

# Consistent3D: Towards Consistent High-Fidelity Text-to-3D Generation with Deterministic Sampling Prior

Zike Wu<sup>1,4</sup> Pan Zhou<sup>\*2,4</sup> Xuanyu Yi<sup>1,4</sup> Xiaoding Yuan<sup>3</sup> Hanwang Zhang<sup>1,5</sup>

<sup>1</sup> Nanyang Technological University <sup>2</sup> Singapore Management University

<sup>3</sup> Johns Hopkins University <sup>4</sup> Sea AI Lab <sup>5</sup> Skywork AI

zike001@e.ntu.edu.sg, panzhou@smu.edu.sg, xuanyu001@e.ntu.edu.sg, xyuan19@jhu.edu, hanwangzhang@ntu.edu.sg



*A ceramic lion*



*An astronaut is riding a horse*



*A panda rowing a boat*



*A skiing penguin wearing a puffy jacket*



*A tiger wearing a tuxedo*



*A baby dragon hatching out of a stone egg*



*The Himeji Castle*



*The Great Wall*

Figure 1. Examples generated by Consistent3D. Our methods can generate detailed, diverse 3D objects and large-scale scenes from a wide range of textual prompts.

## Abstract

Score distillation sampling (SDS) and its variants have greatly boosted the development of text-to-3D generation, but are vulnerable to geometry collapse and poor textures yet. To solve this issue, we first deeply analyze the SDS and find that its distillation sampling process indeed corresponds to the trajectory sampling of a stochastic differential equation (SDE): SDS samples along an SDE trajectory to yield a less noisy sample which then serves as a guidance to optimize a 3D model. However, the randomness in SDE sampling often leads to a diverse and unpredictable

sample which is not always less noisy, and thus is not a consistently correct guidance, explaining the vulnerability of SDS. Since for any SDE, there always exists an ordinary differential equation (ODE) whose trajectory sampling can deterministically and consistently converge to the desired target point as the SDE, we propose a novel and effective “Consistent3D” method that explores the ODE deterministic sampling prior for text-to-3D generation. Specifically, at each training iteration, given a rendered image by a 3D model, we first estimate its desired 3D score function by a pre-trained 2D diffusion model, and build an ODE for trajectory sampling. Next, we design a consistency distillation sampling loss which samples along the ODE trajectory to generate two adjacent samples and uses the less noisy

\*Corresponding author.

sample to guide another more noisy one for distilling the deterministic prior into the 3D model. Experimental results show the efficacy of our Consistent3D in generating high-fidelity and diverse 3D objects and large-scale scenes, as shown in Fig. 1. The codes are available at <https://github.com/sail-sg/Consistent3D>.

## 1. Introduction

Diffusion models (DMs) have recently garnered significant attention in the realm of image synthesis, as evidenced by their remarkable capabilities [34, 48]. This notable progress can be largely attributed to the integration of large-scale image-text pair datasets and the evolution of scalable generative model architectures [30, 35]. This recent success has seamlessly transcended into the domain of text-to-3D generation by leveraging the pre-trained 2D diffusion models [33, 34] to guide the 3D generation process, regardless of the absence of large-scale 3D generative models [3, 45].

The pivotal breakthrough in this field stems from the finding that one can use the score function predicted by pre-trained 2D diffusion models, such as Stable Diffusion [34], to estimate the 3D score function [7, 14, 45]. Since this score function indicates the direction of the higher data density [10, 40], one can first use it to build a stochastic differential equation (SDE) [41], and then sample along the SDE solution trajectory (*i.e.*, SDE reverse process) to iteratively improve a learnable 3D model (*e.g.*, NeRF [26] or Mesh [37]). This is also the underlying mechanism behind the prevalent and leading text-to-3D approach, Score Distillation Sampling (SDS) [31]. In each training iteration, SDS follows the forward SDE to inject noise into a rendered image by a learnable 3D model, and then samples a more realistic pseudo-image along the SDE solution trajectory, where the 3D score function of the SDE is estimated by a pre-trained diffusion model [45]. Next, SDS pulls its rendered image closer to the pseudo-image via optimizing the learnable 3D model.

However, as illustrated in Fig. 2, the high randomness inherent in the SDE solution distribution [11, 41] leads to a highly diverse and unpredictable next point in the solution trajectory, *e.g.*, the pseudo-image [38, 55] in SDS. Although this trajectory may eventually converge to a specific target, *e.g.*, the desired realistic image in SDS, the sampled next point does not always provide the correct guidance in each iteration. This lack of reliability also applies to SDS, significantly increasing the optimization difficulty of the 3D model. It also helps to explain why SDS is so vulnerable and often suffers from geometry collapse and poor fine-grained texture in practice [38, 43, 47].

To address this critical issue, in this paper, we propose a novel and effective method, dubbed “Consistent3D”, which guides text-to-3D generation using deterministic sampling




Figure 2. Comparison between the (reverse) trajectory samplings in the stochastic differential equation (SDE) and ordinary differential equation (ODE).

prior. As illustrated in Fig. 2, for these unpredictable and uncontrollable trajectories sampled from SDE solution distribution, there theoretically always exists a corresponding ordinary differential equation (ODE) whose trajectory shares the same marginal distributions with the SDE solution [41]. Importantly, this ODE trajectory is deterministic and consistently converges to the same target point as the SDE. Sampling along the ODE trajectory guarantees a predictable and deterministic next point, which always directs towards the desired target and thus providing a reliable and consistent guidance. This motivates us to explore the text-to-3D generation from the ODE deterministic sampling perspective.

Specifically, during each training iteration, we begin by estimating the desired 3D score function from the rendered images produced by the learnable 3D model using pre-trained 2D diffusion models. Subsequently, we build a corresponding ODE for solution trajectory sampling. To effectively optimize the underlying 3D representations, we then introduce a Consistency Distillation Sampling loss (CDS), which leverages deterministic sampling prior along the ODE flow. In detail, for each rendered image, we first inject a *fixed* noise to the rendered image so that the corresponding noisy sample lies in the ODE solution distribution and thus can be well denoised. Following this, we sample two *adjacent* points from the ODE trajectory given the noisy sample. The less-noisy sample is then used to guide its more-noisy counterpart, thereby distilling the deterministic prior of the ODE trajectory into the 3D model. Here we use fixed noise to ensure that the samplings from the ODE trajectory for all rendered images converge to the same targeted realistic image, thereby offering more consistent guidance and enhancing the optimization of the 3D model.

Extensive experimental results showcase the efficacy of Consistent3D in generating high-fidelity and diverse 3D objects, along with large-scale scenes, as shown in Fig. 1 and Fig. 4. Comparative evaluations against existing methods, including DreamFusion [31], Magic3D [20] and ProlificDreamer [47], demonstrate the superiority of Consistent3D

in both qualitative and quantitative terms. The proposed approach effectively addresses the challenges associated with randomness in the SDE solution distribution, offering a more reliable and consistent framework to guide the text-to-3D generation process.

## 2. Related Works

**Diffusion Models** [11, 41, 50] are powerful tools for complex data modeling and generation. Their robust and stable capabilities for complex data modeling have also led to their successful application in various domains, such as image [4, 8, 48], video [12, 13, 18], and 3D [31, 45], *etc.* Regarding improving the sampling efficiency, there are two main approaches: learning-free sampling and learning-based sampling. Learning-free sampling typically involves discretizing reverse-time SDE [5, 41] or ODE [16, 21, 23, 39, 52], while learning-based sampling is mainly based on knowledge distillation [25, 36, 42]. This paper is driven by recent progress in learning-based sampling, particularly in distilling knowledge from ODE sampling [36, 42].

**Text-to-3D Generation** stands for generating 3D contents from a given text description. Current 3D generative models [15, 29], usually work in a single object category and suffer from limited diversity due to the lack of large-scale 3D datasets. To achieve open-vocabulary 3D generation, pioneered by DreamFusion [31], several approaches propose to lift text-image diffusion models [34] for 3D generation [47, 54, 55]. The key mechanism of such approaches is the score distillation sampling (SDS), where diffusion priors are used to supervise the optimization of a 3D representation. The following works continue to further improve the stability and fidelity of generation of various aspects, *e.g.*, advanced 3D representation [3, 43, 46, 51], coarse-to-fine training strategy [20, 47, 55] and 3D-aware diffusion priors [19, 22, 38, 44, 54].

## 3. Preliminaries

**Diffusion Models (DMs).** They consist of a forward diffusion process and a reverse sampling process. During the forward process, DMs gradually add Gaussian noise to the vanilla sample  $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$  and generate a series of noisy samples  $\mathbf{x}_t$  according to the distribution:

$$p_t(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \mathbf{x}_0, \sigma_t^2 \mathbf{I}), \quad (1)$$

where  $\sigma_t$  varies along time-step  $t$ . Accordingly, one can easily sample a noisy sample at any time step  $t$  by  $\mathbf{x}_t = \mathbf{x}_0 + \sigma_t \epsilon_t$ , where  $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ .

The reverse process from a Gaussian noise  $\mathbf{x}_T$  to a realistic sample  $\mathbf{x}_0$  is also called trajectory sampling, and can be formally formulated into a reverse SDE [16]:

$$d\mathbf{x} = -\dot{\sigma}_t \sigma_t \nabla \log p_t(\mathbf{x}) dt + \sqrt{\dot{\sigma}_t \sigma_t} d\mathbf{w}, \quad (2)$$

where  $\mathbf{w}$  is the standard Wiener process,  $\dot{\sigma}_t$  is the time derivative of  $\sigma_t$ , and  $\nabla \log p_t(\mathbf{x})$  is the *score function* which indicates the direction of the higher data density [41]. Meanwhile, there exists a corresponding reverse ordinary deterministic equation (ODE) which is defined as follows:

$$d\mathbf{x} = -\dot{\sigma}_t \sigma_t \nabla \log p_t(\mathbf{x}) dt. \quad (3)$$

For this ODE, its trajectory shares the same marginal probability density as the SDE which ensures the same convergence point of ODE and SDE.

Given any noise  $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \sigma_T^2 \mathbf{I})$ , one can either solve the reverse SDE in Eq. (2) or the reverse ODE in Eq. (3) via any numerical solver [1, 23, 53] to generate a real sample  $\hat{\mathbf{x}}_0 \sim p_{\text{data}}(\mathbf{x})$ . In practice, the pre-trained diffusion models [16, 34] are used to estimate the score function, thereby guiding the sampling process.

**Text-to-3D Generation via Score Distillation Sampling (SDS).** Given a camera pose  $\pi$ , SDS distills 2D priors of a pre-trained diffusion model  $D_\phi(\cdot)$  into a 3D model (*e.g.*, NeRF, Mesh) parameterized by  $\theta$ . Formally, SDS applies a denoising training objective to the rendered image  $\mathbf{x}_\pi = g(\theta, \pi)$  where  $g(\cdot)$  is a differentiable renderer and  $\pi$  is a camera pose, and computes the gradient as

$$\nabla_\theta \mathcal{L}_{\text{SDS}}(\theta) = \mathbb{E}_{t, \epsilon} \left[ \lambda(t) (\mathbf{x}_\pi - D_\phi(\mathbf{x}_t, t, y)) \frac{\partial \mathbf{x}_\pi}{\partial \theta} \right], \quad (4)$$

where  $\lambda(t)$  denotes the loss weight,  $\mathbf{x}_t = \mathbf{x}_\pi + \sigma_t \epsilon_t$  denotes the noisy sample and  $y$  is the text condition.

## 4. Method

Here we elaborate on our proposed Consistent3D for effective text-to-3D generation. In Sec. 4.1, we reveal the underlying mechanism of Score Distillation Sampling (SDS) which aims to approximate the SDE sampling process and motivates our proposed methods. Then in Sec. 4.2, we introduce Consistency Distillation Sampling (CDS), a loss designed to efficiently distill deterministic sampling priors for text-to-3D generation, accompanied with a theoretical justification regarding the error bound. Finally, we present how to use our CDS to build our text-to-3D generation framework, dubbed ‘‘Consistent3D’’, in Sec. 4.3.

### 4.1. Revisit Score Distillation Sampling

Before introducing our proposed method, we first connect SDE in Eq. (2) with the leading text-to-3D generation approach SDS in Eq. (4), since this connection directly motivates us to use ODE in Eq. (3) for text-to-3D generation.

First, we discretize the reverse SDE in Eq. (2) and perform the stochastic sampling process following Ho et al. [11], which results in the SDE solution trajectory defined as

$$\begin{aligned} \mathbf{x}_{t_i} &= \hat{\mathbf{x}}^{i-1} + \sigma_{t_i} \epsilon_{t_i} \quad \text{with} \quad \epsilon_{t_i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ \hat{\mathbf{x}}^i &= D_\phi(\mathbf{x}_{t_i}, t_i), \end{aligned} \quad (5)$$

---

**Algorithm 1:** Text-to-3D Generation with CDS

---

**Input** : initial 3D model parameter  $\theta$ , pre-trained diffusion model  $D_\phi$ , text prompt  $y$ , training iteration  $N$ , time-step range  $[t_{\min}, t_{\max}]$ , learning rate  $\eta$

**Output:**  $\theta$

**Sample**  $\epsilon^* \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$  // Fixed noise

**foreach**  $i \in \{0, \dots, N\}$  **do**

$t_2 \leftarrow t_{\max} - (t_{\max} - t_{\min})\sqrt{i/N}$

**Sample** camera pose  $\pi$

**Sample**  $t_1 \in \mathcal{U}[t_2 + \delta, t_2 + \Delta]$

$\mathbf{x}_\pi \leftarrow g(\theta, \pi)$

$\mathbf{x}_{t_1} \leftarrow \mathbf{x}_\pi + \sigma_{t_1}\epsilon^*$

$\mathbf{d}_i \leftarrow (\mathbf{x}_{t_1} - D_\phi(\mathbf{x}_{t_1}, t_1, y)) / \sigma_{t_1}$

$\hat{\mathbf{x}}_{t_2} \leftarrow \mathbf{x}_{t_1} + (\sigma_{t_2} - \sigma_{t_1})\mathbf{d}_i$

$\hat{\mathbf{x}}_0 \leftarrow \mathbf{x}_\pi + \text{sg}[\sigma_{t_1}(\epsilon^* - \mathbf{d}_i)]$

$\mathcal{L}_{\text{CDS}}(\theta; \pi) \leftarrow \lambda(t_2)\|\hat{\mathbf{x}}_0 - \text{sg}[D_\phi(\hat{\mathbf{x}}_{t_2}, t_2, y)]\|_2^2$

$\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{CDS}}(\theta; \pi)$

**end**

---

where  $\hat{\mathbf{x}}^0 = \mathbf{0}$  is to ensure  $\hat{\mathbf{x}}_T = \hat{\mathbf{x}}^0 + \sigma_T \epsilon_T \sim \mathcal{N}(\mathbf{0}, \sigma_T^2 \mathbf{I})$ , and time step schedule  $\{t_i\}$  satisfies  $T = t_1 > t_2 > \dots > t_N = 0$ . To build a connection between SDE and SDS, in Eq. (5), we follow SDS to approximate the score function  $\nabla \log p_t(\mathbf{x})$  in vanilla SDE with a score network  $D_\phi(\mathbf{x}_t, t)$ . Then by iteratively running Eq. (5) from  $t_1$  to  $t_N$ , one can eventually compute the desired SDE solution in expectation, e.g., a real sample  $\hat{\mathbf{x}}_N \sim p_{\text{data}}(\mathbf{x})$  if the network  $D_\phi(\cdot)$  is a well-trained diffusion model like Stable Diffusion [34].

On the other hand, by fixing the camera pose  $\pi$ , by fixing the camera pose  $\pi$ , for a rendered image  $\mathbf{x}_\pi$  by a learnable 3D model  $\theta$ , the optimization process of SDS introduced in Sec. 3 can be formulated as:

$$\begin{aligned} \mathbf{x}_{t_i} &= \mathbf{x}_\pi^{i-1} + \sigma_{t_i} \epsilon_{t_i} \text{ with } \epsilon_{t_i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \\ \mathbf{x}_\pi^i &= g(\theta^i, \pi) \text{ with } \theta^i = \arg \min_{\theta} \|g(\theta, \pi) - D_\phi(\mathbf{x}_{t_i}, t_i)\|, \end{aligned} \quad (6)$$

where  $\mathbf{x}_\pi^0 = g(\theta^0, \pi)$  in which  $\theta^0$  denotes the randomly initialized 3D model [20, 31], and  $g(\cdot)$  is a differentiable renderer [6, 26]. Compared the stochastic sampling process in Eq. (5) with SDS process in Eq. (6), one can observe that if for each iteration  $i$ , one can ideally optimize 3D model  $\theta^i$  so that  $\|g(\theta^i, \pi) - D_\phi(\mathbf{x}_{t_i}, t_i)\| = 0$ , then one can have

$$\mathbf{x}_\pi^i = g(\theta^i, \pi) = D_\phi(\mathbf{x}_{t_i}, t_i). \quad (7)$$

In this case, the SDS optimization process becomes exactly the same as the stochastic sampling process with  $\mathbf{x}_\pi^i$  replaced by  $\hat{\mathbf{x}}^i$ .

However, as illustrated in Fig. 2, sampling along the SDE solution trajectory according to Eq. (5) results in an unpredictable and highly variable next point  $\hat{\mathbf{x}}^i$ , which does not

guarantee the correct direction. This issue also extends to the SDS optimization process, which is equivalent to the SDE trajectory in Eq. (6) when the 3D model is ideally trained in each iteration (i.e., Eq. (7) holds). Consequently, such inherent randomness in SDS leads to less accurate and reliable guidance throughout all training iterations. This could also help explain why SDS is so vulnerable and often suffers from geometry collapse and poor fine-grained texture as observed in many works [43, 47, 55].

## 4.2. Consistency Distillation Sampling

**3D Deterministic Sampling.** Given the stochastic and unpredictable nature of SDS, we are motivated to explore the potential of the ODE deterministic process which can provide consistent and more accurate guidance than SDE for 3D generation as shown by Fig. 2. We start by focusing on the ODE sampling process for a 3D model  $\theta$ :

$$d\theta = -\dot{\sigma}_t \sigma_t \nabla \log p_t(\theta) dt, \quad (8)$$

where  $\theta$  is randomly initialized according to a certain distribution. Following Poole et al. [31] and Wang et al. [45], one can derive the 3D score function  $\nabla_\theta \log p_t(\theta)$  from the 2D score function using the chain rule:

$$\nabla_\theta \log p_t(\theta) = \mathbb{E}_\pi \left[ \nabla_{\mathbf{x}_\pi} \log p_t(\mathbf{x}_\pi) \frac{\partial \mathbf{x}_\pi}{\partial \theta} \right], \quad (9)$$

where the 2D score function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$  can be estimated as  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}) = (D_\phi(\mathbf{x}, t) - \mathbf{x}) / \sigma_t^2$  by a pre-trained diffusion model  $D_\phi(\mathbf{x}, t)$ . Therefore, the key to generating a satisfactory 3D model is to accurately perform the 3D ODE sampling in Eq. (8) using the pre-trained diffusion model.

Unfortunately, unlike the forward SDE process in which a noisy sample can be easily sampled from the perturbation kernel by  $\mathbf{x}_t \sim p_t(\mathbf{x}_t | \mathbf{x}_\pi)$  in Eq. (1), the forward ODE requires iterative simulation of the ODE flow, such as DDIM inversion [27] which is complex and time-consuming. This makes the approximation of the ODE flow with conventional SDS less efficient and is often impractical. Thus, directly applying SDS loss to the ODE flow is practically prohibited.

Inspired by recent advances in diffusion model distillation techniques that facilitate approximation of this deterministic flow without extensive simulation [42], we develop a simple yet effective Consistency Distillation Sampling loss (CDS) tailored for general text-to-3D generation tasks. Further detailed discussions can be found in Appendix A.

**Optimization objective.** We aim to enforce the optimization process of the 3D model  $\theta$  to match the deterministic flow between two *adjacent* ODE sampling steps. Specifically, we always use a *fixed* Gaussian noise  $\epsilon^*$  to perturb the sample, analogous to setting a fixed starting point at the final diffusion time step. This approach ensures a consistent perturbation in all iterations, similar to the technique

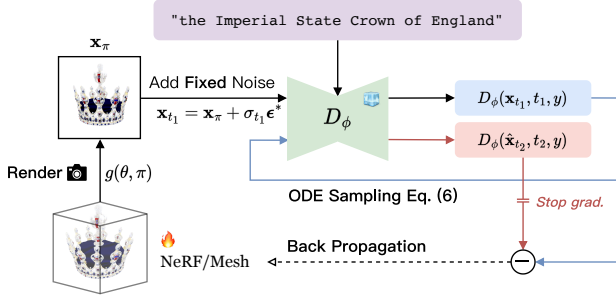


Figure 3. Overview of CDS. In each training iteration, the rendered image is perturbed by a fixed noise and then served as a start point of the deterministic flow for computing the CDS loss.

used in Consistency Training [42]. Next, we optimize  $\theta$  by minimizing the following Consistency Distillation Sampling (CDS) loss:

$$\mathbb{E}_{\pi} [\lambda(t_2) \|D_{\phi}(\mathbf{x}_{t_1}, t_1, y) - \text{sg}(D_{\phi}(\hat{\mathbf{x}}_{t_2}, t_2, y))\|_2^2], \quad (10)$$

where  $\text{sg}(\cdot)$  is a stop-gradient operator,  $t_1 > t_2$  are two adjacent diffusion time steps,  $\mathbf{x}_{t_1} = \mathbf{x}_{\pi} + \sigma_{t_1} \epsilon^*$ , and  $\hat{\mathbf{x}}_{t_2}$  is a less noisy sample derived from deterministic sampling by running one discretization step of a numerical ODE solver from  $\mathbf{x}_{t_1}$ . Particularly, we adopt the Euler solver to compute  $\hat{\mathbf{x}}_{t_2}$  by:

$$\hat{\mathbf{x}}_{t_2} = \mathbf{x}_{t_1} + \frac{\sigma_{t_2} - \sigma_{t_1}}{\sigma_{t_1}} (\mathbf{x}_{t_1} - D_{\phi}(\mathbf{x}_{t_1}, t_1, y)). \quad (11)$$

In practice, we follow Poole et al. [31] and reparameterize the first component in Eq. (10) to skip the CDS gradient directly to  $\mathbf{x}_{\pi}$  and  $\theta$  without computing the U-Net Jacobian. **Time step schedule.** As our target is to match the probability flow ODE of the reverse sampling process, we follow the conventional DMs [16, 41] and set the time steps to decrease monotonically along with the training iteration of the 3D models. This approach redefines our 3D generation process more as a deterministic sampling rather than a mere training process as previous SDS-based approaches [31, 45], thus allowing us to take full advantage of deterministic sampling prior.

Specifically, we define the time step schedule of  $t_2$  according to the current training iteration:

$$t_2 := t_{\max} - (t_{\max} - t_{\min}) \sqrt{i/N}, \quad (12)$$

where  $i$  and  $N$  denotes the current iteration and total iteration, respectively. For the initial time step  $t_1$  which indicates the perturbation level to the rendered image, we empirically uniformly sample it within  $[t_2 + \delta, t_2 + \Delta]$ , which is different from the predetermined time step schedule in Consistency Distillation [42]. This is because we empirically find that the random sampled time step  $t_1$  within a *small* interval collaborated with the deterministic anchor  $t_2$

exhibits self-calibration behaviors, which can actively correct the cumulative error made in earlier steps and alleviate issues such as floaters and Janus faces [14, 38]. We delve deeper into this phenomenon in Sec. 5.4. For more clarity, we summarize our entire text-to-3D generation procedure with the proposed CDS in Algorithm 1.

**Justification.** In the following, we offer a theoretical justification to demonstrate that, upon achieving convergence, our Consistency Distillation Sampling is capable of generating a high-fidelity 3D model.

**Theorem 1.** Assume that the diffusion model  $D_{\phi}(\cdot)$  satisfies the Lipschitz condition. Define  $\Delta := \sup |t_1 - t_2|$ . For any given camera pose  $\pi$ , if convergence is achieved according to Eq. (10), then there exists a corresponding real image  $\mathbf{x}^* \sim p_{\text{data}}(\mathbf{x})$  such that

$$\|\mathbf{x}_{\pi} - \mathbf{x}^*\|_2 = \mathcal{O}(\Delta), \quad (13)$$

where  $\mathbf{x}_{\pi} = g(\theta, \pi)$  denotes the rendered image for pose  $\pi$ .

*Proof.* The proof is based on the truncation error of the Euler solver. We provide the full proof in Appendix C.  $\square$

For a 3D model optimized using the CDS, Theorem 1 guarantees that images rendered from any viewpoint of this model are realistic and closely align with the corresponding real-world scenes.

### 4.3. Consistent3D

Now we are ready to introduce our proposed Consistent3D. As illustrated in Fig. 3, we present a clear design space for our Consistent3D generation framework using our proposed Consistency Distillation Sampling (CDS).

Following previous work [20, 47], Consistent3D is a coarse-to-fine approach consisting of two stages. Specifically, in the coarse stage, we optimize a low-resolution Neural Radiance Field (NeRF) [2, 28]. For the refinement stage, we further optimize a high-resolution textured 3D mesh [37] from the neural field initialization converting from the coarse stage. For these two stages, we always use our proposed CDS.

**NeRF Optimization Stage.** We adopt multi-resolution hash grids, Instant NGP [28] to parameterize the scene by density and color with MLPs, which improves training and rendering efficiency. We follow Magic3D on density bias initialization, camera and light augmentation. In addition, we use orientation loss [31] and 2D normal smooth loss [24]. At this stage, we render  $64 \times 64$  images and use our proposed CDS as guidance. We set  $t_{\max} = 0.7$ ,  $t_{\min} = 0.1$ ,  $\delta = 0.1$ , and  $\Delta = 0.2$ .

**Mesh Refinement Stage.** We convert the neural field into Signed Distance Field (SDF) by subtracting it with a fixed threshold and then optimizing a high-resolution

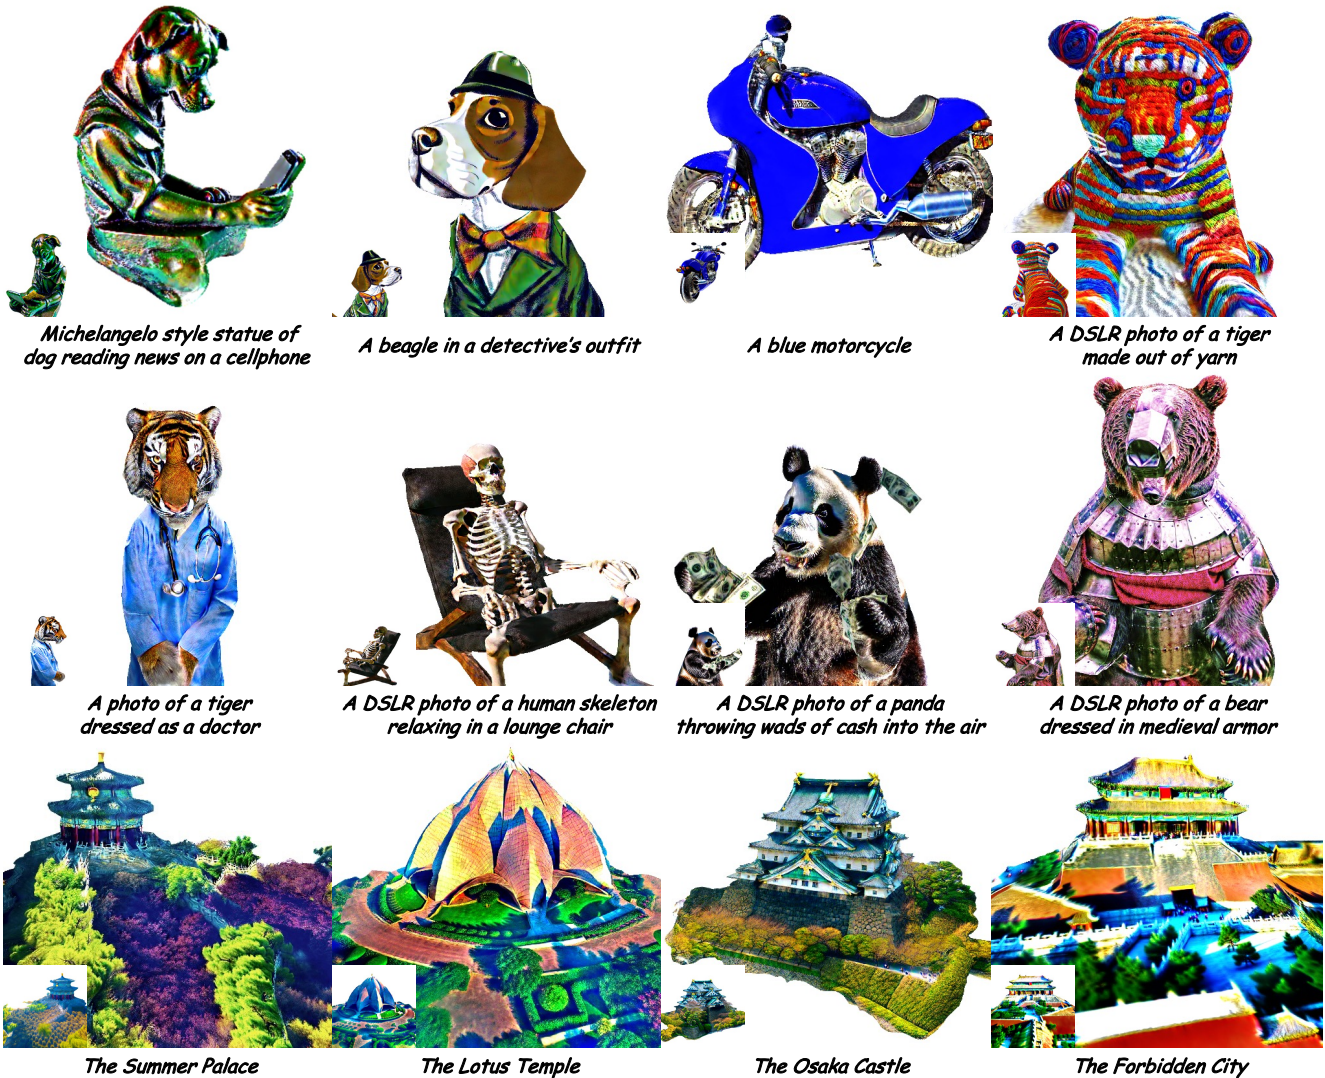


Figure 4. Consistent3D can generate diverse and high-fidelity objects or large-scale scenes highly correlated with the given text prompts.

DMTet [37]. We also initialize the volume texture field directly with the color field from the coarse stage. In addition, we use normal consistency loss and Laplacian smoothness loss. In the refinement stage, we render  $512 \times 512$  images and set  $t_{\max} = 0.5$ ,  $t_{\min} = 0.02$ ,  $\delta = 0.1$ , and  $\Delta = 0.1$ .

## 5. Experiment

### 5.1. Implementation Details

Consistent3D is implemented in PyTorch with a single NVIDIA A100 GPU based on *threestudio* [9] with Stable Diffusion v2.1 [34]. We use the Adan [49] optimizer with a learning rate of 0.05 for grid encoder and 0.005 for other parameters, and a weight decay of  $2 \times 10^{-8}$ . Further implementation details are provided in Appendix B.1.

### 5.2. Text-guided 3D Generation

As illustrated in Fig. 4, our Consistent3D demonstrates versatility in generating high-fidelity 3D objects. Its generated images are not only realistic but also maintain consistency from various viewpoints. Furthermore, it is capable of generating large-scale scenes in  $360^\circ$  with remarkable detail. See more qualitative results in Appendix B.2.

### 5.3. Comparison with the State-of-The-Art

In this section, we present comprehensive qualitative and quantitative experiments to evaluate the efficacy of our Consistent3D framework in text-to-3D generation. We compare our generation performance with *DreamFusion* [31], *Magic3D* [20], and *ProlificDreamer* [47]. For a fair comparison, we use the implementations of all the baseline methods from the open-source repository *threestudio* [9].

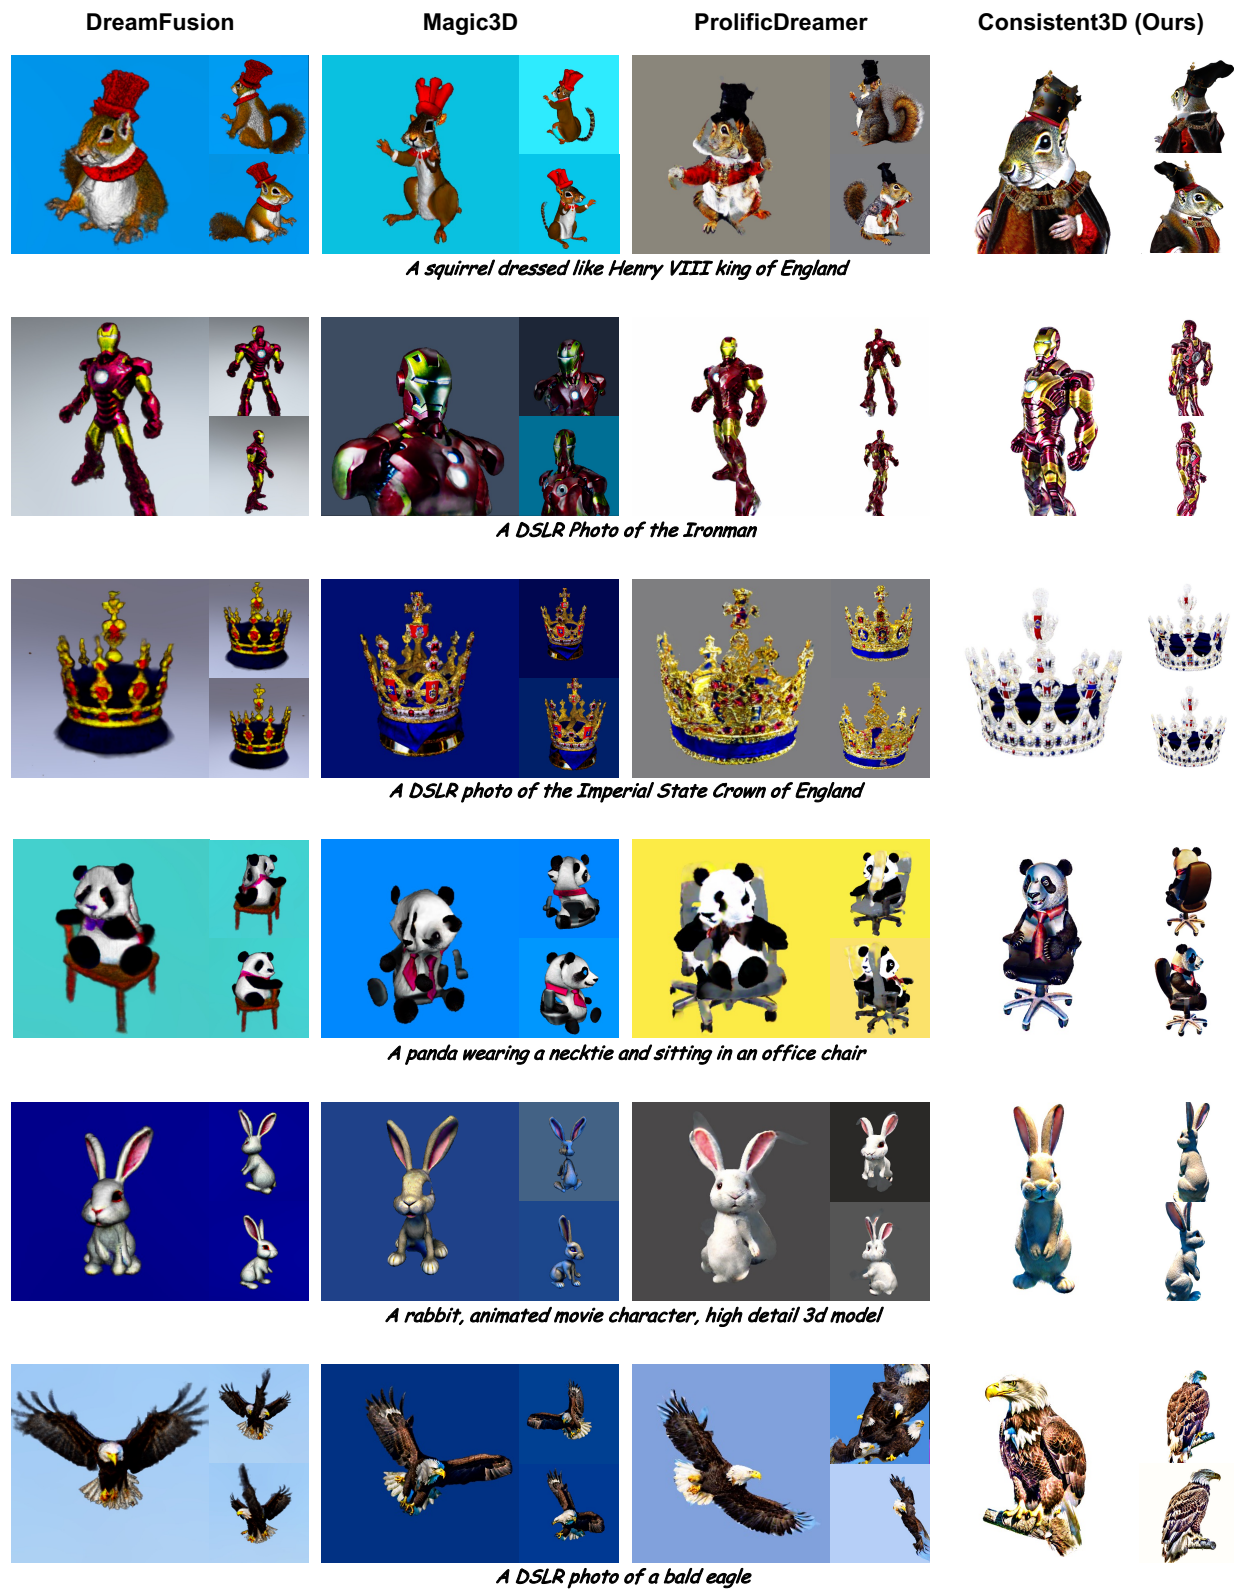


Figure 5. **Qualitative Comparisons of Text-to-3D Generation.** Our approach yields results with enhanced fidelity and more robust geometry.



Figure 6. **Ablation study** of component-wise contribution of Consistent3D: (a) random time step schedule; (b) predetermined time step schedule; (c) random noise in each iteration; (d) our proposed configuration.

Method	Loss	CLIP-R
DreamFusion [31]	SDS	0.310
Magic3D [20]	SDS	0.311
ProlificDreamer [47]	VSD	0.336
Consistent3D (Ours)	CDS	<b>0.348</b>

Table 1. **Quantitative Comparisons** of CLIP R-Precision. Scores were averaged from 40 prompts in the DreamFusion gallery.

**Qualitative Results.** In Fig. 5, we provide qualitative comparisons with the baseline methods. Our approach exhibits more photorealistic details and geometry than both the SDS-based approaches like DreamFusion and Magic3D and the VSD-based approach ProlificDreamer. This improvement mainly stems from our Consistency Distillation Sampling (CDS) which effectively leverages the full potential of large-scale diffusion models by accurately distilling deterministic sampling priors into the 3D model.

**Quantitative Results.** In Tab. 1, we report the results of CLIP R-Precision [32] for 3D objects generated using 40 randomly selected text prompts from the DreamFusion gallery. See more details in Appendix B.1. Each 3D object is rendered from 120 viewpoints with a uniform azimuth angle. The CLIP R-Precision score is computed by averaging the similarity scores between each rendered view and the corresponding text prompt. Additionally, we also conduct a head-to-head user study in Tab. 3. Our quantitative analysis shows the superior performance of our method.

#### 5.4. Ablation Study

We present an ablation study to evaluate the effects of various components in our approach in Fig. 6 and Tab. 2. We conduct experiments with the following configurations: (a) a random time step schedule in DreamFusion [31]; (b) a predetermined time step schedule from Consistency Distillation [42]; (c) varied random noise in each iteration; and

(d) our proposed method incorporating all components. The results in Fig. 6(a) reveal that a random time step schedule detrimentally affects both geometry and texture modeling, since it disrupts established rules of sampling process. Fig. 6(b) suggests that a predetermined time-step schedule is suboptimal for optimization-based methods, since gradient descent does not ensure monotonic optimization progress. This implies that minor randomness helps to accommodate these variations. Fig. 6(c) shows that fixed noise aids in better convergence by providing a consistent perturbation in each iteration.

## 6. Conclusion

In this work, we first connect Score Distillation Sampling (SDS), a leading text-to-3D generation approach, with the solution trajectory sampling of a stochastic differential equation (SDE). This connection helps us to understand the vulnerability in SDS, since the randomness in SDE sampling often provides a highly diverse sample, which is not always less noisy, and could guide the 3D model in the wrong direction. Then motivated by the fact that an ordinary differential equation (ODE) of an SDE can provide a deterministic and consistent sampling trajectory, we propose a novel and effective “Consistent3D” by designing a consistency distillation sampling loss to distill the deterministic sampling prior into a 3D model for text-to-3D generation. Extensive experimental results show that our Consistent3D surpasses state-of-the-art methods in generating high-fidelity and diverse 3D objects and large-scale scenes.

## Acknowledgement

Pan Zhou was supported by the Singapore Ministry of Education (MOE) Academic Research Fund (AcRF) Tier 1 grant.



## References

- [1] Kendall Atkinson. *An Introduction to Numerical Analysis*. John Wiley & Sons, 1991. 3
- [2] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 5
- [3] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 2, 3
- [4] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 3
- [5] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-based generative modeling with critically-damped langevin diffusion. In *International Conference on Learning Representations*, 2021. 3
- [6] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14346–14355, 2021. 4
- [7] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as plug-and-play priors. In *Advances in Neural Information Processing Systems*, 2022. 2
- [8] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10696–10706, 2022. 3
- [9] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023. 6
- [10] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 2
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, pages 6840–6851. Curran Associates, Inc., 2020. 2, 3
- [12] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 3
- [13] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. 3
- [14] Susung Hong, Donghoon Ahn, and Seungryong Kim. Debiasing scores and prompts of 2d diffusion for robust text-to-3d generation. *arXiv preprint arXiv:2303.15413*, 2023. 2, 5
- [15] Heewoo Jun and Alex Nichol. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*, 2023. 3
- [16] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. In *Advances in Neural Information Processing Systems*, 2022. 3, 5, 12
- [17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (ToG)*, 42(4):1–14, 2023. 15
- [18] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. *arXiv preprint arXiv:2303.13439*, 2023. 3
- [19] Weiyu Li, Rui Chen, Xuelin Chen, and Ping Tan. Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. *arXiv preprint arXiv:2310.02596*, 2023. 3
- [20] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023. 2, 3, 4, 5, 6, 8, 12
- [21] Luping Liu, Yi Ren, Zhijie Lin, and Zhou Zhao. Pseudo numerical methods for diffusion models on manifolds. In *International Conference on Learning Representations*, 2021. 3
- [22] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008*, 2023. 3
- [23] Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems*, 2022. 3
- [24] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Realfusion: 360deg reconstruction of any object from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8446–8455, 2023. 5
- [25] Chenlin Meng, Ruiqi Gao, Diederik P Kingma, Stefano Ermon, Jonathan Ho, and Tim Salimans. On distillation of guided diffusion models. *arXiv preprint arXiv:2210.03142*, 2022. 3
- [26] 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. *Communications of the ACM*, 65(1):99–106, 2021. 2, 4
- [27] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023. 4
- [28] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 5
- [29] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751*, 2022. 3
- [30] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 2
- [31] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 2, 3, 4, 5, 6, 8, 12
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 8, 12
- [33] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 2, 3, 4, 6
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 2
- [36] Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*, 2022. 3
- [37] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems*, 34:6087–6101, 2021. 2, 5, 6
- [38] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*, 2023. 2, 3, 5
- [39] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021. 3, 15
- [40] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 2
- [41] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. 2, 3, 5, 12
- [42] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. 3, 4, 5, 8
- [43] Jiayang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 2, 3, 4, 15
- [44] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhchikov, Joshua B Tenenbaum, Frédo Durand, William T Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *arXiv preprint arXiv:2306.11719*, 2023. 3
- [45] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023. 2, 3, 4, 5
- [46] Yiming Wang, Qin Han, Marc Habermann, Kostas Daniilidis, Christian Theobalt, and Lingjie Liu. Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3295–3306, 2023. 3
- [47] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213*, 2023. 2, 3, 4, 5, 6, 8, 12
- [48] Zike Wu, Pan Zhou, Kenji Kawaguchi, and Hanwang Zhang. Fast diffusion model. *arXiv preprint arXiv:2306.06991*, 2023. 2, 3
- [49] Xingyu Xie, Pan Zhou, Huan Li, Zhouchen Lin, and Shuicheng Yan. Adan: Adaptive nesterov momentum algorithm for faster optimizing deep models. *arXiv preprint arXiv:2208.06677*, 2022. 6
- [50] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022. 3
- [51] Taoran Yi, Jiemin Fang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu, Qi Tian, and Xinggang Wang. Gaussian-dreamer: Fast generation from text to 3d gaussian splatting with point cloud priors. *arXiv preprint arXiv:2310.08529*, 2023. 3
- [52] Qinsheng Zhang, Molei Tao, and Yongxin Chen. gDDIM: Generalized denoising diffusion implicit models. *arXiv preprint arXiv:2206.05564*, 2022. 3
- [53] Qinsheng Zhang, Jiaming Song, and Yongxin Chen. Improved order analysis and design of exponential integrator for diffusion models sampling. *arXiv preprint arXiv:2308.02157*, 2023. 3
- [54] Minda Zhao, Chaoyi Zhao, Xinyue Liang, Lincheng Li, Zeng Zhao, Zhipeng Hu, Changjie Fan, and Xin

Yu. Efficientdreamer: High-fidelity and robust 3d creation via orthogonal-view diffusion prior. *arXiv preprint arXiv:2308.13223*, 2023. [3](#)

- [55] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023. [2](#), [3](#), [4](#)