

SKED: Sketch-guided Text-based 3D Editing

Aryan Mikaeili¹ Or Perel² Mehdi Safaei¹ Daniel Cohen-Or³ Ali Mahdavi-Amiri¹

¹Simon Fraser University ²NVIDIA ³Tel Aviv University



Figure 1: Examples of our **Sketch-guided, Text-based 3D editing** method. Taking a pretrained Neural Radiance Field as input, multiview sketches determining the coarse region of edit and a text-prompt, our method is able to generate a localized, meaningful edit.

Abstract

Text-to-image diffusion models are gradually introduced into computer graphics, recently enabling the development of Text-to-3D pipelines in an open domain. However, for interactive editing purposes, local manipulations of content through a simplistic textual interface can be arduous. Incorporating user guided sketches with Text-to-image pipelines offers users more intuitive control. Still, as state-of-the-art Text-to-3D pipelines rely on optimizing Neural Radiance Fields (NeRF) through gradients from arbitrary rendering views, conditioning on sketches is not straightforward. In this paper, we present SKED, a technique for editing 3D shapes represented by NeRFs. Our technique utilizes as few as two guiding sketches from different views to alter an existing neural field. The edited region respects the prompt semantics through a pre-trained diffusion model. To ensure the generated output adheres to the provided sketches, we propose novel loss functions to generate the desired edits while preserving the density and radiance of the base instance. We demonstrate the effectiveness of our proposed method through several qualitative and quantitative experiments. <https://sked-paper.github.io/>

1. Introduction

Art is a reflection of the figments of human imagination. While many are limited in their practical creative capabilities, by pushing the boundaries of digital media, new ways can be created for casual artists and experts alike to express their ideas. At the same time, current neural generative art takes away much of the control from humans. In this work, we attempt to take a step towards restoring some of that control, enabling neural networks to complement users and naturally extend their skills rather than taking hold over the generative process.

The field of image synthesis has been significantly propelled by neural generative models, particularly by the latest text-to-image models that predominantly rely on large language-image models [3, 54, 55, 56]. These models have revolutionized the field of computer vision, as they can produce astonishing visual outcomes from text prompts alone.

The ability of text-to-image models has sparked a wave of editing methods that utilize these models. Many of these techniques rely on prompt editing [14, 18, 27, 36, 44, 52]. Nevertheless, simplifying the interface to text alone means users lack the necessary level of granularity to produce their exact desired outcomes. Sketch-guided editing, on the other

hand, provides intuitive control that aligns with user’s conventional drawing and painting skills. By incorporating user-guided sketches into text-to-image models, powerful editing systems can be created, offering a high degree of flexibility and fine-grained control for manipulating visual content [83, 73].

Although sketch-guided and text-driven methods have proven successful in generating and manipulating 2D images [40, 73, 9], it immediately raises the intriguing question of whether a similar approach could be developed to edit 3D shapes. Since direct text-to-3D models require an abundance of data to scale, state-of-the-art 3D generative models, such as DreamFusion [52] and Magic3D [36], which build on the capabilities of text-to-image models, may be considered as an alternative. However, maintaining control via conditioning with such models remains a challenging task, as these generative pipelines optimize a Neural Radiance Field (NeRF) [42] by amortizing gradients from a multitude of 2D views. In particular, providing consistent sketches across all possible views presents a hurdle for users. Instead, a plausible user interface should act with guidance from as few views as possible, e.g. up to two or three.

In this paper, we present **SKED**, a **SK**etch-guided **3D** **ED**iting technique. Our method acts on reconstructed or generated NeRF models. We assume a text prompt and a minimum of two sketches as input and provide output edits over the neural field faithful to the input conditions. Meeting all input requirements can be challenging as the text prompt may not match the sketch’s semantics, and sketch views may lack coherence. To undertake this complex task, we conceptually break it down into two subtasks that are easier to handle: one that depends on pure geometric reasoning and the other that exploits the rich semantic knowledge of the generative model. These two subtasks work together, with the former providing a coarse estimate of location and boundary, and the latter adding and refining geometric and texture details through fine-grained operations.

Our experiments highlight the effectiveness of our approach for editing various pretrained NeRF instances. We introduce assorted accessories, objects, and artifacts, which are generated and blended into the original neural field seamlessly. Finally, we validate our method through quantitative evaluations and ablation studies to assert the contribution of individual components in our method.

2. Related work

Sketch-Based 3D Modeling. Since its inception in the late twenties, traditional 2D animation has been concerned with creating believable depictions of 3D forms. Highly acclaimed art guidebook, *The Illusion of Life*, [23], advocated for solid three-dimensional drawings to practice “weight, depth and balance.” With the advancement of com-

puter animation, these principles have been widely adopted [32]. Sketch-based modeling is typically concerned with stitching and inflating, and user-drawn sketches to 3D meshes [78]. Starting with Teddy [20], early works focused on converting scribbles of 2D contours to intermediate forms such as closed polylines [26] or implicit functions [25, 65, 1, 57, 6]. Since lifting a single-view sketch to 3D is an under-constrained problem, additional constraints are usually introduced, such as correlating the inflated thickness with chordal axis depth of curved shapes [20], inferring shape and depth from user annotations [57, 26, 16, 49, 7, 80, 33, 22], using existing reference models like human figures [70], and solving a system based on user constraints [46, 24, 64, 13, 12]. More recently, data-driven approaches were suggested to lift and reconstruct objects from multi-view sketches with Conv-nets [39, 11, 34]. Our work departs from the former line of research by limiting a generative model to operate within the boundaries of a sketched region. By utilizing the strength of pretrained diffusion models conditioned on language, we avoid the intricacies of explicitly tuning inflation parameters or collecting large-scale train sets while being able to predict texture and shading simultaneously.

Diffusion Models. Diffusion models [60, 19, 61, 62] have emerged as an increasingly popular technique for generating diverse, high-quality images. More recently, they’ve been used to form state-of-the-art text-to-image models [55] by introducing language embeddings trained on massive amounts of data [56, 3]. Diffusion models are amicable for other conditioning modalities. Related to our work, [40] introduced conditioning on sketches. [73] trained a differentiable edge detector used to compute the edge loss per diffusion step, and [9] allowed finer granularity of control of generated images by distinguishing between sketch, stroke and level of realism. [83] is a contemporary publication which enables additional input conditions by augmenting large-diffusion models on small task-specific datasets.

Another line of works experimented with applying diffusion directly in 3D for generating point clouds [47] or 2D projections forming the feature representation of neural fields [59, 2]. While showing potential, the difficulty lies in scaling them due to the large amount of 3D data required.

Neural Fields. Neural Radiance Fields (NeRF) [42] have generated massive interest as means of representing 3D scenes using deep neural networks. Since then, a flurry of works has improved various aspects of optimized neural fields, yielding higher reconstruction quality [4, 76, 5, 71]. Neural field backbones, in particular, have become more structured and compressed [51, 67, 37, 8, 66]. The pivotal work of [45] introduced an efficient hash-based representation that allows NeRF optimizations to converge within seconds, effectively paving the way for interactive research directions on neural radiance fields. Recent works have ex-

explored interactive editing of neural radiance fields through manipulation of appearance latents [38, 48, 50, 8], by interacting with proxy representations [81, 10], through segmented regions and masks [30, 35, 43, 31] and text-based stylization [17, 75, 82, 43, 15].

Neural fields are an intriguing way to fully generate 3D models because, unlike meshes, they don’t depend on topological properties such as genus and subdivision [28, 15]. Initial generative text-to-3D attempts with NeRF backbones took advantage of a robust language model [53] to align each rendered view on some textual condition [21, 74]. However, without a geometric prior, [53] failed to produce multiview-consistent geometry. [79] used language guidance to generate 3D shapes based on shape embeddings, but their approach still requires large 3D datasets to generate template geometry.

DreamFusion [52] and [77] avoid the scarcity of 3D data by harnessing 2D diffusion models pretrained on large-scale, diverse datasets [58]. Their idea optimizes NeRF representations with score function losses from pretrained diffusion-models [55, 56], where the 2D diffusion process provides gradients for neural fields rendered from random views, and the process is amortized on many different views until an object is formed. Magic3D [36] further improved the quality and performance of generated 3D shapes with a 2-step pipeline which fine-tunes an extracted mesh. Their pipeline also allows for global prompt-based editing and stylization. Concurrently, Latent-NeRF [41] suggested optimizing neural fields in the diffusion latent space. Their work also suggested 3D bounded volumes as an additional constraint for guiding the generation process.

Our work builds on a simplified framework of [52], which operates in color space and aims to combine traditional sketch-based modeling constraints with the generative power of recent advances in the field. Our pipeline is a zero-shot generative setting, requiring no dataset and only text and sketch inputs from the user.

3. Method

In this section, we present our approach to performing text-based NeRF editing, controlled by a few given sketches. Our approach to addressing this issue involves dividing the problem into two substantially easier tasks. First, a rough 3D space that necessitates adjustment is defined using the provided sketches, which helps in guiding the geometry modifications. Second, we use the score distillation sampling method [52] on a text-to-image latent diffusion model to generate fine-detailed and realistic edits based on the text prompt given to the model (see Fig. 2).

To produce meaningful edits that adhere to the sketches, we design two novel objective functions: one to preserve the original density and radiance fields, and the second to alter the added mass in a way that respects the given sketches.

In the following, we describe our loss functions and provide details on how to apply sketch-based text-guided editing. We include a background on Latent diffusion models [55] and Score distillation sampling (SDS) [52] which our method is built upon, in the supplementary material.

3.1. SKED

As demonstrated in Fig. 2, the starting point of our algorithm is a base NeRF model, $F_o : (\mathbf{p}, \hat{\mathbf{r}}; \theta) \rightarrow (\mathbf{c}_o, \sigma_o)$, which maps 3D coordinates and unit rays to color and volumetric density. F_o is obtained through either reconstruction from multiview images [45], or a text-to-3D pipeline [52]. One can use F_o to render multiple views of the neural field and sketch over them to specify spatial cues of their desired edits. Let $\{C\}_{i=1}^N$ be renderings of F_o from N different views, on which we have drawn sketches. These sketches could be masks specifying the region of edit, or closed curves specifying the outer edges of the region of edit. Either way, the input to our algorithm would be preprocessed to masks $\{M\}_{i=1}^N$ where $M_i = \{m_1, m_2, \dots, m_S\}$ are a set of pixels which are inside the sketch region. We call these renderings *sketch canvases*, and the views they were rendered from *sketch views*. Additionally, the algorithm takes as input a text prompt T which defines the semantics of the edit. Similar to DreamFusion, our method is an iterative algorithm. We begin by initializing an editable copy of the base neural field, $F_e = F_o$. At each iteration, we sample a random view and use F_e to render the 3D object from that view. We use the rendered image to calculate the gradient of the SDS loss and push F_e to the high-density regions of the probability distribution function $\mathbb{P}(F_e|T)$ i.e. a NeRF model which its underlying 3D object adheres to text T . However, through our experiments, we found that simply using this process with only text as guidance would drastically change the original field outside of the sketched region. Therefore through two novel loss functions, we attempt to control the editing process in a way that the final output corresponds to the sketches is semantically meaningful and is faithful to the base neural field.

Preservation Loss: One of the main criteria of a good 3D editing algorithm is that the geometry and color of the base object are preserved through the editing process. We encourage this by utilizing an objective we call the preservation loss \mathcal{L}_{pres} . At each iteration of the algorithm, we render an image with F_e from a random camera viewpoint. We modify the raymarching algorithm of NeRF such that when sampling points $\mathbf{p}_i \in \mathbb{R}^3$, to query F_e for density and color values, we also compute a per-coordinate sketch-weight denoted as w_i . The key idea of our algorithm is that for each point \mathbf{p}_i , we decide whether it should be changed by calculating a distance from the point to the sketch masks. We aim to modify the density and radiance only when \mathbf{p}_i is in the proximity of the sketched regions while retaining the

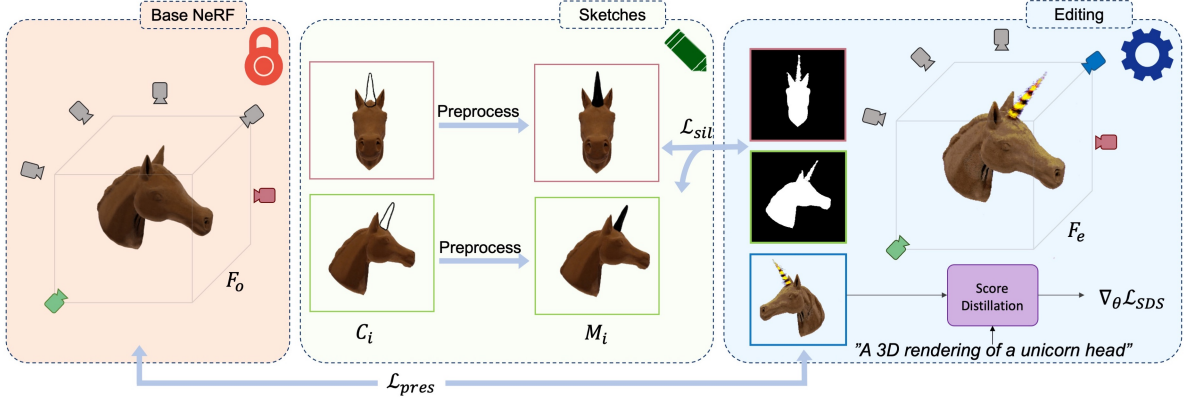


Figure 2: An overview of SKED. We render the base NeRF model F_o from at least two views and sketch over them (C_i). The input to the editing algorithm is these sketches preprocessed to masks (M_i) and a text prompt. In each iteration similar to DreamFusion [52], we render a random view and apply the Score Distillation Loss to semantically align with the text prompt. Additionally, we compute \mathcal{L}_{pres} to preserve the base NeRF by constraining F_e 's density and color output to be similar to F_o away from the sketch regions. Finally, we use the object mask renderings of the sketch views to define \mathcal{L}_{sil} . This loss ensures that the object mask occupies the sketch regions.

original density and radiance for points that are far from the sketched region. Therefore, we first need to define a method for computing the distance of a 3D point to multi-view 2D sketches. We do so by projecting \mathbf{p}_i to each of the sketch views and computing a per-view distance of projected points to sketch regions as:

$$d_j(\mathbf{p}_i) = \min_k \left\| \left[\Pi(\mathbf{p}_i, C_j) + \frac{1}{2} \right] - m_k \right\|^2, \quad (1)$$

where $d_j(\mathbf{p}_i)$ is the per-view distance function and $\Pi(\mathbf{p}_i, C_j)$ is the projection of 3D point p_i to sketch view C_j , rounded to the nearest integer (Fig. 3). This expression computes the minimum distance of the projection of point p_i to the sketch regions in view j . By taking the mean of all the per-view distances, we can define the distance of point p_i to the multiview sketches $D(\mathbf{p}_i) = \frac{1}{N} \sum_{j=1}^N d_j(\mathbf{p}_i)$. We use the mean of distances, as it relaxes the constraint when sketches are not fully aligned, and introduces additional smoothness to the function.

Now that we have established the distance function, we can define \mathcal{L}_{pres} . Using F_o , the base NeRF instance with frozen parameters, we define \mathcal{L}_{pres} as:

$$\mathcal{L}_{pres} = \frac{1}{K} \sum_{i=1}^K w_i [CE(\alpha_e, \bar{\alpha}_o) + \lambda_c \bar{\alpha}_o \|\mathbf{c}_e - \mathbf{c}_o\|^2], \quad (2)$$

where $\bar{\alpha}_o$ is the occupancy of the base object derived from F_o by thresholding the ground truth density of $F_o(\mathbf{p}_i; \theta)$. Following [42], we have $\alpha_e = 1 - \exp(-\sigma_e \delta)$, such that σ_e is the edited neural field density $F_e(\mathbf{p}_i; \phi)$, and δ is the step distance between samples along a ray. CE is the cross entropy loss. Furthermore, \mathbf{c}_e and \mathbf{c}_o are the color and ground

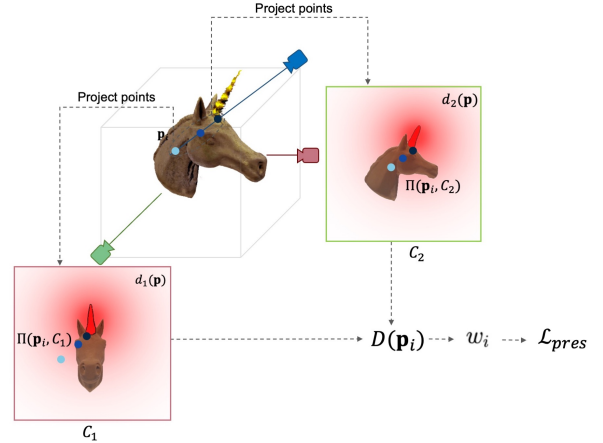


Figure 3: 3D points \mathbf{p}_i sampled at random views are projected to the sketch views C_j to obtain $\Pi(\mathbf{p}_i, C_j)$. In each C_j , distance d_j , between projected points and the pixels containing the sketch masks is computed. The red color in C_1 and C_2 demonstrates $d_1(\mathbf{p})$ and $d_2(\mathbf{p})$ in image space respectively. Finally, for each 3D point, $d_j(\mathbf{p}_i)$ s are averaged to get the distance $D(\mathbf{p}_i)$ to all sketch views. $D(\mathbf{p}_i)$ is used as the weights of the points w_i in \mathcal{L}_{pres} .

truth color of \mathbf{p}_i derived from F_e and F_o respectively, and λ_c is a hyperparameter controlling the importance of color preservation in the editing process. We limit color preservation to the occupied region of F_o . Tightly constraining the color may drive the model to diverge, hence λ_c is chosen such that density preservation takes higher priority. The preservation strength of each coordinate is modulated by:

$$w_i = 1 - \exp\left(-\frac{D(\mathbf{p}_i)^2}{2\beta^2}\right). \quad (3)$$

Intuitively, w_i controls the importance of the loss for each

point based on the distance $D(\mathbf{p}_i)$. The sensitivity of w_i to $D(\mathbf{p}_i)$ is controlled by the hyperparameter β , where lower values tighten the constraint on the optimization such that only the sketch region is modified. Higher values allow for a softer falloff region, which allows the optimized volume to better blend with the base model.

Silhouette Loss: Another essential criterion is to respect the sketched regions, i.e. the new density mass added to F_e should occupy the regions specified by the sketches. We enforce this by rendering the object masks of all sketch views. We then maximize the values of the object masks in the sketched regions by minimizing the following loss:

$$\mathcal{L}_{sil} = \frac{1}{H.W.N} \sum_{j=1}^N \sum_{i=1}^{H.W} -\mathbb{I}_{M_j}(\mathbf{x}_i) \log C_j^\alpha(\mathbf{x}_i). \quad (4)$$

In this equation, H and W are the dimensions of the rendered object masks, \mathbb{I}_{M_j} is an indicator function that is equal to 1 if pixel $\mathbf{x}_i \in \mathbb{R}^2$ is in a sketched region and 0 otherwise and C_j^α is the alpha object mask rendered with $F_e(\mathbf{p}, \hat{\mathbf{r}}; \phi)$ from each sketch view.

Optimization: Similar to prior generative works on NeRFs [21, 36, 52] we use an additional objective \mathcal{L}_{sp} to enforce sparsity of the object by minimizing the entropy of the object masks in each view. Therefore the final objective of our editing process is:

$$\mathcal{L}_{total} = \mathcal{L}_{SDS} + \lambda_{pres} \mathcal{L}_{pres} + \lambda_{sil} \mathcal{L}_{sil} + \lambda_{sp} \mathcal{L}_{sp}, \quad (5)$$

where λ_{pres} , λ_{sil} and λ_{sp} are the weights of the different terms in our objective.

We use Instant-NGP [45] as our neural renderer for its performance and memory efficiency. To avoid sampling empty spaces, this framework keeps an occupancy grid for tracking empty regions. The occupancy grid is used during raymarching to efficiently skip samples in empty spaces. In addition, the grid is periodically pruned during training to keep it aligned with the hashed feature structure. In our editing process, if the occupancy grid of F_o is used without change, the model may initially avoid sampling points in the sketch regions, preventing correct gradient flow and forcing our framework to rely on random alterations of the occupancy grid.

To alleviate this problem, we find the bounding boxes of the sketch masks M and intersect them in \mathbb{R}^3 to define a coarse editing region in 3D space. We manually turn on the occupancy grid bits of F_e within the sketch intersection region. In addition, we define a warm-up period at the beginning of optimization, where we avoid pruning the occupancy grid to help the model solidify the edited region, and prevent it from culling it as empty space.

4. Results and Evaluation

4.1. Implementation details

We use Stable-DreamFusion’s open-source GitHub repository [69] and integrate with kaolin-wisp’s renderer for an interactive UI [68]. In all our experiments, unless stated otherwise, we set λ_{pres} , λ_{sil} , λ_{sp} and λ_c to 5×10^{-6} , 1, 5×10^{-4} and 5 respectively. The guidance scale for classifier free guidance in the Stable Diffusion model is set to 100 and timesteps for noise scheduling are uniformly sampled in the range of (20, 980) in each iteration. The warm-up period for the occupancy grid pruning is set to 1,000 iterations. We also set the camera pose range to ensure that the sketch region remains visible in all sampled views. For large sketch regions, we use a Kd-tree [72] to implement Equation 3 with efficient queries. We use the ADAM [29] optimizer with a learning rate of 0.005 and apply the exponential scheduler with a decay of 0.1 by the end of the optimization. We run our algorithm for 10,000 iterations, taking approximately 30-40 minutes on a single NVIDIA RTX 3090 GPU. For our experiments, we use both publicly available 3D assets, and artificial assets generated by Stable-DreamFusion [69] guided only by text. We use the v1.4 version of the Stable diffusion [55] model. Unless specified otherwise, we use the same default hyperparameters settings throughout all experiments depicted in the paper.

4.2. Qualitative Results

Sketch and Text Control. Fig. 4 demonstrates examples of SKED on a variety of objects and shapes. Evidently, our method is able to satisfy the coarse geometry defined by user sketches, and at the same time naturally blend semantic details according to the given text-prompts. Note that the sketch is not required to be accurate or tight: by making the contour curve more complex, the user can further force the pipeline to generate a specific shape.

Next, we show that given a fixed pair of multiview sketches, our method produces semantic details to fit a diverse set of text-prompts (Fig. 6). Note that our method can generate details within the sketch boundary (e.g. Nutella jar) even if the sketch doesn’t match the text-prompt description. In Fig. 4 we also present the complementary case, by re-using the same text prompts and switching through different sketch sets, our method has the flexibility to produce localized edits (i.e., ”cherry on top of a sundae”).

Additionally, In Fig. 5 we demonstrate the ability of our method in performing various types of edits. We are able to perform both additive edits (crown on teddy’s head or whipped cream on pancake) and replacement edits where we overwrite a part of the object with a different part (tree to cactus or the long white skirt).

Table 1: Fidelity of base field. To assess a method’s ability to preserve the original content, we measure the **PSNR** \uparrow of the method’s output against renderings from the base model. SKED (*no-preserve*) refers to a variant of our method which doesn’t apply \mathcal{L}_{pres} . Text-Only refers to a public re-implementation of [52].

Method	Cat +chef hat		Cupcake +candle		Horse +horn		Sundae +cherry		Plant +flower		Mean
	View 1	View 2	View 1	View 2	View 1	View 2	View 1	View 2	View 1	View 2	
	SKED	31.05	34.13	23.73	25.98	32.45	31.46	26.47	25.99	21.71	
SKED (<i>no-preserve</i>)	15.58	16.59	20.12	19.47	18.02	16.52	17.39	17.44	10.16	10.12	16.14
Text-Only [69]	15.63	16.78	17.38	17.15	16.69	15.09	20.57	20.74	13.75	12.68	16.65



Figure 4: Examples of using SKED to edit various objects reconstructed with InstantNGP [45] (anime girl) or generated with DreamFusion [52] (plant, sundae, teddy bear, sundae, cupcake). All examples were edited using two sketch views and the text prompt.

Although not text based, we are also able to perform simple carving edits by masking off the 3D grid in the sketched regions (inset).



Base Model Distribution.

Our method assumes edits are applied on top of a base NeRF model. We explore both pretrained reconstructions from multiview images, generated with [45], and generated outputs from DreamFusion [52] using the same diffusion model we use for editing [55]. The renderings of reconstructed objects are assumed to follow the distribution of the

underlying diffusion model. Our method performs successfully on both reconstructed examples: Anime Girl, Plate, Cat, Horse (Fig. 4, 6, 7, 8) and generated ones: Plant, Teddy, Cupcake, Sundae (Fig. 4).

Progressive Editing. Our method can be used as a sequential editor where one reconstructs or generates a NeRF, and then progressively edits it. In Fig. 7 we exhibit a two-step editing by first reconstructing a cat object with [45], then generating a red tie, followed by adding a chef hat.

Preservation Sensitivity. We also demonstrate the overlay of the sketch masks with the edited NeRF rendered from the sketch views and the effect of changing β (Eq. 3) in



Figure 5: Various types of edits. SKED is capable of overwriting parts the base model (Cactus, Skirt), as well as adding new details (Pancake, Teddy).



Figure 6: Examples with a single set of sketches and a variety of text prompts. Our method is able to respect the geometry of the sketches while adding details to fit the different prompts’ semantics.




Figure 7: Progressive editing. The cat is first edited by adding a red tie, and then a chef hat is added in a subsequent edit.

Fig. 8. It is evident that increasing β changes the base NeRF more and more edits appear outside the sketch regions.

4.3. Quantitative Results

To the best of our knowledge, our method is the first to perform sketch-guided editing on neural fields. Hence, at the absence of an existing benchmark for systematic comparison, we also suggest a series of tests to quantify various aspects of our method. We conduct our evaluation using a set of five representative samples using the setting from Sec-

tion 4. Each sample includes a base shape, a pair of hand drawn sketches and a guiding text prompt. Comparisons to ”Text-only” ignore the input sketches and apply prompt editing. For a fair comparison, all experiments use the same diffusion model and implementation framework of [69]. In the following, we establish that all three metrics are necessary to quantify the method’s value.

Base Model Fidelity. We quantify our method’s ability to preserve the base field outside of the sketch area using PSNR (Table 1). As ground truth, we use $\{C \setminus M\}_{i=1}^{N=2}$, the renderings from the base model, excluding the filled sketch regions. We measure the PSNR w.r.t. to the output sketch view rendered with edited field F_e , and the output from [52] using the same camera view. Our results show that across all inputs, our method consistently preserves the original base field content, compared to the Text-only method which lacks this ability. Note that a method may obtain a perfect PSNR if it does not change the original neural field. Therefore, we further measure the quality of change as well.

Sketch Filling. To gauge whether our method respects the user sketches, we measure the ratio of sketch area filled with generated mass (Table 2). We denote the metric we use as *Intersection-over-Sketch*, and define it as $IoS = \sum_{i=1}^N |M_i \cap C_i^\alpha| / |M_i|$. Here, M_i is the sketch

Table 2: Sketch alignment score. We measure the similarity between the user input and generated result by Intersection-over-Sketch (IoS \uparrow). The IoS is calculated using the intersection between two views of filled sketches, M_i , and the alpha mask of generated edit C_i^α . See Section 4.3 for elaborate details of this metric. The SKED (*no-silh*) variant, which runs with \mathcal{L}_{pres} and without \mathcal{L}_{sil} avoids generating content in the sketch region (see also Fig. 9)

Method	Cat		Cupcake		Horse		Sundae		Plant		Mean
	+chef hat		+candle		+horn		+cherry		+flower		
	View 1	View 2	View 1	View 2	View 1	View 2	View 1	View 2	View 1	View 2	
SKED	0.9384	0.9689	0.8364	0.8875	0.6423	0.5363	0.7817	0.9096	0.9388	0.7801	0.8220
SKED (<i>no-silh</i>)	0.0196	0.0176	0.0263	0.0209	0.0090	0.0077	0.0541	0.0506	0.0024	0.0028	0.0211

Table 3: Semantic alignment score. We measure the CLIP-similarity [53] \uparrow of the rendered method output with the clip embedding of the input text prompt. Text-Only refers to a public re-implementation of [52]. The qualitative equivalents of Cat and Plant examples are depicted in Fig. 7 from the main paper: Compared with [52] which changes the structure of the base model to satisfy the text semantics, our method preserves the base model, while also maintaining semantic correlation with the text.

Method	Cat		Cupcake		Horse		Sundae		Plant		Mean
	+chef hat		+candle		+horn		+cherry		+flower		
SKED	0.2336 \pm 1.1e-3	0.2849 \pm 4.2e-3	0.2943 \pm 4.8e-3	0.2635 \pm 2.4e-3	0.2933 \pm 4.0e-3	0.2739					
Text-Only [69]	0.2744 \pm 4.4e-3	0.2818 \pm 6.2e-3	0.2928 \pm 8.4e-3	0.2674 \pm 4.9e-3	0.2865 \pm 3.3e-3	0.2806					




Figure 8: Sensitivity control. Depending on the sensitivity value determined by β in Eq. 3, our method can either edit only the sketched region and minimally modify the rest of the neural field, or produce larger edits outside the sketch regions (softer blending). We display the overlay of sketches v.s. edited output.

region and C_i^α is the thresholded alpha mask rendered with F_e from each sketch view. To ensure the metric is resilient to alpha thresholding values, the score we report is averaged over 9 thresholding values we apply to C_i^α , ranging at [25, 225]. We point out that a high *IoS* score by itself does not guarantee a high quality output, i.e. a method can cheat by simply filling the sketch region with some fixed color.

Semantic Alignment. We assess if our method generates semantically meaningful content aligning with the text-prompt using Clip-similarity [53]. In Table 3, we present the evaluations which demonstrate that although we don't optimize for CLIP performance directly, our method achieves comparable results with Text-Only. We perform this experiment by sampling forty views around each edited object and averaging the CLIP similarity of each view to the corresponding texts.

4.4. Ablation Studies

To perform sketch-based, local editing, we use two loss terms, \mathcal{L}_{pres} , which preserves the base object, and \mathcal{L}_{sil} which generates the desired edit according to the input sketches. We visually analyze the effect of these loss terms on two examples: one reconstructed with [45] (Fig. 9 top row) and another generated by Stable-DreamFusion [69] (Fig. 9 bottom row). We differentiate between these two examples as editing neural fields generated by Stable-DreamFusion ensures the rendered base model input is within the diffusion model distribution, which leads to fewer adversarial artifacts (see discussion in Section 4.2).

Qualitative ablation results are presented in Fig. 9. Text-Only, is equivalent to applying DreamFusion [52] initialized with a pretrained NeRF. This method employs neither of the two geometric losses and adheres to the semantics of the text prompt but drastically alters the base neural field, F_o . We also experimented with lowering the learning rate to avoid steering from the base model with high gradients, but that did not help in mitigating this effect.

When only \mathcal{L}_{sil} is employed, the sketch regions are edited according to the text-prompts, but the base region also changes drastically. The flower, a generated NeRF, changes more meaningfully compared to the cat. When only \mathcal{L}_{pres} is applied, no explicit constraint exists to respect the sketches. Therefore, our method yields a color artifact in the proximity of the sketch regions. When both constraints are simultaneously applied (our method), the edits respect both sketches and text prompt, and preserve the base NeRF.

We further validate our claim with a quantitative ablation




Figure 9: Ablation Study. We demonstrate of the effect of our suggested losses on editing neural fields (zoom in for details). The prompts used are "A cat wearing a *chef hat*" and "A *red flower stem* rising from a potted plant". All methods were initialized with the same base models (leftmost column), and optimized with the same diffusion model [55]. Text-only uses the public re-implementation of [52]. The rightmost method shows our full pipeline, compared to ablated versions of it omitting \mathcal{L}_{pres} and \mathcal{L}_{sil} respectively.



Figure 10: Effect of Diffusion Model Backbone. SKED is compatible with any diffusion models applicable with [52], e.g. the diffusion model should be trained to generate images in the editing domain, and respond to directional prompts, as designated by [52].

which repeats the experiments in Section 4.3 for the \mathcal{L}_{sil} and \mathcal{L}_{pres} variants (see Table 1 and Table 2). Evidently, both variants are inferior to the full pipeline.

5. Conclusion, Limitations, and Future Work

We presented SKED, a NeRF editing method conditioned on text and sketch. Using novel loss functions, our framework allows for local editing of neural fields. Similar to previous works [52, 36, 41], our approach utilizes the SDS Loss and may be vulnerable to the well-known "multiface issue" (inset figure) depending on the choice of diffusion model and prompt. Our method supports a single set of prompt and sketch views at a time. A simple workaround is to apply our method multiple times progressively (Fig. 7). Our results rely on the publicly available Stable-Diffusion



model [55], which is less amenable to directional text prompts and produces lower quality 3D generated outputs compared to commercial diffusion models used by previous works [52, 36]. In Fig 10 we show that it is possible to get better results by using the Deepfloyd-IF model [63].

Future directions may expand our method to better support for non-opaque materials, or condition on other modalities, possibly through the diffusion model. More research may further extend the usage of sketch scribbles for animation, similar to [12].

Acknowledgements: We thank Rinon Gal, Masha Shugrina, Roy Bar-On, and Janick Martinez Esturo for proofreading and helpful comments and discussions. This work was funded in part by NSERC Discovery grant (RGPIN-2022-03111), NSERC Discovery Launch Supplement (DGECR-2022-00359), and the Israel Science Foundation under Grant No. 2492/20.

References

- [1] A. Alexe, V. Gaidrat, and L. Barthe. Interactive modelling from sketches using spherical implicit functions. In *Proceedings of the 3rd International Conference on Computer Graphics, Virtual Reality, Visualisation and Interaction in Africa*, AFRIGRAPH '04, page 25–34, New York, NY, USA, 2004. Association for Computing Machinery. 2
- [2] Titas Anciukevicius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J. Mitra, and Paul Guerrero. RenderDiffusion: Image diffusion for 3D reconstruction, inpainting and generation. *arXiv*, 2022. 2
- [3] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. ediff-i: Text-to-image diffusion models with ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 1, 2
- [4] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*, 2021. 2
- [5] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 2
- [6] Adrien Bernhardt, Adeline Pihuit, Marie-Paule Cani, and Loïc Barthe. Matisse : Painting 2d regions for modeling free-form shapes. *SBIM 08*, 06 2008. 2
- [7] Minh Tuan Bui, Junho Kim, and Yunjin Lee. 3d-look shading from contours and hatching strokes. *Computers & Graphics*, 51:167–176, 2015. International Conference Shape Modeling International. 2
- [8] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *arXiv*, 2021. 2, 3
- [9] Shin-I Cheng, Yu-Jie Chen, Wei-Chen Chiu, Hung-Yu Tseng, and Hsin-Ying Lee. Adaptively-realistic image generation from stroke and sketch with diffusion model. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2023. 2
- [10] Chong Bao and Bangbang Yang, Zeng Junyi, Bao Hujun, Zhang Yinda, Cui Zhaopeng, and Zhang Guofeng. Neumesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [11] Johanna Delanoy, Mathieu Aubry, Phillip Isola, Alexei A Efros, and Adrien Bousseau. 3d sketching using multi-view deep volumetric prediction. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 1(1):1–22, 2018. 2
- [12] Marek Dvorožňák, Daniel Sýkora, Cassidy Curtis, Brian Curless, Olga Sorkine-Hornung, and David Salesin. Monster Mash: A single-view approach to casual 3D modeling and animation. *ACM Transactions on Graphics (proceedings of SIGGRAPH ASIA)*, 39(6), 2020. 2, 9
- [13] Lele Feng, Xubo Yang, Shuangjiu Xiao, and Fan Jiang. An interactive 2d-to-3d cartoon modeling system. In *International Conference on E-learning and Games*, 2016. 2
- [14] Rao Fu, Xiao Zhan, Yiwen Chen, Daniel Ritchie, and Srinath Sridhar. Shapecrafter: A recursive text-conditioned 3d shape generation model. *Neurips*, 2022. 1
- [15] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances In Neural Information Processing Systems*, 2022. 3
- [16] Yotam Gingold, Takeo Igarashi, and Denis Zorin. Structured annotations for 2D-to-3D modeling. *ACM Transactions on Graphics (TOG)*, 28(5):148, 2009. 2
- [17] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Learning Representations*, 2022. 3
- [18] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 1
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2
- [20] Takeo Igarashi, Satoshi Matsuoka, and Hidehiko Tanaka. Teddy: A sketching interface for 3d freeform design. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, page 409–416, USA, 1999. ACM Press/Addison-Wesley Publishing Co. 2
- [21] Ajay Jain, Ben Mildenhall, Jonathan T. Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. 2022. 3, 5
- [22] Pradeep Kumar Jayaraman, Chi-Wing Fu, Jianmin Zheng, Xueting Liu, and Tien-Tsin Wong. Globally consistent wrinkle-aware shading of line drawings. *IEEE Transactions on Visualization and Computer Graphics*, 24(7):2103–2117, 2018. 2
- [23] Thomas F Johnston and O. Disney. *Disney Animation: The Illusion of Life*. Abbeville Press, New York, 1st edition, 1981. 2
- [24] Pushkar Joshi and Nathan A. Carr. Repoussé: Automatic Inflation of 2D Artwork. In Christine Alvarado and Marie-Paule Cani, editors, *Eurographics Workshop on Sketch-Based Interfaces and Modeling*. The Eurographics Association, 2008. 2
- [25] Olga Karpenko, John Hughes, and Ramesh Raskar. Free-form sketching with variational implicit surfaces. *Computer Graphics Forum*, 21, 05 2002. 2
- [26] Olga A. Karpenko and John F. Hughes. Smoothsketch: 3d free-form shapes from complex sketches. *ACM Trans. Graph.*, 25(3):589–598, jul 2006. 2
- [27] Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-based real image editing with diffusion models. *arXiv preprint arXiv:2210.09276*, 2022. 1

- [28] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Popa Tiberiu. Clip-mesh: Generating textured meshes from text using pretrained image-text models. December 2022. [3](#)
- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. [5](#)
- [30] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. In *Advances in Neural Information Processing Systems*, volume 35, 2022. [3](#)
- [31] Zhengfei Kuang, Fujun Luan, Sai Bi, Zhixin Shu, Gordon Wetzstein, and Kalyan Sunkavalli. Palettenerf: Palette-based appearance editing of neural radiance fields. *ArXiv*, abs/2212.10699, 2022. [3](#)
- [32] John Lasseter. Principles of traditional animation applied to 3d computer animation. *SIGGRAPH Comput. Graph.*, 21(4):35–44, aug 1987. [2](#)
- [33] Changjian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. Bendsketch: Modeling freeform surfaces through 2d sketching. *ACM Trans. Graph.*, 36(4), jul 2017. [2](#)
- [34] Changjian Li, Hao Pan, Yang Liu, Xin Tong, Alla Sheffer, and Wenping Wang. Robust flow-guided neural prediction for sketch-based freeform surface modeling. *ACM Trans. Graph.*, 37(6), dec 2018. [2](#)
- [35] Gang Li, Heliang Zheng, Chaoyue Wang, Chang Li, Changwen Zheng, and Dacheng Tao. 3ddesigner: Towards photorealistic 3d object generation and editing with text-guided diffusion models, 2022. [3](#)
- [36] 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. *arXiv preprint arXiv:2211.10440*, 2022. [1](#), [2](#), [3](#), [5](#), [9](#)
- [37] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. [2](#)
- [38] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. *arXiv preprint arXiv:2105.06466*, 2021. [3](#)
- [39] Zhaoliang Lun, Matheus Gadelha, Evangelos Kalogerakis, Subhransu Maji, and Rui Wang. 3d shape reconstruction from sketches via multi-view convolutional networks. pages 67–77, 10 2017. [2](#)
- [40] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2022. [2](#)
- [41] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. *arXiv preprint arXiv:2211.07600*, 2022. [3](#), [9](#)
- [42] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. [2](#), [4](#)
- [43] Ashkan Mirzaei, Yash Kant, Jonathan Kelly, and Igor Gilitschenski. LaTeRF: Label and text driven object radiance fields. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [3](#)
- [44] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. *arXiv preprint arXiv:2211.09794*, 2022. [1](#)
- [45] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. [2](#), [3](#), [5](#), [6](#), [8](#)
- [46] Andrew Nealen, Takeo Igarashi, Olga Sorkine, and Marc Alexa. FiberMesh: Designing freeform surfaces with 3D curves. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)*, 26(3):article no. 41, 2007. [2](#)
- [47] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts, 2022. [2](#)
- [48] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. [3](#)
- [49] Luke Olsen, Faramarz Samavati, and Joaquim Jorge. Natursketch: Modeling from images and natural sketches. *Computer Graphics and Applications, IEEE*, 31:24 – 34, 01 2012. [2](#)
- [50] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A high-dimensional representation for topologically varying neural radiance fields. *ACM Trans. Graph.*, 40(6), dec 2021. [3](#)
- [51] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision (ECCV)*, 2020. [2](#)
- [52] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [9](#)
- [53] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, and Jack Clark. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning, PMLR*, 8748–8763, page 8748–8763, 2021. [3](#), [8](#)
- [54] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical textconditional image generation with clip latents, 2022. [1](#)
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *arXiv preprint arXiv:2112.10752*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [9](#)
- [56] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image

- diffusion models with deep language understanding, 2022. 1, 2, 3
- [57] R. Schmidt, B. Wyvill, M. C. Sousa, and J. A. Jorge. Shapeshop: Sketch-based solid modeling with blobtrees. In *ACM SIGGRAPH 2006 Courses*, SIGGRAPH '06, page 14–es, New York, NY, USA, 2006. Association for Computing Machinery. 2
- [58] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 3
- [59] J. Ryan Shue, Eric Ryan Chan, Ryan Po, Zachary Ankner, Jiajun Wu, and Gordon Wetzstein. 3d neural field generation using triplane diffusion, 2022. 2
- [60] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2256–2265, Lille, France, 07–09 Jul 2015. PMLR. 2
- [61] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020. 2
- [62] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS'20, Red Hook, NY, USA, 2020. Curran Associates Inc. 2
- [63] StabilityAI. Deepfloyd if: A modular cascaded diffusion model, 2023. <https://www.deepfloyd.ai/deepfloyd-if>. 9
- [64] Daniel Šýkora, Ladislav Kavan, Martin Čadík, Ondřej Jamříška, Alec Jacobson, Brian Whited, Maryann Simmons, and Olga Sorkine-Hornung. Ink-and-ray: Bas-relief meshes for adding global illumination effects to hand-drawn characters. *ACM Trans. Graph.*, 33(2), apr 2014. 2
- [65] Chiew-Lan Tai, Hongxin Zhang, and Jacky Fong. Prototype modeling from sketched silhouettes based on convolution surfaces. *Comput. Graph. Forum*, 23:71–84, 03 2004. 2
- [66] Towaki Takikawa, Alex Evans, Jonathan Tremblay, Thomas Müller, Morgan McGuire, Alec Jacobson, and Sanja Fidler. Variable bitrate neural fields. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH '22, New York, NY, USA, 2022. Association for Computing Machinery. 2
- [67] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. 2021. 2
- [68] Towaki Takikawa, Or Perel, Clement Fuji Tsang, Charles Loop, Joey Litalien, Jonathan Tremblay, Sanja Fidler, and Maria Shugrina. Kaolin wisp: A pytorch library and engine for neural fields research. <https://github.com/NVIDIAGameWorks/kaolin-wisp>, 2022. 5
- [69] Jiayang Tang. Stable-dreamfusion: Text-to-3d with stable-diffusion, 2022. <https://github.com/ashawkey/stable-dreamfusion>. 5, 6, 7, 8
- [70] Emmanuel Turquin, Jamie Wither, Laurence Boissieux, Marie-paule Cani, and John F. Hughes. A sketch-based interface for clothing virtual characters. *IEEE Computer Graphics and Applications*, 27(1):72–81, 2007. 2
- [71] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-NeRF: Structured view-dependent appearance for neural radiance fields. *CVPR*, 2022. 2
- [72] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. 5
- [73] Andrey Voynov, Kfir Aberman, and Daniel Cohen-Or. Sketch-guided text-to-image diffusion models. 2022. 2
- [74] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. *arXiv preprint arXiv:2112.05139*, 2021. 3
- [75] Can Wang, Ruixiang Jiang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Nerf-art: Text-driven neural radiance fields stylization. *arXiv preprint arXiv:2212.08070*, 2022. 3
- [76] Chen Wang, Xian Wu, Yuan-Chen Guo, Song-Hai Zhang, Yu-Wing Tai, and Shi-Min Hu. Nerf-sr: High-quality neural radiance fields using supersampling. *arXiv*, 2021. 2
- [77] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *arXiv preprint arXiv:2212.00774*, 2022. 3
- [78] L. Williams. Shading in two dimensions. In *Proceedings of Graphics Interface '91*, GI '91, pages 143–151, Toronto, Ontario, Canada, 1991. Canadian Man-Computer Communications Society. 2
- [79] 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. *arXiv preprint arXiv:2212.14704*, 2022. 3
- [80] Chih-Kuo Yeh, Shicun Huang, Pradeep Kumar Jayaraman, Chi-Wing Fu, and Tong-Yee Lee. Interactive high-relief reconstruction for organic and double-sided objects from a photo. *IEEE Transactions on Visualization and Computer Graphics*, 23:1796–1808, 2017. 2

- [81] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. [3](#)
- [82] Kai Zhang, Nick Kolkin, Sai Bi, Fujun Luan, Zexiang Xu, Eli Shechtman, and Noah Snavely. Arf: Artistic radiance fields, 2022. [3](#)
- [83] Lvmin Zhang and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. [2](#)