zhong et al. 2020a] and sketch preprocessing [Zhang et al. 2021a] are also utilized to improve the robustness of generation from hand-drawn sketches. Despite the successful generation results, the above approaches are hard to achieve detailed local control for out-of-domain examples. Zheng et al. [2023] propose a diffusion-based tri-plane generation framework to generate SDF models.

Different from the above works, our method utilizes NeRF as a 3D representation and thus synthesizes 3D shapes with realistic textures instead of geometry models only. Besides, our method is not restricted to single or multiple object categories and aims for a more general framework for sketch-based 3D model generation.

Text- and Image-based 3D Generation. With the development of diffusion models, text-based 3D content generation has become popular in recent years. One category of methods (e.g., [Wang et al. 2023d]) utilizes diffusion models to generate 3D representations like tri-plane features directly, but their performance and generalization are limited by 3D training datasets. Another category of methods employs pre-trained 2D diffusion models [Rombach et al. 2022] to optimize 3D representations by the SDS loss [Poole et al. 2023] or SJC (Score Jacobian Chaining) loss [Wang et al. 2023a]. The subsequent works try to improve the performance based on the above framework. Some works divide the generation process into two stages: geometry optimization and texture optimization, and utilize diverse 3D representations like DM Tet [Lin et al. 2023], BRDF [Chen et al. 2023a], and Gaussian Splatting [Tang et al. 2023a]. Other works modify the supervision during generation by designing, for example, latent space optimization [Metzer et al. 2023], VSD (Variational Score Distillation) loss [Wang et al. 2023c], interval score matching [Liang et al. 2023], and normal-depth supervision [Qiu et al. 2023]. MVDream [Shi et al. 2023b] further proposes a four-view generation model and solves the Janus problem. Please refer to insightful surveys [Liu et al. 2024; Sun et al. 2024; Xia and Xue 2024] for comprehensive understanding of the 3D generation approaches. Although the above works generate high-quality results, users cannot precisely control the geometry shapes and texture details. Compared

with the above works, our method adds easily drawn sketches as an additional condition and achieves more detailed control during the text-based generation.

To improve the controllability of 3D generation, many works utilize single-view images to replace or combine with text prompts as input. Make-It-3D [Tang et al. 2023b] and RealFusion [Melas-Kyriazi et al. 2023] add image constraints to 3D generation, but still fail to generate optimal results in the back view. Follow-up works [Liu et al. 2023b,a; Shi et al. 2023a] modify Stable Diffusion [Rombach et al. 2022] to infer other views of a single-view image. Their pre-trained models are utilized in Magic123 [Qian et al. 2023] to achieve better geometry results in novel views. DreamCraft3D [Sun et al. 2023] and HyperDreamer [Wu et al. 2023] respectively introduce DreamBooth [Ruiz et al. 2023] and image super-resolution to enhance the texture quality. Concurrently, ImageDream [Wang and Shi 2023] adds image condition into a multi-view generation model to improve the generation robustness. However, in real applications, obtaining a desirable 2D image with detailed control is nontrivial. Additionally, it is difficult for the above works to support local modification because of the global dependence on the input images. Apart from real images, Control3D [Chen et al. 2023b] utilizes sketches as conditions for diffusion-based 3D generation but tends to generate fuzzy details due to the dependence of 2D ControlNet. A concurrent work of our method, MVControl [Li et al. 2023] extracts edge maps from images as an additional condition, but do not deal with hand-drawn sketches and local editing. In contrast, our method liberates the reliance on images, enabling high-quality 3D generation from scratch, and additionally facilitates sketch-based local editing.

3D Content Editing. Compared with 3D generation from scratch, 3D content editing requires additional considerations on the relationship between the edited and original contents and the preservation of unedited regions. Previous methods only support geometry editing [Gao et al. 2023b; Garbin et al. 2022; Xu and Harada 2022; Yuan et al. 2022], require laborious interaction operations [Liu et al. 2021; Yang et al. 2022, 2021], or only support global manipulation [Liu et al. 2023c; Nguyen-Phuoc et al. 2022; Zhang et al. 2021b]. To achieve more detailed and semantic-driven editing, many approaches [Gao et al. 2023a; Wang et al. 2022a, 2023b] utilize the pre-trained CLIP model [Radford et al. 2021] to compute the text guidance loss, which limits their performance because of CLIP’s training for text-image alignment instead of generation. To further improve the editing performance, the subsequent works utilize Stable Diffusion for 3D editing. These approaches utilize a user-defined 3D geometry box [Cheng et al. 2023; Li et al. 2024] or an inferred 3D mask by multi-modal attention [Sella et al. 2023; Zhuang et al. 2023] to label local editing regions, supporting editing manipulations like adding or replacing objects while preserving unedited regions in scenes. However, it is still difficult to control the shape or position of the edited content precisely by using text prompts. Most similar to our work, SKED [Mikaeili et al. 2023] also supports sketch-based control of 3D editing results. However, SKED requires multi-view sketches as input and cannot control texture details because of the transformation of sketches into texture-less masks. Compared with existing works, our method supports single-view sketch-based editing and achieves

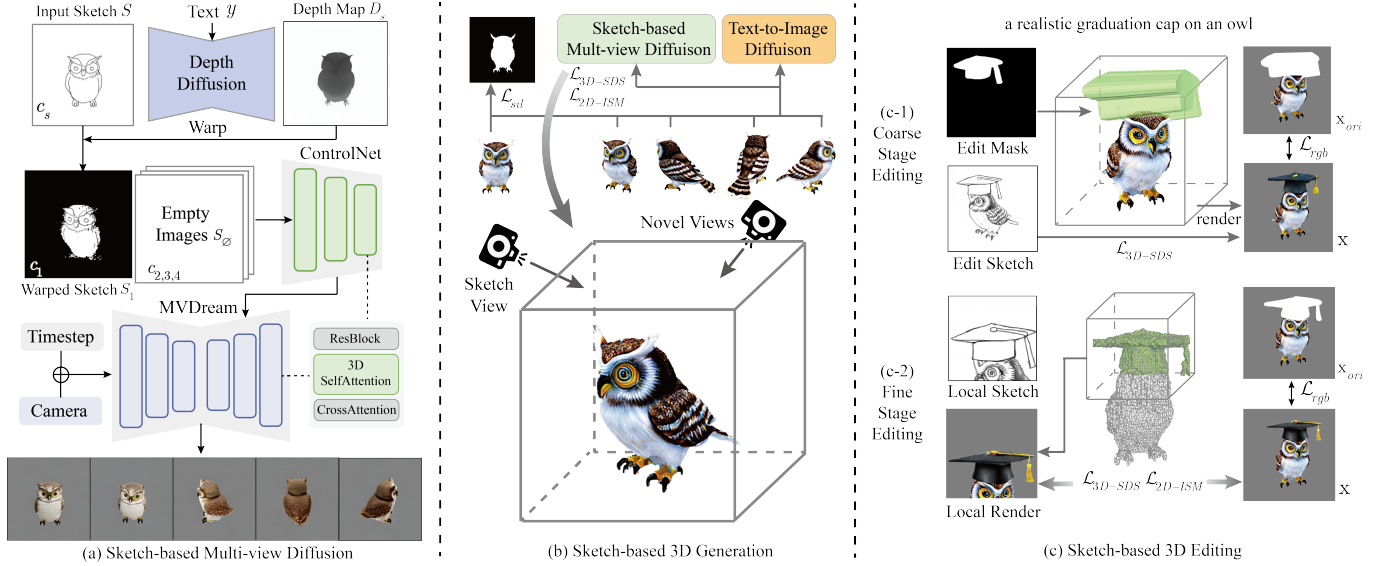


Fig. 2. The overview of our SketchDream for sketch-based generation and editing. Given an input sketch S and a text prompt y , we design a sketch-based multi-view diffusion model (a), which takes S , depth-warped sketch S_1 , and white images S_0 as conditions and generates multi-view images in the sketch view c_s and novel views c_{mv} . In order to generate realistic 3D contents (b), we render images under five views corresponding to those in the multi-view diffusion model and then optimize a NeRF by 3D score distillation and 2D interval score matching. For sketch-based 3D editing (c), we design a two-stage editing framework. In the coarse stage, we build a coarse 3D mask and generates a coarse editing result, which is used to get precise 3D masks for high-quality local editing in the fine stage.

more complicated editing operations like controllable geometry modification.

3 PRELIMINARY

MVDream. To avoid the Janus problem and generate correct geometry, we build our framework on MVDream [Shi et al. 2023b]. With text inputs, MVDream simultaneously generates images under four views, which have uniformly distributed view angles, with the same elevation. To control view angles, the absolute camera extrinsic matrix is encoded and added to the time-step embedding in a UNet. To ensure cross-view consistency, MVDream utilizes a 3D attention module that shares the queries Q , keys K , and values V in all views. The diffusion model is fine-tuned from 2D Stable Diffusion on the Objaverse [Deitke et al. 2023] dataset.

Score Distillation Sampling (SDS). First proposed in DreamFusion [Poole et al. 2023], SDS has been widely used in text-to-3D generation based on 2D diffusion models. To mitigate the color saturation problem, we utilize an SDS version with x_0 -reconstruction loss:

$$\mathcal{L}_{SDS}(\theta, x = g(\theta, c)) = \mathbb{E}_{t,c,\epsilon} [\|x - \hat{x}_0\|_2^2], \quad (1)$$

where $g(\theta, c)$ is a rendered image of 3D representation θ with camera condition c . \hat{x}_0 is the estimated x_0 (images without noise) based on the UNet's output $\epsilon_\phi(x_t; y, c, t)$, where y is a text condition, x_t is the diffusion forward process results at time step t with Gaussian noise ϵ . In each optimization step, t and ϵ are randomly sampled, constraining the rendered images x to satisfy the distribution of pre-trained diffusion models ϕ . More details can be found in [Shi et al. 2023b].

4 METHODOLOGY

In this section, we present our framework for sketch-based text-to-3D content generation and editing, as illustrated in Fig. 2. Our approach addresses the challenge of synthesizing realistic 3D contents from sparse 2D sketches and textual inputs by completing missing appearance details and extending single-view information into the 3D space. To accomplish this, we introduce a sketch-based multi-view diffusion model in Sec. 4.1. This model predicts depth maps for sketch warping to establish spatial correspondence and generates realistic multi-view images. We detail our sketch-based 3D content generation to achieve a seamless 3D representation in Sec. 4.2. The sketch-based multi-view diffusion model collaborates with SDS to apply 3D constraints and sketch control. Simultaneously, a pre-trained 2D text-to-image diffusion model enhances appearance details. Sec. 4.3 delves into our sketch-based editing framework, showcasing its ability to facilitate effective local editing while preserving the original features in unedited regions.

4.1 Sketch-based Multi-View Diffusion Model

To achieve sketch-based 3D generation, the missing texture and material should be added, and the geometry information contained in 2D sketches should be propagated into 3D space. To achieve these goals, we design a sketch-based multi-view image generation model, which is built on the pre-trained MVDream [Shi et al. 2023b] because of its powerful 3D understanding ability. To achieve sketch control, we modify the MVDream backbone to generate four novel-view images with an additional image in the sketch view. Formally, given an input sketch S with a corresponding camera condition

c_s and a text condition y , our multi-view diffusion model generates a sketch-view image x_s and novel-view images $\{x_1, x_2, x_3, x_4\}$ corresponding to novel-view camera conditions $\{c_1, c_2, c_3, c_4\}$. We design a depth-guided warping strategy to explicitly build the spatial correspondence and utilize a 3D attention module to ensure 3D consistency.

Depth-guided Warping. Adding an additional control like sketch into pre-trained 2D diffusion models has become successful in previous works [Mou et al. 2024; Zhang et al. 2023]. However, different from 2D image translation with precise pixel alignment, sketch-based novel view image synthesis requires a complicated understanding of 3D geometry to generate correct results. We observe that depth maps can serve as an intermediate geometry representation to solve the sketch’s ambiguity and improve the faithfulness of sketches to generated models.

We build a depth map generation diffusion model, which takes the sketch S and text y as input and generates the corresponding depth map D_s . Since the success of 2D ControlNet comes from the spatial alignment, we build this correspondence in 3D space by depth-guided pose warping [Fehn 2004; Somraj 2020]. Specifically, given the source depth D_s and camera parameters c_s , we reproject the pixels of the sketch image into 3D space. Instead of warping the input sketch S to all four novel views, we only warp it to the nearest view c_1 because the quality of the warped sketch in a far viewpoint can degrade significantly. The warped target sketch S_1 is generated by interpolating the source pixel values. Formally, we denote the depth-guidance warping process as:

$$S_1 = \text{Warp}(S, D_s, c_s, c_1). \quad (2)$$

3D Attention Control Module. We design a specific 3D attention control module to apply ControlNet [Zhang et al. 2023] to the pre-trained MVDream. In the sketch view c_s , the input of the control branch is the sketch S itself. For novel view images, the input in the nearest view c_1 is the warped sketch S_1 , while the inputs of the other views $\{c_2, c_3, c_4\}$ are empty images S_\emptyset . To ensure 3D consistency, the 3D ControlNet utilizes the 3D attention module [Shi et al. 2023b]. In this module, as discussed in Sec. 3, the Q, K, and V features in the self-attention layers are shared across multiple views, thus generating consistent multi-view images. It should be noticed that although the inputs of the ControlNet in far views are empty, the 3D attention shares the sketch condition information to achieve effective control.

Training. We trained the depth generation model and the sketch-based multi-view generation model separately. The former was finetuned based on the 2D ControlNet. During the training, we fixed the VAE and ControlNet branch and trained the text-to-image UNet. We treated the depth map as a color image and set the resolution as 256×256 . To train the networks, we utilize a subset of Objaverse [Deitke et al. 2023] containing 150k 3D objects with 24 rendered images for each object, wherein 8 images with random viewpoints serve as input views and the other 16 images serve as target images. For the images of input views, we extract paired sketches by the method in [Chan et al. 2022].

To train the multi-view ControlNet, we utilized the zero convolution to initialize the newly added layers as in the 2D ControlNet. Given the dataset of paired multi-view images and sketches, the

training samples $\{x, s, y, c\}$ contain images $x = \{x_s, x_1, x_2, x_3, x_4\}$, sketch conditions $s = \{S, S_1, S_\emptyset\}$, text conditions y , and camera conditions $c = \{c_s, c_1, c_2, c_3, c_4\}$. We define the controllable multi-view diffusion loss as:

$$\mathcal{L}_{MV-ctrl}(\phi) = \mathbb{E}_{x,t,y,s,c,\epsilon} [\|\epsilon - \epsilon_\phi(x_t; t, y, c, s)\|_2^2], \quad (3)$$

where t is a time step and x_t is the noisy image generated from x .

4.2 Sketch-based 3D Generation

Since the generated 4-view images in Sec. 4.1 are view-sparse and not strictly 3D consistent, it is hard to directly optimize NeRF. Additionally, randomly sampling the 4 orthogonal views N times generates $4 \times N$ images, which, however, might have different geometry, color, and texture, thus lacking 3D coherence. Therefore, we utilize SDS optimization to generate 3D contents. Specifically, we render five view images x : the sketch view image to control the geometry and four randomly sampled view images to optimize the 3D NeRF representation. We calculate the 3D SDS \mathcal{L}_{SDS}^{3D} as defined in Equation (1) while utilizing the sketch-based diffusion network $\epsilon_\phi(x_t; t, y, c, s)$. Similar to MVDream [Shi et al. 2023b], we use the classifier free guidance (CFG) rescale trick to mitigate the color over-saturation.

Apart from the 3D SDS constraint, we also utilize the 2D pre-trained text-to-image model (Stable Diffusion [Rombach et al. 2022]). We use the Interval Score Matching (ISM) [Liang et al. 2023] loss, which is more robust and generates more realistic results in our framework than the original SDS. To further improve the sketch faithfulness, we apply a 2D silhouette loss in the sketch view:

$$\mathcal{L}_{sil} = \|M_s - C_s^\alpha\|_2^2, \quad (4)$$

where C_s^α is the rendered object alpha mask by 3D NeRF, and M_s is the sketched region obtained by background removal. The overall objective of our 3D generation process is:

$$\mathcal{L}_{total}(\theta) = \lambda_1 \mathcal{L}_{SDS}^{3D} + \lambda_2 \mathcal{L}_{ISM}^{2D} + \lambda_3 \mathcal{L}_{sil} + \lambda_4 \mathcal{L}_{orient}, \quad (5)$$

where \mathcal{L}_{orient} is the regular orientation loss proposed by [Poole et al. 2023]. We also turn on the point lighting [Poole et al. 2023] and soft shading [Lin et al. 2023] as in MVDream to smooth the geometry.

4.3 Sketch-based 3D Editing

Our framework supports sketch-based local editing of the generated or reconstructed NeRFs of 3D models. To achieve detailed control, users are provided with sketches [Chan et al. 2022] synthesized from rendered images. Subsequently, they can modify these sketches and draw an additional mask to label edited regions. Since it is challenging to directly infer the interaction between object components and generate high-quality results, we design a two-stage editing strategy. In the coarse stage, we get an initial editing result, which is used to predict a detailed 3D mask representing the component interaction. In the fine stage, the framework generates realistic, high-quality editing results while precisely preserving the unedited regions.

4.3.1 Coarse-Stage Editing. To enable effective 3D local editing, our approach involves the transformation of a 2D mask into 3D space to label editing regions in novel views. In the coarse stage, we utilize a cylinder mesh model for this purpose. Beginning with

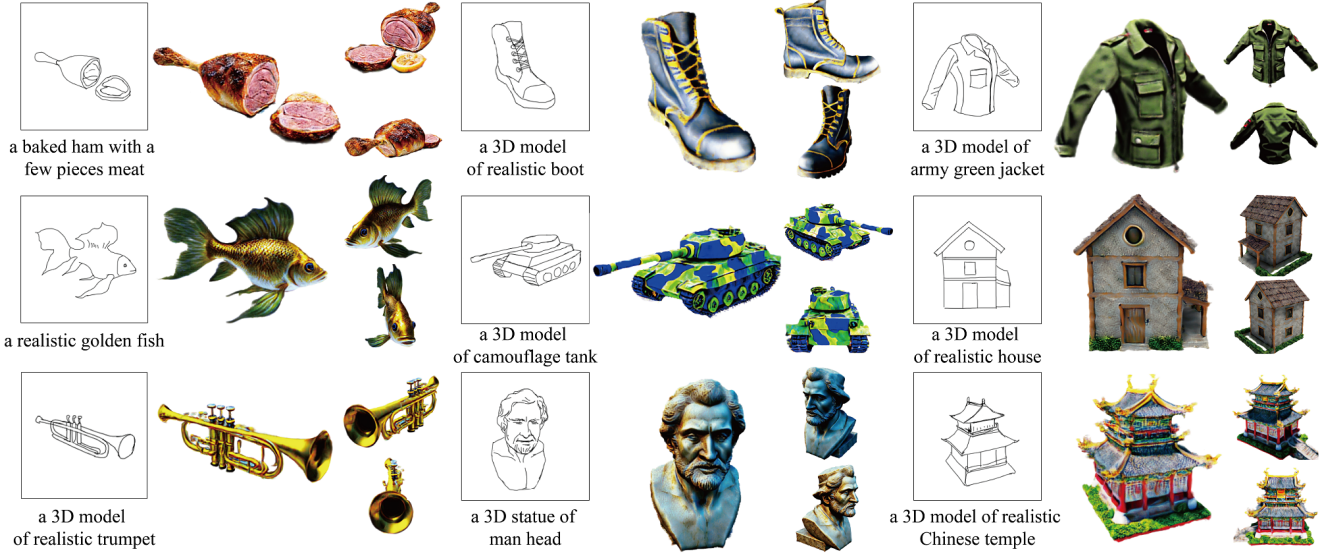


Fig. 3. Sketch-based generation results. Given hand-drawn sketches, our method generates high-quality 3D results, which are faithful to the input sketches and texts. Our method can generate models under diverse categories, including clothes, food, animals, humanoid objects, etc. It can be seen that the shape and pattern details can be controlled by sketches.




Fig. 4. Sketch-based multi-view image generation results. Given hand-drawn sketches (a) and texts shown above the images, our method generates realistic multi-view images (b), which are faithful to the input sketches and text prompts.

a hand-drawn 2D mask, users define the minimum and maximum depth values, which are utilized to construct the top and bottom surfaces of the cylinder. Similar to sketch-based generation in Sec. 4.2, we render five view images x to optimize the NeRF. To maintain the unedited regions, we render the 3D mask model in the camera conditions c to generate M_{2D} . We define an image loss to preserve the color features:

$$\mathcal{L}_{rgb} = \left\| x \odot \overline{M}_{2D} - x_{ori} \odot \overline{M}_{2D} \right\|_2^2, \quad (6)$$




Fig. 5. Sketch-based generation results with different text prompts. Our method generates diverse and realistic results, whose geometry is controlled by the input sketch while appearance being controlled by the text prompts.

where x_{ori} are the multi-view rendered images before editing. Since we only need to get coarse editing results in this stage, the 2D loss is not used. The overall objective is:

$$\mathcal{L}_{total}^{coarse}(\theta) = \alpha_1 \mathcal{L}_{SDS}^{3D} + \alpha_2 \mathcal{L}_{rgb} + \alpha_3 \mathcal{L}_{sil} + \alpha_4 \mathcal{L}_{orient}. \quad (7)$$

4.3.2 Fine-Stage Editing. Utilizing the coarse 3D mask generates an initial editing result, but some undesirable regions are also mistakenly included or undesirably changed, as shown in Fig. 12. To address this issue, we get a more precise 3D mask to improve the editing performance based on the results of the coarse stage. Since NeRF lacks segmentation information and is computationally expensive, we translate the coarse NeRF results into a mesh model. This 3D model has the newly generated components and generally maintains the original geometry in the unedited regions. Then, we label local regions of the 3D mesh to represent the edited regions by setting the vertices within the coarse mesh as the editing vertices, followed by manual refinement. It is convenient to paint on the mesh model by modifying vertex colors. In the case of editing/removing existing components, we extract the mesh model from the original




Fig. 6. Sketch-based editing results. Given real 3D models, users can select arbitrary views to render images and edit the local regions by inputting texts, modifying sketches (extracted from original render images), and drawing masks. Our method supports the change of local components of real models, such as changing the lion’s head orientation (a) opening the treasure chest (b), and changing clothes (c). Our method also supports adding new high-quality components with natural interactions with the original components.

NeRF and apply the same labeling strategy, unifying the results with the newly generated one. During the optimization, we render the newly labeled 3D mask regions to generate a precise 2D mask M_{2D} , which is used to calculate \mathcal{L}_{rgb} to preserve the unedited regions.

Globally optimizing the whole edited object can correctly determine the interaction between object components. However, we find that it tends to generate fuzzy results and have low faithfulness to the input sketches. To further improve the editing performance, we propose a local enhancement strategy that separately adds diffusion constraints in the local editing regions. Specifically, we construct the bounding sphere of the refined 3D mask. The sphere center defines the camera viewpoint while its radius defines the camera position. We utilize the local camera parameters to render the local editing regions, focusing the network attention into the interested components. During editing optimization, we randomly render the global images or local regions and optimize the following objective:

$$\mathcal{L}_{total}^{fine}(\theta) = \beta_1 \mathcal{L}_{SDS}^{3D} + \beta_2 \mathcal{L}_{ISM}^{2D} + \beta_3 \mathcal{L}_{rgb} + \beta_4 \mathcal{L}_{sil} + \beta_5 \mathcal{L}_{orient}. \quad (8)$$

We utilize the 2D diffusion loss in the fine stage to improve the details. \mathcal{L}_{rgb} is calculated similarly to that in the coarse stage, with M_{2D} rendered by the precise 3D mask.

5 EVALUATION

In this section, we conduct a series of qualitative and quantitative experiments to demonstrate the superiority of our framework to the alternative solutions. In Sec. 5.1, we present the sketch-based generation and editing results. In Sec. 5.2, we show a comparison with state-of-the-art methods. In Sec. 5.3, we conduct an ablation study to validate the effectiveness of the key components of our

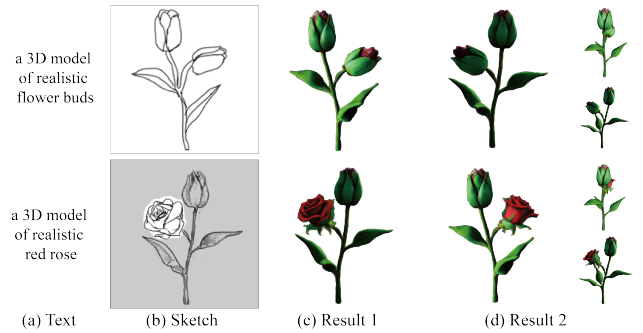


Fig. 7. Sketch-based generation and editing results. In the top row, given the input text (a) and sketch (b), our method generates realistic 3D results, as shown in (c) and (d) with different views. Users can edit the local regions in a selected view by modifying the text (a) and sketch (b) in the bottom row. Our method generates high-quality local modification results.

framework. We also conduct a user study in Sec. 5.4 to further support the better performance and interaction of our approach.

Implementation Details. Our networks are trained and tested on two NVIDIA RTX A6000 GPUs. During the training of the depth generation model, the learning rate is $1e-5$, the batch size is 64, and the number of training steps is 50k. To train 3D ControlNet, we set the learning rate to $1e-5$, batch size to 4 with gradient accumulation of 8, and the number of training steps to 30k. For the sketch-based generation, we optimize the NeRF representation with 12k steps. For the first 10k steps, $\lambda_1 = 1$, $\lambda_2 = 0$, $\lambda_3 = 1e2$, and λ_4 linearly increased from 10.0 to 1000.0 with 20k to 50k steps. For the remaining steps, there is a probability of 0.9 to set $\lambda_1 = 0$, $\lambda_2 = 1$, and otherwise




Fig. 8. Sketch-based generation comparison. For existing approaches, we first utilize ControlNet [Zhang et al. 2023] to generate 2D images and then use these approaches to generate 3D contents from the 2D images. Magic123 [Qian et al. 2023] generates realistic results in the sketch view but has weird geometry in other views. DreamCraft3D [Sun et al. 2023] generates better results in geometry and texture but still has obvious artifacts. With multi-view information, ImageDream [Wang and Shi 2023] generates correct geometry but has too light appearance. In contrast, our method generates better results with correct geometry and realistic appearance.

$\lambda_1 = 1, \lambda_2 = 0$. The resolution of images is set to 64×64 for the first 50k steps and changed into 256×256 afterward. For the sketch-based editing, in the coarse stage, we optimize the NeRF by 50k steps, with $\alpha_1 = 1.0, \alpha_2 = 1e5, \alpha_3 = 1e2$, and λ_4 being linearly increased from 10.0 to 1000.0 with 10k to 30k steps. In the fine stage, the hyperparameters are the same as those for generation, and $\beta_3 = 1e5$ for the first 10k steps and $\beta_3 = 1e6$ for the remaining steps. We implement sketch-based generation and editing on a single NVIDIA RTX A100. It takes 1.0-1.3 hours to generate a single example depending on the area size of the objects in sketches.

5.1 Results

Our method supports high-quality 3D generation from sketch and text inputs. As shown in Fig. 3, given text prompts and hand-drawn sketches, our approach generates high-quality 3D results respecting the input sketches and texts. Thanks to our sketch-based multi-view generation model, the generated results not only have good quality in the input views of sketches but also have abundant and realistic details in the back view, as shown in the jacket and dolls. Our generated results are in large scopes controlled by texts and not restricted to limited categories. As shown in Fig. 5, given the same hand-drawn sketches, users can input different text prompts to generate diverse and realistic results. It can be seen that sketches and texts provide complementary information: sketches control the geometry of results, while texts control their appearance.

Our method further supports sketch-based local region editing. As shown in Fig. 6, given the NeRFs reconstructed from the real models, users can select arbitrary views and then modify the rendered

sketches and provide text descriptions. Local editing regions are labeled by 2D masks. Our method generates realistic editing results with good quality in local editing regions, natural interactions with the original objects (e.g., the new dress and opening treasure chest), and unedited regions well preserved. Utilizing our method, users can modify the existing components (Fig. 6 (Left)) or add new components (Fig. 6 (Right)). Additionally, our method supports further editing of generated 3D contents. Fig. 7 (Top) shows a generated result given an input sketch and text prompt. Then, users can select a view and modify the rendered sketch and text prompt to achieve fine-grained editing of the results.

5.2 Comparison

Sketch-based Generation. Our method generates realistic 3D contents given single-view sketches. Since there are no existing approaches to synthesize high-quality results based on sketches and texts directly, we compared our method with an intuitive baseline: first utilizing a 2D sketch-to-image generation approach ControlNet [Zhang et al. 2023] to synthesize 2D images and then utilizing state-of-the-art image-to-3D approaches to generate 3D contents. We compared our approach with Magic123 [Qian et al. 2023], DreamCraft3D [Sun et al. 2023], and ImageDream [Wang and Shi 2023], which take single-view images and texts as input for 3D content generation. As shown in Fig. 8, Magic123 generates realistic results in the same views as the input sketches because of the high-quality generation results of 2D ControlNet and faithful input view reconstruction. However, the results exhibit weird geometry in other views, such as missing wings in the dragon example and distorted

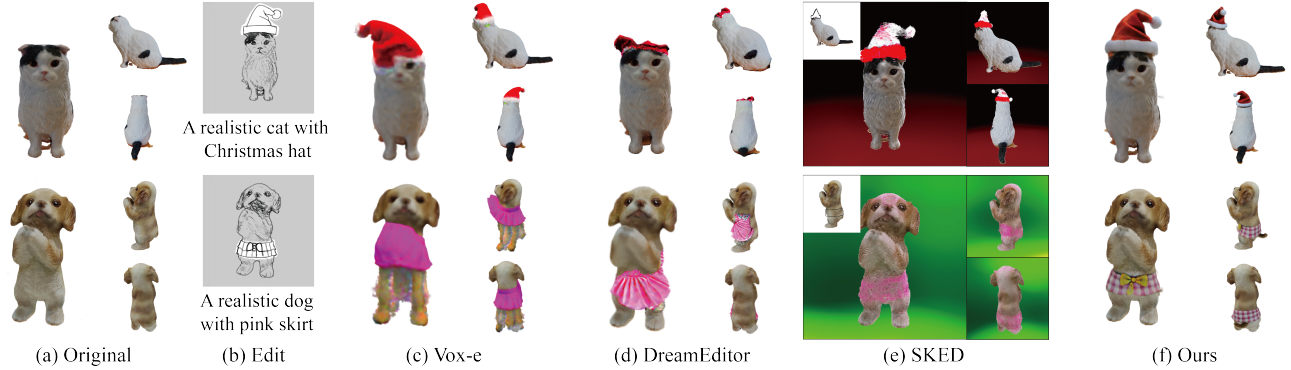


Fig. 9. Sketch-based editing comparison. Users can edit the reconstructed real 3D models by changing the texts and editing the rendered sketches. Vox-E [Sella et al. 2023] and DreamEditor [Zhuang et al. 2023] only support content editing via text prompts and thus cannot control the shape and editing regions. SKED [Mikaeili et al. 2023] allows editing of local regions with sketches in additional views (shown in the top-left corner), achieving detailed control, but its results are less realistic than ours.

Table 1. The quantitative comparisons with sketch-based generation methods, including Magic123 [Qian et al. 2023], DreamCraft3D [Sun et al. 2023], and ImageDream [Wang and Shi 2023]. The abbreviations “TF”, “SF”, “GQ”, and “TQ” mean text faithfulness, sketch faithfulness, geometry quality, and texture quality, respectively. The methods are evaluated in terms of the mean value of the CLIP score (CLIP space cosine similarity \times 100) and those metrics in the user study. We further report the standard deviation for the CLIP score. For all the metrics except *STD*, the higher, the better.

Method	CLIP		Human Preference				Overall
	Mean	STD	TF	SF	GQ	TQ	
Magic123	27.61	3.18	3.18	3.19	2.78	2.86	2.76
DreamCraft3D	29.34	2.56	3.35	3.32	2.95	2.91	2.87
ImageDream	27.61	2.75	3.40	3.49	3.06	2.89	2.97
Ours	30.90	2.13	4.71	4.52	4.67	4.72	4.72

legs in the piano example. Utilizing a better pre-trained model (i.e., stable-zero123) and DreamBooth for texture enhancement, DreamCraft3D generates more reasonable 3D geometry and more realistic textures. Still, it leads to obvious artifacts, such as the missing wings and blurry details. ImageDream utilizes a pre-trained multi-view diffusion model and solves the structure errors but still has an oblique problem, as shown in the piano example. In contrast, our method does not rely on intermediate images and directly generates 3D models from sketches. Due to the depth information analysis and multi-view 3D constraint, our method generates more realistic results in terms of both geometry and appearance.

Sketch-based Editing. Our method supports local editing of real objects based on sketches and text. To show its advantages, we compare it with existing local editing approaches, including Vox-E [Sella et al. 2023], DreamEditor [Zhuang et al. 2023], and SKED [Mikaeili et al. 2023]. Vox-E utilizes a voxel grid to reconstruct real objects and achieves text-based local editing. However, since the editing region has to be predicted by texts, it does not support accurate control of editing regions and results. For example, the left eye of the cat is undesirably changed, and the shape and texture

Table 2. The quantitative comparisons with SKED [Mikaeili et al. 2023] for sketch-based editing. The abbreviations “TF”, “SF”, “PU”, and “EQ” mean text faithfulness, sketch faithfulness, preservation of unedited regions, and editing component quality, respectively. The methods are evaluated in terms of the mean value of the CLIP score and those metrics in the user study. The standard deviation for the CLIP score is also included.

Method	CLIP		Human Preference				Overall
	Mean	STD	TF	SF	PU	EQ	
SKED	32.69	3.05	2.46	2.50	3.02	2.24	2.33
Ours	34.16	2.50	4.86	4.86	4.84	4.82	4.84

patterns of the dog example are unexpected. DreamEditor [Zhuang et al. 2023] also utilizes text to predict editing regions and thus might have mistaken editing regions such as the back view of dogs. Additionally, due to the NeRF definition on original meshes, it cannot support large-scale geometry editing, such as adding a Christmas hat. To achieve a more detailed control, SKED takes multi-view sketches with texts as input. However, it generates blurry results in the edited regions and slightly changes the unedited regions. Additionally, since SKED translates the input sketches into binary masks, it cannot control the texture details like the wrinkles of the hat and the decorative patterns of the skirt. In contrast, our method generates the most realistic editing results.

Quantitative Study. We assess our method and existing baseline approaches using CLIP Score [Radford et al. 2021]. This metric evaluates the correlation between the text prompts and rendering images by embedding them into a shared latent space and calculating the cosine similarity. The reasonableness of this metric is discussed in the supplemental material. For sketch-based generation, we compare our method against existing approaches on 20 examples. These examples involve a wide set of categories, ranging from animals, architecture, food, vehicles, and instruments. We personally crafted the prompts for these examples, providing concise descriptions of simple, everyday objects. As shown in Table 1,

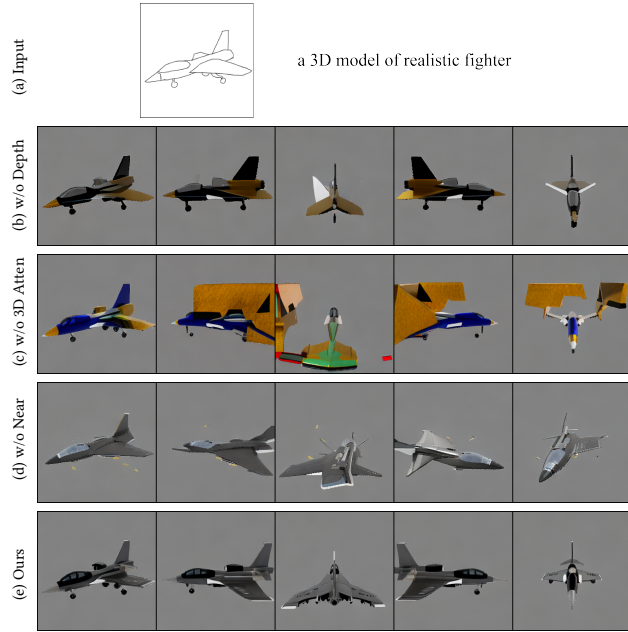


Fig. 10. Ablation study of the Multi-View ControlNet. Given a hand-drawn sketch and a text prompt (a), without the depth-guided warping (b), the generated results have good quality in the sketch view but are inconsistent across views. Without the 3D attention module (c), the results have strange components in the views other than the sketch view. If the input sketch is warped into all views (d), the generated images might suffer from twisted geometry. Our full method (e) generates the most realistic results in the input and novel views with good 3D consistency.

compared with the combination of 2D ControlNet and Image-to-3D baseline approaches, our method achieves the highest scores, thus validating the better alignment of our results with text prompts. For editing, we compare with the sketch- and text-based editing approach, SKED, on 15 editing examples. Similarly, these examples exhibit good diversity involving different categories. As shown in Table 2, our method also has the highest scores, validating the better editing alignment for text prompts. The examples used to calculate the CLIP scores are given in the supplement materials.

5.3 Ablation Study

We conduct ablation studies to prove the effectiveness of the key components in our framework. Specifically, we disable the key components of the sketch-based multi-view diffusion model to show their impacts. Then, we show the effectiveness of each loss term for the sketch-based 3D generation. For sketch-based editing, we also remove the two-stage strategy and local enhancement to prove their necessity.

In the sketch-based multi-view image generation model, we predict depth maps to warp the input sketches, explicitly building the spatial correspondence. As shown in Fig. 10 (b), without the depth warp strategy, the generated results have low 3D consistency, leading to different colors in the aircraft wings and additional parts in the back view. In order to add sketch control into the pre-trained

MVDream, we construct ControlNet with 3D attention. If we remove it and feed conditional inputs independently, the generated results tend to have strange parts in novel views due to the lack of 3D information. Since the predicted depth maps have slight errors, we only warp the input sketch into the nearest view, whose results are most reliable. If we warp it into all views, the warped sketches in far views have low quality, which largely affects the generation results, as shown in row (d). In contrast, our full model generates the best results with realistic appearance and good 3D consistency.

For sketch-based 3D generation, to add a sketch constraint in novel views, we utilize the SDS loss based on the sketch-based multi-view diffusion model. As shown in Fig. 11 (c), if we remove the 3D Loss, the generated results show good realism but have no relationship with the input sketches. We utilize the depth warp strategy to connect novel views with the input sketch view explicitly. As shown in Fig. 11 (d), without this strategy, sketches sometimes cannot control the viewpoints of the generated 3D contents. The control of details is also affected, e.g., causing an undesired tiger pose. During optimization, we utilize a silhouette loss to improve the detail control further. Without this silhouette loss (Fig.11 (e)), local details tend to misalign with sketches, such as the bigger frog feet and a tiger tail with a different shape. Additionally, if we remove the 2D loss during optimization, the generated results (Fig.11 (f)) have more fuzzy details, such as the dim frog texture and blurry tiger hair. Our full approach generates the most realistic results with fine details and the most faithfulness to the input sketches.

For sketch-based editing, the unedited regions should be retained while the edited components should be effectively modified. As shown in Fig. 12 (c), if the image loss with the original objects is removed, the whole objects are totally changed and have obvious floaters due to the lack of constraint with the original objects. When the local enhancement strategy is removed, the quality of edited regions is degraded. The faithfulness with the sketch is also affected. See the missing clothes of the bear and the smiles of the snowman in Fig. 12 (d). Our editing process includes two-stage optimization. As shown in (e), directly using the coarse 3D masks without the refined 3D masks has limited influences in the sketch view. However, in a novel view, unedited regions in the generated results (h) are undesirably modified, such as the strange bark and new cake pattern, which are included in coarse 3D masks. Compared with all the baseline approaches, our full method generates the most realistic results on the editing components while preserving the unedited regions.

5.4 Perception Study

We conducted a perception study to validate our method using human eyes. For sketch-based generation, we compare with the same set of approaches in qualitative comparison. The same 20 cases are also utilized, each containing an input sketch, a text prompt, and the synthesized 3D contents. Users evaluated the results on a scale of 1-5, where ‘1’ means very poor, while ‘5’ indicates excellent, across five aspects: text faithfulness, sketch faithfulness, geometry quality, texture quality, and overall quality. All methods presented video results of generated models, showcasing a full rotation in azimuth and a fixed elevation of 15 degrees. For each participant, we randomly

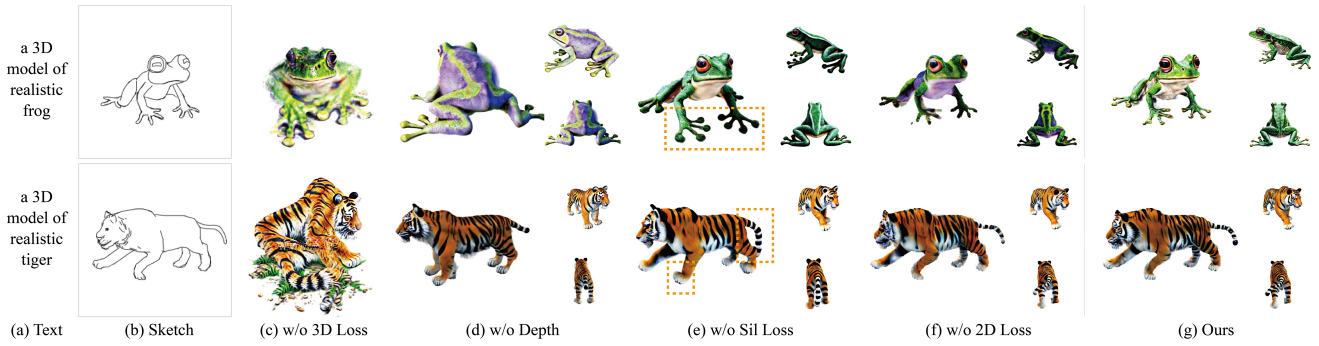


Fig. 11. Ablation study of sketch-based generation. The results without the 3D loss (c) are realistic but have a low correlation with the input sketches (b). Without the depth-guided warping (d), the generated results are low in faithfulness to the input sketches. Without the silhouette loss (e), the generated results misalign with sketches in local details, such as the feet and tail. Without the 2D loss (f), the results tend to be fuzzy.

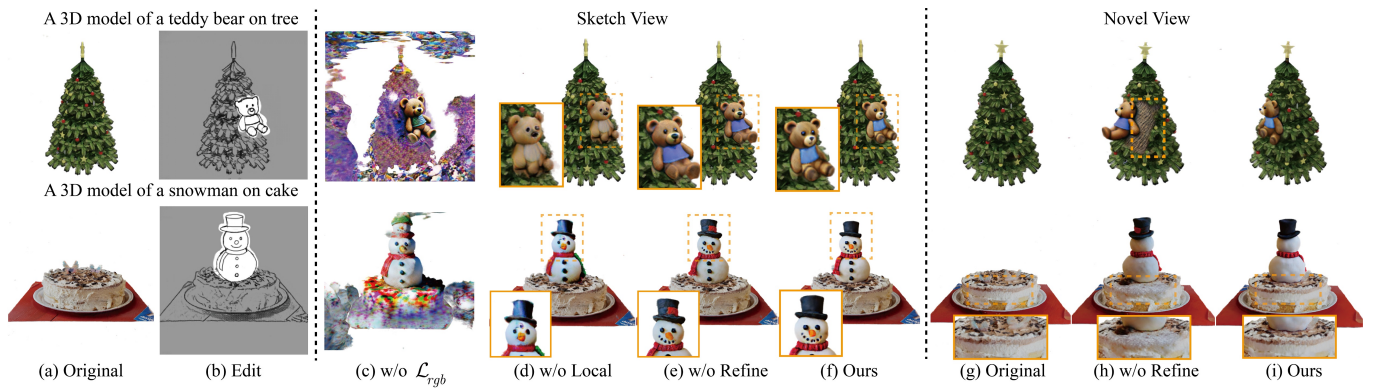


Fig. 12. Ablation study of sketch-based editing. The original 3D models (a) and the editing texts and sketches (b) are shown in the left regions. Without the image loss (c), unedited regions are totally changed. Without the local enhancement (d), the generated components are fuzzy and have obvious artifacts like the strange eyes. Replacing the fine stage mask with the coarse mask, the editing results are acceptable in view of sketches (e), but the unedited regions are changed in novel views (h).

selected 5 cases from the entire set, resulting in $5 \times 5 = 25$ answers per user. In total, 41 participants aged 18-40 without professional drawing skills contributed, yielding a dataset of $25 \times 41 = 1025$ answers. Our method outperformed existing baseline approaches, achieving the highest scores in all aspects according to Table 1. Notably, our approach excelled in text faithfulness but showed a slight weakness in sketch faithfulness. In the supplementary materials, box plots are drawn to show the evaluation scores. We also report the one-way ANOVA tests and paired T-tests to validate the superior performance of our approach.

For sketch-based editing, we employed a similar setup to the generation study but compared with SKED. Users also evaluated the results using a 1-5 scale across five aspects: text faithfulness, sketch faithfulness, preservation of unedited regions, editing component quality, and overall quality. For each user, we collected $5 \times 2 = 10$ answers, resulting in a total of $10 \times 41 = 410$ answers. As shown in Table 2, our method achieved better results compared with the baseline approach in all aspects. Similar to the generation study, our approach excelled in text faithfulness. However, it showed a slight weakness in the quality of newly generated components, a key

challenge in the editing process. Similar to sketch-based generation, we further report the box plot, one-way ANOVA tests, and paired T-test in the supplementary materials.

6 CONCLUSION

In this paper, we have presented a sketch-based 3D generation and editing method to create realistic high-quality 3D models from the input text prompts and hand-drawn sketches. In order to supplement the missing appearance and propagate the single-view sketch into 3D space, we design a sketch-based multi-view image generation diffusion model. To improve the sketch faithfulness and generation quality, we translate input sketches into depth maps, which are utilized to warp sketches into novel views to build 3D correspondence. The original and warped sketches serve as input conditions of 3D ControlNet, which has a 3D attention module to generate 3D consistent multi-view images. To generate realistic 3D contents based on sketches, we utilize the 3D SDS of sketch-based multi-view diffusion and 2D ISM of text-to-image diffusion, optimizing high-quality NeRF models. We further propose a sketch-based local editing method, which includes a coarse stage to get a precise




Fig. 13. The failure cases of our method. For sketch-based generation, given sketches (a) and text (above the images), our method generates results (b) that are misaligned with the sketches in detail. Additionally, our method cannot edit examples that very rare in dataset.

editing 3D mask, and a fine stage with local enhancement to generate high-quality editing results. Extensive experiments show the advantages of our method over baseline approaches.

Thanks to the good generalization of the pre-trained diffusion model, our method produces generated results that are not constrained to limited categories. However, as illustrated in Fig. 13, challenges arise for cases that are rare in the training dataset. In such instances, while our method generates a coarse result, the sketch faithfulness and realism may be influenced. Similarly, for sketch-based editing, our method faces limitations in handling what can be deemed ‘too strange cases,’ often resulting in degraded editing results with fuzzy details. For future work, although our method supports the appearance control based on text prompts, it is still difficult for users to achieve detail control of lighting, color, and material. Additionally, hand-drawn sketches loosely control the shape and layout while lacking too detailed control like the mushroom dots in Fig. 1. Adding additional conditions, such as color strokes, can partly solve this problem. Besides, our current system currently requires about 1 hour for generation and 1.5 hours for editing, and thus does not support interactive generation or editing. A potential solution to accelerate the sketch-based generation and editing process is to design 3D native generation pipelines [Hong et al. 2023] trained by sketch constraints.

ACKNOWLEDGMENTS

This work was supported by National Natural Science Foundation of China (No. 62322210), Beijing Municipal Natural Science Foundation for Distinguished Young Scholars (No. JQ21013), and Beijing Municipal Science and Technology Commission (No. Z231100005923031).

REFERENCES

- Caroline Chan, Frédo Durand, and Phillip Isola. 2022. Learning to generate line drawings that convey geometry and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 7915–7925.
- Ding-Yun Chen, Xiao-Pei Tian, Yu-Te Shen, and Ming Ouhyoung. 2003. On visual similarity based 3D model retrieval. In *Computer Graphics Forum*, Vol. 22. Wiley Online Library, 223–232.
- Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. 2023a. Fantasia3D: Disentangling Geometry and Appearance for High-quality Text-to-3D Content Creation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 22189–22199.
- Tao Chen, Ming-Ming Cheng, Ping Tan, Ariel Shamir, and Shi-Min Hu. 2009. Sketch2Photo: internet image montage. *ACM Transactions on Graphics* 28, 5 (2009), 124.
- Yang Chen, Yingwei Pan, Yehao Li, Ting Yao, and Tao Mei. 2023b. Control3d: Towards controllable text-to-3d generation. In *Proceedings of the 31st ACM International Conference on Multimedia*. 1148–1156.

- Xinhua Cheng, Tianyu Yang, Jianan Wang, Yu Li, Lei Zhang, Jian Zhang, and Li Yuan. 2023. Progressive3D: Progressively local editing for text-to-3D content creation with complex semantic prompts. *arXiv preprint arXiv:2310.11784* (2023).
- Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmid, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. 2023. Objaverse: A universe of annotated 3D objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13142–13153.
- Johanna Delanoy, Mathieu Aubry, Phillip Isola, Alexei A Efros, and Adrien Bousseau. 2018. 3D sketching using multi-view deep volumetric prediction. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 1 (2018), 1–22.
- Christoph Fehn. 2004. Depth-image-based rendering (DIBR), compression, and transmission for a new approach on 3D-TV. In *Stereoscopic displays and virtual reality systems XI*, Vol. 5291. 93–104.
- Thomas Funkhouser, Patrick Min, Michael Kazhdan, Joyce Chen, Alex Halderman, David Dobkin, and David Jacobs. 2003. A search engine for 3D models. *ACM Transactions on Graphics* 22, 1 (2003), 83–105.
- Chenjian Gao, Qian Yu, Lu Sheng, Yi-Zhe Song, and Dong Xu. 2022. SketchSampler: Sketch-Based 3D Reconstruction via View-Dependent Depth Sampling. In *European Conference on Computer Vision*. 464–479.
- Lin Gao, Feng-Lin Liu, Shu-Yu Chen, Kaiwen Jiang, Chun-Peng Li, Yu-Kun Lai, and Hongbo Fu. 2023b. SketchFaceNeRF: Sketch-based Facial Generation and Editing in Neural Radiance Fields. *ACM Transactions on Graphics* 42, 4 (2023), 159:1–159:17.
- William Gao, Noam Aigerman, Thibault Groueix, Vova Kim, and Rana Hanocka. 2023a. TextDeformer: Geometry Manipulation Using Text Guidance. In *ACM SIGGRAPH 2023 Conference Proceedings*. 82:1–82:11.
- Stephan J. Garbin, Marek Kowalski, Virginia Estellers, Stanislaw Szymanowicz, Shideh Rezaeifar, Jingjing Shen, Matthew Johnson, and Julien Valentin. 2022. VolTeMorph: Realtime, Controllable and Generalisable Animation of Volumetric Representations. *CoRR abs/2208.00949* (2022).
- Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. 2023. LRM: Large Reconstruction Model for Single Image to 3D. *CoRR abs/2311.04400* (2023).
- Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. 2006. Teddy: a sketching interface for 3D freeform design. In *ACM SIGGRAPH 2006 Courses*. 11–es.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5967–5976.
- Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. 2020. Sketch2CAD: Sequential CAD modeling by sketching in context. *ACM Transactions on Graphics* 39, 6 (2020), 164:1–164:14.
- Changjian Li, Hao Pan, Adrien Bousseau, and Niloy J Mitra. 2022. Free2CAD: Parsing freehand drawings into CAD commands. *ACM Transactions on Graphics* 41, 4 (2022), 93:1–93:16.
- Changjian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. 2018. Robust flow-guided neural prediction for sketch-based freeform surface modeling. *ACM Transactions on Graphics* 37, 6 (2018), 238:1–238:12.
- Yuhan Li, Yishun Dou, Yue Shi, Yu Lei, Xuanhong Chen, Yi Zhang, Peng Zhou, and Bingbing Ni. 2024. FocalDreamer: Text-Driven 3D Editing via Focal-Fusion Assembly. In *Conference on Artificial Intelligence*. 3279–3287.
- Zhiqi Li, Yiming Chen, Lingzhe Zhao, and Peidong Liu. 2023. MVControl: Adding Conditional Control to Multi-view Diffusion for Controllable Text-to-3D Generation. *arXiv preprint arXiv:2311.14494* (2023).
- Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. 2023. LucidDreamer: Towards High-Fidelity Text-to-3D Generation via Interval Score Matching. *arXiv preprint arXiv:2311.11284* (2023).
- Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3D: High-resolution text-to-3D content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 300–309.
- Jian Liu, Xiaoshui Huang, Tianyu Huang, Lu Chen, Yuenan Hou, Shixiang Tang, Ziwei Liu, Wanli Ouyang, Wangmeng Zuo, Junjun Jiang, and Xianming Liu. 2024. A Comprehensive Survey on 3D Content Generation. *CoRR abs/2402.01166* (2024).
- Kunhao Liu, Fangneng Zhan, Yiwen Chen, Jiahui Zhang, Yingchen Yu, Abdulmotaleb El Saddik, Shijian Lu, and Eric P Xing. 2023c. StyleRF: Zero-shot 3D Style Transfer of Neural Radiance Fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8338–8348.
- Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023b. Zero-1-to-3: Zero-shot one image to 3D object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9298–9309.
- Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. 2021. Editing conditional radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5773–5783.
- Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. 2023a. SyncDreamer: Generating Multiview-consistent Images from a Single-view Image. *CoRR abs/2309.03453* (2023). <https://doi.org/10.48550/arXiv.2309.03453>

- Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhansu Maji, and Rui Wang. 2017. 3D shape reconstruction from sketches via multi-view convolutional networks. In *International Conference on 3D Vision (3DV)*. 67–77.
- Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. 2023. RealFusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8446–8455.
- Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. 2023. Latent-NeRF for shape-guided generation of 3D shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12663–12673.
- Midjournal. 2022. *Midjournal*. <https://www.midjourney.com/>
- Aryan Mikaeili, Or Perel, Mehdi Safaei, Daniel Cohen-Or, and Ali Mahdavi-Amiri. 2023. SKED: Sketch-guided text-based 3D editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 14607–14619.
- Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024. T2I-Adapter: Learning Adapters to Dig Out More Controllable Ability for Text-to-Image Diffusion Models. In *Conference on Artificial Intelligence*. 4296–4304.
- Thu Nguyen-Phuoc, Feng Liu, and Lei Xiao. 2022. SNeRF: stylized neural implicit representations for 3D scenes. *ACM Transactions on Graphics* 41, 4 (2022), 142:1–142:11.
- Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *International Conference on Learning Representations*.
- Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, and Bernard Ghanem. 2023. Magic123: One Image to High-Quality 3D Object Generation Using Both 2D and 3D Diffusion Priors. *CoRR abs/2306.17843* (2023).
- Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi Zuo, Mutian Xu, Yushuang Wu, Weihao Yuan, Zilong Dong, Liefeng Bo, and Xiaoguang Han. 2023. RichDreamer: A Generalizable Normal-Depth Diffusion Model for Detail Richness in Text-to-3D. *arXiv preprint arXiv:2311.16918* (2023).
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. 8748–8763.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. 2023. DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 22500–22510.
- Etai Sella, Gal Fiebelman, Peter Hedman, and Hadar Averbuch-Elor. 2023. Vox-E: Text-guided voxel editing of 3D objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 430–440.
- Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. 2023a. Zero123++: a Single Image to Consistent Multi-view Diffusion Base Model. *CoRR abs/2310.15110* (2023).
- Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. 2023b. MV-Dream: Multi-view diffusion for 3D generation. *arXiv preprint arXiv:2308.16512* (2023).
- Nagabhushan Somraj. 2020. Pose-Warping for View Synthesis / DIBR. <https://github.com/NagabhushanSN95/Pose-Warping>
- Jia-Mu Sun, Tong Wu, and Lin Gao. 2024. Recent advances in implicit representation-based 3D shape generation. *Vis. Intell.* 2, 1 (2024).
- Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. 2023. Dreamcraft3D: Hierarchical 3D generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818* (2023).
- Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. 2023a. DreamGaussian: Generative Gaussian splatting for efficient 3D content creation. *arXiv preprint arXiv:2309.16653* (2023).
- Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. 2023b. Make-It-3D: High-fidelity 3D Creation from A Single Image with Diffusion Prior. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 22819–22829.
- Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2022a. CLIP-NeRF: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 3835–3844.
- Can Wang, Ruixiang Jiang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2023b. NeRF-Art: Text-driven neural radiance fields stylization. *IEEE Transactions on Visualization and Computer Graphics* (2023).
- Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. 2023a. Score jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12619–12629.
- Jiayun Wang, Jierui Lin, Qian Yu, Runtao Liu, Yubei Chen, and Stella X Yu. 2022b. 3D shape reconstruction from free-hand sketches. In *European Conference on Computer Vision*. 184–202.
- Lingjing Wang, Cheng Qian, Jifei Wang, and Yi Fang. 2018. Unsupervised learning of 3D model reconstruction from hand-drawn sketches. In *Proceedings of the 26th ACM international conference on Multimedia*. 1820–1828.
- Peng Wang and Yichun Shi. 2023. ImageDream: Image-Prompt Multi-view Diffusion for 3D Generation. *arXiv preprint arXiv:2312.02201* (2023).
- Tengfei Wang, Bo Zhang, Ting Zhang, Shuyang Gu, Jianmin Bao, Tadas Baltrusaitis, Jingjing Shen, Dong Chen, Fang Wen, Qifeng Chen, et al. 2023d. RODIN: A generative model for sculpting 3D digital avatars using diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4563–4573.
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. 2023c. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation with Variational Score Distillation. In *Advances in Neural Information Processing Systems*.
- Tong Wu, Zhibing Li, Shuai Yang, Pan Zhang, Xingang Pan, Jiaqi Wang, Dahua Lin, and Ziwei Liu. 2023. HyperDreamer: Hyper-Realistic 3D Content Generation and Editing from a Single Image. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.
- Weihao Xia and Jing-Hao Xue. 2024. A Survey on Deep Generative 3D-aware Image Synthesis. *ACM Comput. Surv.* 56, 4 (2024), 90:1–90:34.
- Nan Xiang, Ruibin Wang, Tao Jiang, Li Wang, Yanran Li, Xiaosong Yang, and Jianjun Zhang. 2020. Sketch-based modeling with a differentiable renderer. *Computer Animation and Virtual Worlds* 31, 4-5 (2020), e1939.
- Tianhan Xu and Tatsuya Harada. 2022. Deforming radiance fields with cages. In *European Conference on Computer Vision*. 159–175.
- Bangbang Yang, Chong Bao, Junyi Zeng, Hujun Bao, Yinda Zhang, Zhaopeng Cui, and Guofeng Zhang. 2022. NeuMesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *European Conference on Computer Vision*. 597–614.
- Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. 2021. Learning object-compositional neural radiance field for editable scene rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13779–13788.
- Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yewen Ma, Rongfei Jia, and Lin Gao. 2022. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 18353–18364.
- Robert C Zeleznik, Kenneth P Herndon, and John F Hughes. 2006. SKETCH: An interface for sketching 3D scenes. In *ACM SIGGRAPH 2006 Courses*. 9–es.
- Bohan Zeng, Shanglin Li, Yutang Feng, Hong Li, Sicheng Gao, Jiaming Liu, Huaxia Li, Xu Tang, Jianzhuang Liu, and Baochang Zhang. 2023. IPDreamer: Appearance-Controllable 3D Object Generation with Image Prompts. *arXiv preprint arXiv:2310.05375* (2023).
- Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. 2021b. Editable free-viewpoint video using a layered neural representation. *ACM Transactions on Graphics* 40, 4 (2021), 149:1–149:18.
- Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3836–3847.
- Song-Hai Zhang, Yuan-Chen Guo, and Qing-Wen Gu. 2021a. Sketch2model: View-aware 3D modeling from single free-hand sketches. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6012–6021.
- Xin-Yang Zheng, Hao Pan, Peng-Shuai Wang, Xin Tong, Yang Liu, and Heung-Yeung Shum. 2023. Locally Attentional SDF Diffusion for Controllable 3D Shape Generation. *ACM Transactions on Graphics* 42, 4 (2023), 91:1–91:13.
- Yue Zhong, Yulia Gryaditskaya, Honggang Zhang, and Yi-Zhe Song. 2020a. Deep sketch-based modeling: Tips and tricks. In *2020 International Conference on 3D Vision (3DV)*. 543–552.
- Yue Zhong, Yonggang Qi, Yulia Gryaditskaya, Honggang Zhang, and Yi-Zhe Song. 2020b. Towards practical sketch-based 3D shape generation: The role of professional sketches. *IEEE Transactions on Circuits and Systems for Video Technology* 31, 9 (2020), 3518–3528.
- Jingyu Zhuang, Chen Wang, Liang Lin, Lingjie Liu, and Guanbin Li. 2023. DreamEditor: Text-driven 3D scene editing with neural fields. In *SIGGRAPH Asia 2023 Conference Papers*. 1–10.