Connecting Consistency Distillation to Score Distillation for Text-to-3D Generation

Zongrui Li^{1,2} [†]⁽⁰⁾, Minghui Hu² [†]⁽⁰⁾, Qian Zheng^{3,4} ^[2]⁽⁰⁾, and Xudong Jiang²⁽⁰⁾

¹ Rapid-Rich Object Search (ROSE) Lab, Nanyang Technological University
 ² School of Electrical and Electronic Engineering, Nanyang Technological University
 ³ College of Computer Science and Technology, Zhejiang University

⁴ The State Key Lab of Brain-Machine Intelligence, Zhejiang University

{ZONGRUI001, e200008}@e.ntu.edu.sg, qianzheng@zju.edu.cn, EXDJiang@ntu.edu.sg

Abstract. Although recent advancements in text-to-3D generation have significantly improved generation quality, issues like limited level of detail and low fidelity still persist, which requires further improvement. To understand the essence of those issues, we thoroughly analyze current score distillation methods by connecting theories of consistency distillation to score distillation. Based on the insights acquired through analvsis, we propose an optimization framework, Guided Consistency Sampling (GCS), integrated with 3D Gaussian Splatting (3DGS) to alleviate those issues. Additionally, we have observed the persistent oversaturation in the rendered views of generated 3D assets. From experiments, we find that it is caused by unwanted accumulated brightness in 3DGS during optimization. To mitigate this issue, we introduce a Brightness-Equalized Generation (BEG) scheme in 3DGS rendering. Experimental results demonstrate that our approach generates 3D assets with more details and higher fidelity than state-of-the-art methods. The codes are released at https://github.com/LMozart/ECCV2024-GCS-BEG.

Keywords: Text-to-3D Generation · Score Distillation Sampling · Consistency Model

1 Introduction

Text-to-3D generation [4, 11, 17, 27, 29, 43, 44] has gained substantial attention due to its great potential and indispensable role in many applications, such as gaming, filmmaking, and architecture. While an end-to-end text-to-3D generative model [11, 27] is often difficult to train and lacks versatility due to the limited 3D assets in the training set, distillation 3D assets from a well-trained 2D generative model (*e.g.*, diffusion model (DM) [9,30]) has become a popular option due to its data-free feature. A typical work, DreamFusion [29], transfers knowledge from the pre-trained diffusion model to a learnable 3D representation through Score Distillation Sampling (SDS), generating text-aligned 3D assets.

^{\dagger} Equal contribution. \boxtimes Corresponding author.



Fig. 1: Text-to-3D generation results of the proposed Guided Consistency Sampling (GCS) and Brightness-equalized Generation. Each 3D asset is distilled from a pre-trained 2D diffusion model and demonstrated with three different views. The results below the dotted line are generated by the fine-tuned diffusion models.

Nevertheless, as highlighted in many subsequent works [17, 39, 43–45], SDS suffers from a low level of detail in generation results. Such an issue stems from the poor generalization ability of the distillation method [43] and the inherent randomness in the sampling process [17, 44]. Although work like [43] is proposed to enhance the generalization ability, it necessitates fine-tuning the diffusion model during training, dramatically extending the training time. By contrast, more recent works [17,44] aim to mitigate the randomness in the sampling process by introducing Probability Flow Ordinary Differential Equations (PF-ODEs) [38], and speed up the distillation process through 3D Gaussian Splatting (3DGS) [13], resulting in cost-efficient, high-quality 3D generation.

Despite the effectiveness of score distillation methods, issues like limited detail and low fidelity remain in the generated 3D asset. To reveal the essence of those issues, we conduct an in-depth analysis to connect consistency distillation [14, 37, 51], a method that distills information from a pre-trained DM through PF-ODEs, to score distillation. Such a novel view helps us identify potential causes of those issues in the score distillation process: 1) the inherent distillation errors of PF-ODEs, 2) ineffective conditional guidance accounting for classifier-free guidance (CFG) effects [10], and 3) lack of constraints in pixel domain to avoid out-of-distribution problem, requiring further improvements.

Except for identified issues in score distillation, there is another noteworthy concern in the 3D rendering process, *i.e.*, the over-saturated views for generated 3D asset [17, 43, 44]. While the cause of such an issue is correlated with types of 3D representation, this paper particularly focuses on one (*i.e.*, 3DGS) due to its remarkable reduction in optimization time. Through experiments, we observe



Fig. 2: An overview of the proposed GCS. We first initialize the 3D representation via the pre-trained 3D generator. For each training epoch, we randomly render a batch of views \mathbf{x}_{π} and diffuse them to $\mathbf{x}_{\{\pi,e\}}$ with a fixed noise ϵ^* . We then apply the ODE diffusion process to gradually add noise to the $\mathbf{x}_{\{\pi,e\}}$ and transfer it to $\bar{\mathbf{x}}_{\{\pi,e\to s\}}$ and $\bar{\mathbf{x}}_{\{\pi,s\to t\}}$. In the denoising path, we conduct conditional and unconditional denoising steps, as shown in the figure. Eventually, we calculate the \mathcal{L}_{GCS} (Eq. 17) to update the parameters of 3D representation (3DGS). Note that we add '*' on $\bar{\mathbf{x}}_{\{\pi,e\to 0;y\}}$ to indicate that it is obtained from different sampling trajectories.

that the brightness in generated 2D views accumulates to the next epoch during distillation, eventually leading to over-saturation.

To achieve better performance and overcome the aforementioned issues, this paper proposes **Guided Consistency Sampling** (GCS) with **Brightnessequalized Generation** (BEG) for 3DGS-based text-to-3D generation (Fig. 2). The proposed GCS includes three components to solve the identified issues in score distillation: a compact consistency loss reduces distillation errors, a conditional guidance loss provides more effective conditional guidance, and a constraint on pixel domain enhances the fidelity of the 3D assets. Additionally, the proposed BEG helps regularize the accumulation of brightness in 3DGS, which significantly alleviates over-saturation issues. Examples of generated 3D assets can be found in Fig. 1. In summary, our contributions are as follows:

- 1. We identify three problems in the PF-ODEs-based score distillation method by connecting consistency distillation to score distillation.
- 2. We propose **Guided Consistency Sampling**, an improved score distillation method, to enhance the details and fidelity of the generated 3D asset.
- 3. We find an accumulated brightness issue in the 3DGS-based rendering process that causes over-saturation and propose **Brightness-equalized Generation** to alleviate it in the 3DGS-based rendering process.

2 Related Work

Text-to-3D Generation by Score Distillation. SDS, introduced in Dream-Fusion [29], and concurrent Score Jacobian Chain (SJC) proposed in [42], have been regarded as milestone works in distilling information from a 2D pre-trained diffusion models to differentiable 3D representations, which inspiring many subsequent Text-to-3D generation methods [4,5,7,12,18,19,24,33,35,39–41,43,44, 47–49,53] to improve the quality of generated 3D assets. Generally, they either endow the DM with the pose-aware generative ability [19, 24, 35, 41, 48, 49] to alleviate the Janus problem [1] for a more consistent 3D asset, or redesign the distillation process and optimization objectives [4, 5, 12, 18, 43, 44, 47, 53] to reduce the distillation error from the 2D pre-trained diffusion model for a higher level of details. While the performance of the former methods may be restricted by the inherent limitations of SDS [17], the latter methods dedicated to overcoming these limitations have recently gained increasing attention. For instance, ProlificDreamer [43] proposes sampling from the parameter distribution instead of seeking the optimal solution for high-fidelity generation. CSD [10] introduces insights into the effects of SDS's components and proposes improved strategies. More recent works like [17, 44] introduce PF-ODEs into the SDS for a stable. deterministic distillation process. After thoroughly investigating recent works, we recognize a connection between consistency distillation and score distillation, which motivates us to analyze and redesign the optimization objectives in view of consistency distillation. Although concurrent work [44] also highlighted the effectiveness of text-to-3D generation using consistency distillation, our analysis delves deeper into such connection from a broader perspective. We also compare with the findings of [44] in Tab. 1. Another concurrent study [23] established a relationship between SDS and bridge matching [6,15], leading to an improvement similar to that seen in our method. Guided by acquired insights, we propose GCS that integrates techniques in CM to improve generation quality and gain state-of-the-art performance.

Differentiable 3D Representations, such as NeRF [17, 25, 29, 43], 3DGS [7, 13, 39], or learnable intrinsic that integrated with differentiable geometry [4, 34] and textures [2, 3, 8, 16], can be optimized by minimizing the loss between rendered and ground truth images, establishing a connection between 3D and 2D representations. Text-to-3D by distillation relies heavily on differentiable 3D representations. While NeRF is well-known to be slow to optimize, we use 3DGS considering its efficiency and comparable performance with NeRF.

3 Preliminary

Diffusion Models (DM) [9,31] apply a forward process to perturb data samples \mathbf{x}_0 drawn from a data distribution $p_{\text{data}}(\mathbf{x})$ by the Gaussian noise $\mathcal{N}(\mathbf{0}, \sigma_t^2 \mathbf{I})$ with time-dependent variance σ_t , where $t \in [0, T]$. As such, the distribution of noisy samples \mathbf{x}_t at time-step t is:

$$p_t\left(\mathbf{x}_t \mid \mathbf{x}_0\right) = \mathcal{N}\left(\mathbf{x}_t; \mathbf{x}_0, \sigma_t^2 \mathbf{I}\right),\tag{1}$$

which can be re-parameterized to $\mathbf{x}_t = \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}_t$, where $\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. A reverse process can denoise the noisy samples from \mathbf{x}_t to \mathbf{x}_0 , expressed as variational inference of Markov processes [9]. For all diffusion processes, there exists a

corresponding deterministic process with trajectory sharing the same marginal distribution, known as PF-ODE [38]:

PF-ODE:
$$d\mathbf{x}_t = \left[\mathbf{f}(\mathbf{x}_t, t) - \frac{1}{2}g(t)^2 \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) \right] dt,$$
 (2)

where $\mathbf{f}(\mathbf{x}_t, t)$ is the drift coefficients, dt is the infinitesimal negative timestep, $\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t)$ is the score function of $p_t(\mathbf{x}_t)$, estimated by a neural network $\epsilon_{\theta}(\mathbf{x}_t, t)$. An ODE solver (*e.g.*, DDIM [36], DPM-Solver [20], etc.) can be used to derive the solution of ODE, identical as a sampling process. The sampling trajectory of ODE is deterministic, and the randomness comes from the initial value, which is more stable than a random sampling process described by the inverse Stochastic Differentiable Equation (SDE) [38].

Consistency Model (CM) [37] is proposed to facilitate a single-step or low number of function evaluations (NFEs) [28] generation by distilling knowledge from pre-trained DM models. It defines a one-step generator $f_{\theta}(.;.)$ with trainable parameters θ that directly predicts the denoised image \mathbf{x}_0 given t and \mathbf{x}_t , constrained by some boundary conditions [37]. $f_{\theta}(.;.)$ is trained by minimizing consistency distillation loss [37] defined as:

$$\mathcal{L}_{\mathrm{CD}}\left(\boldsymbol{\theta},\boldsymbol{\theta}^{-}\right) = \mathbb{E}\left[\omega(t)\left\|\boldsymbol{f}_{\boldsymbol{\theta}}\left(\mathbf{x}_{t_{n+1}};t_{n+1}\right) - \boldsymbol{f}_{\boldsymbol{\theta}^{-}}\left(\bar{\mathbf{x}}_{\left\{t_{n+1}\to t_{n}\right\}};t_{n}\right)\right\|_{2}^{2}\right],\quad(3)$$

where, $0 = t_1 < t_2 \cdots < t_N = T$, $\bar{\mathbf{x}}_{\{t_{n+1} \to t_n\}}$ is calculated given ODE solver $\Phi(.)$ as $\bar{\mathbf{x}}_{\{t_{n+1} \to t_n\}} = \Phi(\mathbf{x}_{t_{n+1}}; t_{n+1}, t_n)$, θ^- is updated during training process through an exponential moving average (EMA) strategy [38]. The ultimate goal of CM is to maintain the self-consistency condition along the trajectory $\{\mathbf{x}_t\}_{t \in [0,T]}$, satisfying,

$$\boldsymbol{f}(\mathbf{x}_t;t) = \boldsymbol{f}(\mathbf{x}_{t'};t') \quad \forall t, t' \in [0,T].$$
(4)

Subsequent works of CM [10,14,51] integrate Classifier-free guidance (CFG) [10] or jumps along PF-ODE trajectory into the CM for high-fidelity samples given low or high NFEs, respectively. Inspired by those works, we seek to improve the performance of SDS further in view of CM given their connection.

4 Methodology

In Sec. 4.1, we analyze current advances [17,29,43,44,47] in a unified framework obtained by connecting consistency distillation [14,37,51] to score distillation. We revisited and identified three problems in the current PF-ODEs-based score distillation from the perspective of consistency distillation. We further proposed our solution, *i.e.*, Guided Consistency Sampling (GCS), and explained in detail in Sec. 4.2. While GCS is integrated with 3DGS, we observe the accumulated brightness causing over-saturation in 3DGS, as explained in Sec. 4.3. We propose Brightness Equalized Generation (BEG) to alleviate this issue.

6 Z. Li et al.

 Table 1: Specific design space employed by proposed GCS and selected 2D diffusionbased score distillation text-to-3D generation method.

	DreamFusion [29]	ProlificDreamer [43]	Consistent3D [44]	
3D Representation	NeRF / GS	NeRF	NeRF / GS	
CFG Weight	-	7.5 +	50+	
Objective	SDS (Eq. 7)	VSD (Eq. 9)	CDS (Eq. 15)	
Objective #1: Generator Loss				
term A	\mathbf{x}_{π}	$F_{\phi}(\mathbf{x}_{\{\pi,t\}};t,y)$	$F_{\theta}\left(\mathbf{x}_{\{\pi,t\}};t,y\right)$	
term B	$F_{\theta}\left(\mathbf{x}_{\{\pi,t\}};t,y\right)$	$F_{\theta}\left(\mathbf{x}_{\{\pi,t\}};t,y\right)$	$F_{\theta}\left(\bar{\mathbf{x}}_{\{\pi,t\to s;y\}};s,y\right)$	
	LucidDreamer [17]	Ours		
3D Representation	NeRF / GS	NeRF / GS	GS w/ BEG	
CFG Weight	7.5 +	100+	7.5+	
Objective	ISD (Eq. 13)	SDS w/ CFG (Eq. 10)	GCS (Eqs. 20, 22, 23)	
Objective #1: Generator Loss				
term A	-	\mathbf{x}_{π}	$G_{\theta}\left(\hat{\mathbf{x}}_{\{\pi,t\}};t,e,\emptyset\right)$	
term B	-	$F_{\theta}\left(\mathbf{x}_{\{\pi,t\}};t,y\right)$	$F_{\theta}\left(\bar{\mathbf{x}}_{\{\pi,t\to s;y\}};s,e,\emptyset\right)$	
Objective #2: Classifier Loss				
term A	$F_{\theta}(\hat{\mathbf{x}}_{\{\pi,s\}}; s, \emptyset)$	$F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,\emptyset)$	$F_{\theta}\left(\hat{\mathbf{x}}_{\{\pi,e\}}; e, \emptyset\right)$	
term B	$F_{\theta}\left(\hat{\mathbf{x}}_{\{\pi,t\}};t,y\right)$	$F_{\theta}\left(\mathbf{x}_{\{\pi,t\}};t,y\right)$	$F_{\theta}\left(G_{\theta}\left(\hat{\mathbf{x}}_{\{\pi,t\}};t,e,y\right);e,y\right)$	
Pixel Constraint		+SDS on pixel domain [53]	$+ \mathcal{L}_{CP}$ (Eq. 23)	

4.1 Connecting Consistency Distillation to Score Distillation

Through theoretical analysis, we find that consistency and score distillation share a similar optimization objective. While CM enforces a *self-consistency* of a PF-ODE sampling trajectory (Eq. 3), SDS and its variants (score distillation) facilitate a *cross-consistency*⁵ between the diffusion trajectory of the rendered view and the denoising trajectory (usually with text condition) of a noisy sample. This cross-consistency is maintained by minimizing the gap between a sample on the diffusion and denoising trajectory. The score distillation method diversifies in choosing those samples to form the optimization objectives (see Tab. 1 for a summary). In the following part, we unify the optimization objective of the score distillation method by sample (\mathbf{x}) instead of score (ϵ) to connect consistency distillation to score distillation based on their similarity.

In vanilla SDS, given a camera poses π , a 3D representation ξ can be projected to a specific 2D view, noted as $\mathbf{x}_{\pi} = g(\pi, \xi)$. The optimization objective of SDS on 3D parameters ξ is given as:

$$\min_{\xi} \mathcal{L}_{\text{SDS}}(\xi) = \mathbb{E}_{t,\pi} \left[\omega(t) \left\| \boldsymbol{\epsilon}_{\theta} \left(\mathbf{x}_{t}; t, y \right) - \boldsymbol{\epsilon} \right\|_{2}^{2} \right],$$
(5)

where $\epsilon_{\theta}(\mathbf{x}_t; t, y)$ is the predicted noise given \mathbf{x}_t , time-step t, and y; $\omega(t)$ is a t related weight function. As suggested in [32, 52, 53], SDS loss is equivalent to:

⁵ We define *cross-consistency* as an alignment between the diffusion trajectory of the rendered view and the text-conditioned sampling trajectory.

$$\mathcal{L}_{\text{SDS}}(\xi) = \mathbb{E}_{t,\pi} \left[\omega(t) \left\| \boldsymbol{\epsilon}_{\theta} \left(\mathbf{x}_{\{\pi,t\}}; t, y \right) - \boldsymbol{\epsilon} \right\|_{2}^{2} \right], \tag{6}$$

$$= \mathbb{E}_{t,\pi} \left\| c(t) \left\| \mathbf{x}_{\{\pi,t\}} - F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y) \right\|_{2}^{2} \right\|,$$
(7)

$$=: \mathcal{L}_{\text{Distill}}(\xi), \tag{8}$$

where $\mathbf{x}_{\{\pi,t\}} = \alpha_t \mathbf{x}_{\pi} + \sigma_t \epsilon$, $F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y) = \frac{\mathbf{x}_{\{\pi,t\}} - \sigma_t \epsilon_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y)}{\alpha_t}$, c(t) is another t related weight function. In CM [9,22], $f_{\theta}(.;.)$ can be parameterize as $F_{\theta}(.;.)$, which makes Eq. 3 similar to Eq. 7 in formula. As stated in [36,44], we interpret minimizing \mathcal{L}_{SDS} as facilitating a cross-consistency of the stochastic and deterministic trajectory. The variance of SDS injects different conditions on samples. For instance, replacing \mathbf{x}_{π} by samples from another fine-tuning diffusion model, noted as $F_{\phi}(\mathbf{x}_{\{\pi,t\}};t,y)$ in Eq. 7, we have Variational Score Distillation (VSD) [43] as:

$$\mathcal{L}_{\text{VSD}}(\xi) = \mathbb{E}_{t,\pi} \left[c(t) \left\| F_{\phi}(\mathbf{x}_{\{\pi,t\}};t,y) - F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y) \right\|_{2}^{2} \right],$$
(9)

which distills information by enforcing cross-consistency of the trajectory parameterized by ϕ and θ . In practice, CFG with a higher guidance weight w (e.g., w = 100) is imperative in SDS. Specifically, the expression of SDS with CFG in the form of distillation loss [10, 47] is:

$$\mathcal{L}_{\text{Distill}}^{\text{CFG}}(\xi) := \mathbb{E}_{t,\pi}[c(t)\| \underbrace{\left[\mathbf{x}_{\pi} - F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y)\right]}_{\text{generator loss}} + w \underbrace{\left[F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,\emptyset) - F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,y)\right]}_{\text{classifier loss}} \|_{2}^{2}\right].$$
(10)

Notably, we extend the definition in [47], noting the generator loss as a guidance to make \mathbf{x}_{π} more close to the prior distribution and the classifier loss as an update direction for \mathbf{x}_{π} to align with the text-condition. As studied in [47], the driving force of SDS is the classifier loss. By omitting the generator loss, they derive *Classifier Score Distillation* (CSD), given:

$$\mathcal{L}_{\text{CSD}}(\xi) := \mathbb{E}_{t,\pi}[c(t) \| [F_{\theta}(\mathbf{x}_{\{\pi,t\}}; t, \emptyset) - F_{\theta}(\mathbf{x}_{\{\pi,t\}}; t, y)] \|_{2}^{2}].$$
(11)

Additionally, as shown in [17], $\mathbf{x}_{\{\pi,t\}}$ could be replaced by the deterministic status derived from DDIM to improve the distillation quality:

$$\hat{\mathbf{x}}_{\{\pi,t\}} = \sum_{k=0}^{s} \alpha_{(k+\delta_k)} \left(\frac{1}{\alpha_k} \mathbf{x}_k - \frac{\sigma_{(k+\delta_k)}}{\alpha_{(k+\delta_k)}} \epsilon_{\theta}(\mathbf{x}_k; k, \emptyset) \right) + \sigma_{(k+\delta_k)} \epsilon_{\theta}(\mathbf{x}_k; k, \emptyset), \quad (12)$$

where t > s with $t = s + \delta_s$. Similar to derive Eq. 8, based on Eq. 12, we can express *Interval Score Distillation* (ISD) [17] as follows:

$$\mathcal{L}_{\text{ISD}}(\xi) = \mathbb{E}_{t,s,\pi} \left[c(t) \left\| F_{\theta}(\hat{\mathbf{x}}_{\{\pi,s\}}; s, \emptyset) - F_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, y) \right\|_{2}^{2} \right],$$
(13)

 $\overline{7}$

8 Z. Li et al.

which can be regarded as a classifier loss. With CFG, an additional generator loss is combined with ISD,

$$\mathcal{L}_{\text{ISD}}^{\text{CFG}}(\xi) = \mathbb{E}_{t,s,\pi}[c(t) \| [F_{\theta}(\hat{\mathbf{x}}_{\{\pi,s\}}; s, \emptyset) - F_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, \emptyset)] + w[F_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, \emptyset) - F_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, y)], \|_{2}^{2}],$$
(14)

we find the generator loss in Eq. 14 is highly correlated to a *Consistency Distillation Sampling* (CDS) studied in [44]:

$$\mathcal{L}_{\text{CDS}}(\xi) = \mathbb{E}_{t,s,\pi} \left[c(t) \| F_{\theta}(\mathbf{x}_{\{\pi,t\}}; t, y) - F_{\theta}(\bar{\mathbf{x}}_{\{\pi,t\to s;y\}}; s, y) \|_{2}^{2} \right], \quad (15)$$

where, $\bar{\mathbf{x}}_{\{\pi,t\to s;y\}} = F_{\theta}(\mathbf{x}_{\{\pi,t\}};t,s,y) = \Phi(\mathbf{x}_{\{\pi,t\}};t,s,y)$. Particularly, CDS is a special case as it mainly guides the generated views to match a particular origin of PF-ODE trajectory from the prior distribution through self-consistency. In such case, an upper bound for the distillation error can be deduced [44]:

$$\|\mathbf{x}_{\pi} - \mathbf{x}_0\|_2 = \mathcal{O}\left((\Delta t)^p\right) T.$$
(16)

where $\mathbf{x}_0 \sim p_{\text{data}}(\mathbf{x})$ is an real image, $\Delta t = \max\{|\delta_k|\}, k \in [0, ..., s]$ and $t = s + \delta_s$. With CFG, CDS can further facilitate a cross-consistency of text and null condition trajectories. However, it performs poorly with low CFG weight [44], which needs further improvement. Inspired by the similarity of CM and the optimization objectives of the score distillation method, we are motivated to reformulate the optimization objective in SDS by extending the theories in CM to improve text-to-3D generation quality. We sum up our solution in three aspects: 1) we reduce the error bound suggested Eq. 16 to improve the distillation quality; 2) we provide more reliable guidance during distillation accounting for CFG effects; 3) we implement constraints on pixel domain to enhance the fidelity.

4.2 Guided Consistency Sampling

In this section, we introduce the Guided Consistency Sampling (GCS), which includes three parts of objectives: a Compact Consistency (CC) loss, a Conditional Guidance (CG) score, and a Constraint on Pixel domain (CP):

$$\mathcal{L}_{\rm GCS}(\xi) = \mathcal{L}_{\rm CC}(\xi) + \mathcal{L}_{\rm CG}(\xi) + \mathcal{L}_{\rm CP}(\xi).$$
(17)

Compact Consistency (CC) Loss aims to improve further the self-consistency of PF-ODE denoising trajectory, which eventually reduces the distillation error bound in Eq. 16 for a more aligned distribution of rendered views. Inspired by [14, 51], we define a solution function:

$$G_{\theta}(\mathbf{x}_t; t, s, y) := \mathbf{x}_t + \int_t^s \frac{\mathbf{x}_u - \mathbb{E}[\mathbf{x}|\mathbf{x}_u]}{u} \mathrm{d}u.$$
(18)

which $G_{\theta}(\mathbf{x}_t; t, s, y)$ solves the PF-ODE from initial time t to a final time s according to exponential integrator [20, 21, 50]. Owing to the fact that G is

intractable as it can only be obtained by calculating s partial derivative at time t, we follow the scheme in [51] and adhere to the first-order definition⁶ of DPM-Solver [20], re-parameterised G as:

$$G_{\theta}(\mathbf{x}_t; t, s, y) = \frac{\sigma_s}{\sigma_t} \mathbf{x}_t - \alpha_s (e^{-h} - 1) \epsilon_{\theta}(\mathbf{x}_t, t, y),$$
(19)

where $h = \lambda_s - \lambda_t$ with the log-SNR λ defined as $\lambda = \log(\alpha/\sigma)$ and ϵ_{θ} is the prediction from the network. According to the previous definition, we mathematically define Compact Consistency Loss, a critical part of GCS, to improve the details of the 3D asset:

$$\mathcal{L}_{\rm CC}(\xi) = \mathbb{E}_{t,s,e,\pi} \left[\| G_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, e, \emptyset) - G_{\theta}(\bar{\mathbf{x}}_{\{\pi,t \to s;y\}}; s, e, \emptyset) \|_2^2 \right],$$
(20)

where t > s > e, and $\hat{\mathbf{x}}_{\{\pi,t\}}$ is obtained by DDIM inversion from $\mathbf{x}_{\{\pi,e\}}$, calculated as $\mathbf{x}_{\{\pi,e\}} = \alpha_e \mathbf{x}_{\pi} + \sigma_e \epsilon^*$ to $\bar{\mathbf{x}}_{\{\pi,e\to s\}}$ and eventually to $\bar{\mathbf{x}}_{\{\pi,s\to t\}}$. A null condition is applied to all DDIM inversion steps. Notably, ϵ^* is the random noise that will only be sampled once and kept fixed in practice⁷. During the training process, the parameters in the pre-trained model are frozen. In this condition, we substantiated Lemma 1, which extends the premises established in [37,44].

Lemma 1 ([14, 37, 44, 51]). Let $\Delta t = \max\{|\delta_k|\}, k \in [0, ..., n_s]$, where n_s is the index of δ at time step s, and $F_{\theta}(\cdot, \cdot)$ is the origin prediction function grounded on the empirical PF-ODE. Assume F_{θ} satisfies the Lipschitz condition, if there is a \mathbf{x}_{π} satisfying $\mathcal{L}_{CC}(\xi) = 0$, given an image $\mathbf{x}_0 \sim p_{data}(\mathbf{x})$, for any $t, s, e \in [0, ..., T]$ with t > s > e, we have:

$$\sup_{t,e,\mathbf{x}_{\pi}} \| \hat{\mathbf{x}}_{\{\pi,e\}}, \hat{\mathbf{x}}_{\{0,e\}} \|_{2} = \mathcal{O}\left((\Delta t)^{p} \right) (T-e),$$
(21)

 $\hat{\mathbf{x}}_{\{0,e\}}$ is the distribution of \mathbf{x}_0 diffused to time e, p is the order of ODE solver.

Proof. The proof is provided in the Appendix for completeness.

Lemma. 1 validates that CC achieves a lower upper limit on error margins than CDS (Eq. 16) for a better distillation from the pre-trained model. Considering the importance of CFG in generating high-quality contents, another issue worth noting is whether it is practical to integrate CFG (or Perp-Neg [1]) into a generator loss $\mathcal{L}_{\rm CC}$. In CDS, it applies CFG in every ODE denoising step, which leads to a large accumulated error magnified by CFG [22, 26]. Instead, we only implement CFG in one step [22] (see Fig. 2), leading to lower accumulated errors. This strategy significantly improves the generation quality with low CFG weight (see Fig. 3). However, we still observe artifacts in the generated 3D model, likely due to a lack of effective classifier loss that provides text-conditional guidance.

⁶ For the effect of higher orders, please refer to the Appendix for more details.

⁷ We use ϵ^* to ensure a valid diffusion process at low time step and reduce the computational cost, compared to [17].



Fig. 3: Generated views by using \mathcal{L}_{CC} with different CFG strategies at a low CFG weight (w = 7.5). While \mathcal{L}_{CC} (left) implements CFG in only one denoising step, \mathcal{L}_{CC}^* (right) applies CFG in every ODE denoising step.

Conditional Guidance (CG) Loss is proposed to accommodate guidance while maintaining a lower accumulated error, which we interpret as conditional guidance. Inspired by [10,17,26], we facilitate an alignment between the unconditional trajectory ($\mathcal{T}_{\{e\to 0, \emptyset\}}$ from a noisy sample at e to 0) and the text-conditional trajectory ($\mathcal{T}_{\{t\to e\to 0, y\}}$ derived from a noisy sample $\hat{\mathbf{x}}_{\{\pi, t\}}$ then sampling down to e and 0), separately. Mathematically, the CG score can be expressed as:

$$\mathcal{L}_{\mathrm{CG}}(\xi) = \mathbb{E}_{t,e,\pi} \left[\|F_{\theta}(\hat{\mathbf{x}}_{\{\pi,e\}}; e, \emptyset) - F_{\theta}(G_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, e, y); e, y)\|_{2}^{2} \right], \qquad (22)$$

where $F_{\theta}(\mathbf{x}; t, y)$ predicts the \mathbf{x}_0 from time step t given the conditional information y as shown in Eq. 8, and $G_{\theta}(\mathbf{x}; t, e, y)$ access the midpoint e of the trajectory from initial timestep t as explained in Eq. 19. In practice, we can integrate CFG into \mathcal{L}_{CG} for better performance. We believe the proposed \mathcal{L}_{CG} provides more reliable guidance with less accumulated error, as it avoids long navigation along the trajectory. Specifically, the gap between $\mathcal{T}_{\{e\to 0,\emptyset\}}$ and $\mathcal{T}_{\{t\to e\to 0,y\}}$ is a more precise guidance since it directly affects $\hat{\mathbf{x}}_{\{\pi,e\}}$, which is a more similar sample to \mathbf{x}_{π} [26] due to the DDIM inversion.

It is also interesting to see that \mathcal{L}_{ISD} is equivalent to \mathcal{L}_{CG} when set e = 0 in computing $F_{\theta}(G_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}};t,e,y);e,y)$. We intuitively explain the effect of a midpoint e to reduce the error of one-step denoising from a noisy sample conditioned at large t. Eventually, we find \mathcal{L}_{CC} and \mathcal{L}_{CG} work together improving the quality of generated 3D asset.

Constraint on Pixel (CP) Domain aims to achieve a closer resemblance in the pixel domain for \mathcal{L}_{CG} . While a satisfactory \mathbf{x}_{π} is obtained in the latent domain, whether such an \mathbf{x}_{π} is equally satisfactory in the pixel domain appears to have been overlooked. In fact, the optimized \mathbf{x}_{π} may not adhere to the prior distribution stipulations of VAE, leading to out-of-distribution artifacts. Consequently, this leads to unrealistic color in the decoded images (*e.g.*, a paint-like color in the generated results of LucidDreamer [17]). Inspired by similar methods proposed in [53], we further calculate \mathcal{L}_{CG} in the pixel domain for enhanced supervision. In this context, given an image decoder \mathcal{D} along with $\hat{\mathbf{x}}_{\{\pi,e\}}$, the definition of CP can be articulated as follows:

$$\mathcal{L}_{\rm CP}(\xi) = \mathbb{E}_{t,e,\pi} \left[\| \mathcal{D}(\dot{\mathbf{x}}_{\pi}) - \mathcal{D}(\tilde{\mathbf{x}}_{\pi}) \|_2^2 \right], \tag{23}$$

$$\dot{\mathbf{x}}_{\pi} = F_{\theta}(G_{\theta}(\hat{\mathbf{x}}_{\{\pi,t\}}; t, e, y); e, y), \tag{24}$$

$$\tilde{\mathbf{x}}_{\pi} = F_{\theta}(\hat{\mathbf{x}}_{\{\pi,e\}}; e, \emptyset). \tag{25}$$

Our empirical observations indicate that optimizing CP yields notable enhancements in color fidelity, albeit with increased computational cost and over-saturation. Consequently, we suggest an additional strategy outlined in Sec. 4.3, aimed at alleviating over-saturation. Extensive experiments lead us to conclude that the CP term is constructive in bolstering the overall quality of 3D assets.

4.3 Brightness-equalized Generation

A common problem in previous works is that the generated 3D assets often suffer from over-saturation, especially in those that use Gaussian Splatting as the 3D representation to be trained. We experimentally find that the high brightness region (highlight) generated in the current training epoch is carried over to the next epoch⁸. That is, the brightness of the highlight points accumulates during training, eventually leading to over-saturation. To alleviate this issue, we propose resetting the brightness of the Gaussian⁹ adaptively according to the exposure status of the generated view. Specifically, in each training epoch, we calculate the $m^{\text{th}}(85)$ percentile of each image \mathbf{x}_{π}^{i} in an image batch $\mathbf{X}_{\pi} =$ $\{\mathbf{x}_{\pi}^{i}, i \in [1, ..., B]\}$, noted as $P_m = \{p_m^{i}, i \in [1, ..., B]\}$, where B is the batch size. We then find the maximum value in P_m and reset the brightness of Gaussian to $T_{\rm B} = 0.8$ of the current brightness if $\max(P_m) > T_{\rm GS}$, $T_{\rm GS} = 0.9$. Such a simple method will significantly alleviate the over-saturation issues and facilitate brightness equalization.

5 Experiments

5.1 Implementation Details

We implement the GCS with PyTorch and train it on an A800 GPU with the Adam optimizer. Some hyperparameters (the learning rates, camera positions, rendering resolution, etc.) are similar to that in [17]. We use 5000 epochs as [43] for total training iterations, which takes about one hour per scene.

Time schedule. In the experiment, t is sampled from a uniform distribution $\mathcal{U}(20, 500 + \delta_{\text{warm}})$, where δ_{warm} is linearly decreasing from 480 to 0 in the first 1500 epochs. $s = t - \delta$, where $\delta = 100$ for most cases. e is sampled from a uniform distribution as $\mathcal{U}(s - \delta, s - \frac{\delta}{10})$.

Initialization. We apply Shap-E [11] and Point-E [27] to initialize the 3D Gaussians, which are widely used in 3DGS-based text-to-3D generation [17,44].

5.2 Text-to-3D Generation

Qualitative Comparison. We have showcased diverse views of the 3D assets generated by our versatile approach, indicating that we can not only generate highly detailed figures but also create reasonable outputs to align with the complex prompt (results shown in the first row). To comprehensively assess the efficacy of our proposed Guided Conditional Sampling (GCS), we conducted a

⁸ A visualization of the brightness accumulation can be found in the Appendix.

⁹ We use the average of RGB channels as the brightness.

12 Z. Li et al.



Fig. 4: Qualitative comparison among the proposed and other methods on text-to-3D generation results. From left to right, results generated by DreamFusion [29], GaussianDreamer [46], ProlificDreamer [43], LucidDreamer [17], and the proposed method, with a CFG weight 100, 100, 7, 7, 7, respectively. For each sub-figure, left: main view, right-top: back view, right-bottom: normal/depth map of the 3D asset.

comparative analysis against the current state-of-the-art in GS-based [17, 46] and baseline NeRF-based [29, 43] methods in text-to-3D generation, as shown in Fig. 4. To ensure a fair comparison, all listed methods implement Stable Diffusion 2.1 as the base model. Other hyperparameters were configured following the default settings of the respective methods. The results indicate that previous methods often struggled with insufficient details, prominent artifacts (*e.g.*, Janus problem [1]), or over-saturation. In contrast, our approach exhibits notable improvements in detail (*e.g.*, the armor of the cat), fidelity, and exposure. Specifically, we can generate highly detailed textures, achieve photorealistic effects (such as lights on the cat's armor), and maintain more consistent colors (such as the cat's fur color) even with a low CFG weight, illustrating the effectiveness of our proposed GCS and the brightness-equalization techniques employed during training. Additional results can be found in the Appendix.

Quantitative Comparison. We present the average CLIP Score, FID Score, and user preference of five comparative methods [17, 29, 39, 43, 46] and the proposed GCS on 120 generated views across 30 prompts selected from the Dream-Fusion gallery¹⁰ in Table 2. For CLIP Score computation, we follow the steps presented in [29]. Results show that our performance is comparable to that of VSD [43] in the evaluation of prompt alignment, with much fewer instances of Janus problems [1], as evidenced in the qualitative comparison shown in Fig. 4. The proposed approach exhibits improved 3D consistency and fidelity for the 3D asset, resulting in a higher user preference in prompt alignment.

¹⁰ https://dreamfusion3d.github.io/gallery.html

Table 2: Quantitative comparison regarding CLIP Score, FID Score, and user preference on three given questions (Q1, Q2, and Q3). The bold (underline) number indicates the best (second-best) results.

	GaussianDreamer	DreamGaussian	DreamFusion	LucidDreamer	ProlificDreamer	Ours
CLIP Score↑	30.29	28.65	30.68	31.30	33.17	<u>32.37</u>
FID Score↓	133.48	215.89	137.22	109.57	119.53	103.40
User Preference: $Q1\uparrow^1$	9.47	1.05	6.32	24.21	12.63	46.32
User Preference: $Q2\uparrow^2$	10.53	1.05	6.32	30.53	7.37	44.21
User Preference: Q3 ^{↑3}	9.47	2.11	6.32	26.32	11.58	44.21

¹Q1: From the perspective of color brightness, which of the following methods produces the most balanced brightness (e.g., no over-saturation, moderate saturation)? ² Q2: From the perspective of **prompt alignment**, which of the following methods align most with the prompt?

³ Q3: From the perspective of fidelity, which of the following methods produces the most realistic objects?



"a DSLR photo of a cat wearing armor."

Fig. 5: Ablation study of proposed components.LucidDreamer [17] serves as the baseline, we demonstrate the results under settings: (a) LucidDreamer [17] + BEG, (b) $\mathcal{L}_{CC}(\xi) + \mathcal{L}_{CG}(\xi)$, (c) $\mathcal{L}_{GCS}(\xi)$, and (d) full mode ($\mathcal{L}_{GCS}(\xi) + BEG$) from left to right.

For FID, we follow similar steps in [43] to calculate FID between 120 views of 3D object¹¹ and 50k Stable Diffusion 2.1 generated images. We implement Perp-Neg when necessary for different objects. The results show that the generated views of the proposed method are more similar to the prior distribution compared to other methods, indicating that the proposed GCS reduces distillation error. User Study. We further conduct a user study among 30 volunteers for a more comprehensive evaluation. Specifically, we asked the volunteers to evaluate the generation quality regarding brightness (Q1), prompt alignment (Q2), and fi*delity* (Q3). Those volunteers are also required to indicate their preference based on rendered 360° videos of 5 objects randomly sampled from 30 examples within six methods. We shuffled the presented order of the generated results to avoid any form of leakage. The results of user preference in percentage form shown in Table 2 indicate that our method performs best in all aspects. An example question can be found in the Appendix.

5.3Ablation Study

We conducted an ablation study to assess the effects of parameter setup (including δ , $T_{\rm GS}$, and $T_{\rm B}$), the proposed Guided Consistency Sampling (GCS) loss, and the brightness-equalized generation method (BEG).

 $^{^{11}}$ Views are generated with fixed elevation (0°, 3.6k views generated in total) and varying azimuth (uniformly covering 360°)



Fig. 6: Ablation study of hyperparameters' effects (left: δ , middle: T_{GS} , right: T_{B})

Effect of time interval δ is evaluated by an ablation study on δ shown in Fig. 6. We find that large δ causes over-smoothness (*e.g.*, low level of detail on the strawberry's surface), while small δ causes color distortion (*e.g.*, unrealistic color in strawberry). 100 is set to generate visual-pleasant results.

Effect of Different Loss and BEG¹². As illustrated in Fig. 5, we compared the full model's ($\mathcal{L}_{GCS}(\xi)$ + BEG) generated results with those from Lucid-Dreamer [17], and two alternative models: $\mathcal{L}_{CC}(\xi) + \mathcal{L}_{CG}(\xi)$ and $\mathcal{L}_{GCS}(\xi)$ without BEG method, with a CFG weight set to 7.5. Compared to LucidDreamer [17], the full model successfully generated more intricate details (*e.g.*, armor's patterns and higher armor coverage). While there was a slight alleviation in color distortion (inconsistency of fur color between the back and front of the cat), the issue may persist. Consequently, we introduced \mathcal{L}_{CP} to mitigate color distortion further (comparing $\mathcal{L}_{CC}(\xi) + \mathcal{L}_{CG}(\xi)$ and $\mathcal{L}_{GCS}(\xi)$).

However, we observed a significant increase in overall brightness during optimization. The implementation of the BEG method notably alleviates this issue, as evidenced in the qualitative comparison between $\mathcal{L}_{GCS}(\xi)$ and the full model in Fig. 5. We find that BEG will only affect the overall brightness, but not the fidelity, as shown in a qualitative comparison between LucidDreamer [17] and LucidDreamer [17] + BEG. The effect of BEG is controlled by $T_{\rm B}$ and $T_{\rm GS}$. Specifically, large (small) $T_{\rm B}$ causes over-saturation (color distortion) results (see the comparison in Fig. 6). $T_{\rm B} = 0.8$ works for general cases. Small $T_{\rm GS}$ will reset the brightness too frequently to cause artifacts, $T_{\rm GS} = 0.9$ is a suitable setup.

6 Conclusion

In this work, we connect consistency distillation to score distillation. From this connection, we propose Guided Consistency Sampling (GCS), an optimization framework that includes three parts: Compact Consistency (CC) loss for improved generator loss, Conditional Guidance (CG) score to enhance conditional guidance, and Constraints in the Pixel domain (CP) for improved fidelity. In addition, we innovate the Brightness-equalized Generation (BEG) to tackle the over-saturation issue in 3DGS training. The proposed approach improves the details, quality, and lighting effects of the generated 3D assets.

¹² We conduct additional ablation studies regarding CC and CG to analyze the individual effect on generated 3D assets in the Appendix.

15

Acknowledgement

This work is supported by Rapid-Rich Object Search (ROSE) Lab, Nanyang Technological University, Singapore, and the State Key Lab of Brain-Machine Intelligence, Zhejiang University, Hangzhou, China.

References

- Armandpour, M., Zheng, H., Sadeghian, A., Sadeghian, A., Zhou, M.: Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3D, alleviate janus problem and beyond. arXiv preprint arXiv:2304.04968 (2023)
- Bensadoun, R., Kleiman, Y., Azuri, I., Harosh, O., Vedaldi, A., Neverova, N., Gafni, O.: Meta 3d texturegen: Fast and consistent texture generation for 3d objects. arXiv preprint arXiv:2407.02430 (2024)
- Burley, B., Studios, W.D.A.: Physically-based shading at disney. In: ACM SIG-GRAPH (2012)
- 4. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3D content creation. Int. Conf. Comput. Vis. (2023)
- Chen, Y., Zhang, C., Yang, X., Cai, Z., Yu, G., Yang, L., Lin, G.: It3D: Improved text-to-3D generation with explicit view synthesis. In: AAAI (2024)
- Chen, Y., Georgiou, T.T., Pavon, M.: On the relation between optimal transport and schrödinger bridges: A stochastic control viewpoint. Journal of Optimization Theory and Applications (2016)
- Chen, Z., Wang, F., Wang, Y., Liu, H.: Text-to-3D using gaussian splatting. In: IEEE Conf. Comput. Vis. Pattern Recog. (2024)
- Deng, K., Omernick, T., Weiss, A., Ramanan, D., Zhu, J.Y., Zhou, T., Agrawala, M.: Flashtex: Fast relightable mesh texturing with lightcontrolnet. arXiv preprint arXiv:2402.13251 (2024)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Adv. Neural Inform. Process. Syst. (2020)
- Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
- Jun, H., Nichol, A.: Shap-e: Generating conditional 3D implicit functions. arXiv preprint arXiv:2305.02463 (2023)
- Katzir, O., Patashnik, O., Cohen-Or, D., Lischinski, D.: Noise-free score distillation. arXiv preprint arXiv:2310.17590 (2023)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. ACM Trans. Graph. (2023)
- Kim, D., Lai, C.H., Liao, W.H., Murata, N., Takida, Y., Uesaka, T., He, Y., Mitsufuji, Y., Ermon, S.: Consistency trajectory models: Learning probability flow ode trajectory of diffusion. Int. Conf. Learn. Represent. (2024)
- 15. Léonard, C.: A survey of the schr\" odinger problem and some of its connections with optimal transport. arXiv preprint arXiv:1308.0215 (2013)
- Li, Z., Zheng, Q., Shi, B., Pan, G., Jiang, X.: Dani-net: Uncalibrated photometric stereo by differentiable shadow handling, anisotropic reflectance modeling, and neural inverse rendering. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
- Liang, Y., Yang, X., Lin, J., Li, H., Xu, X., Chen, Y.: Luciddreamer: Towards high-fidelity text-to-3D generation via interval score matching. arXiv preprint arXiv:2311.11284 (2023)

- 16 Z. Li et al.
- Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3D: High-resolution text-to-3D content creation. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
- Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3D object. In: Int. Conf. Comput. Vis. (2023)
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. Adv. Neural Inform. Process. Syst. (2022)
- Lu, C., Zhou, Y., Bao, F., Chen, J., Li, C., Zhu, J.: Dpm-solver++: Fast solver for guided sampling of diffusion probabilistic models. arXiv preprint arXiv:2211.01095 (2022)
- Luo, S., Tan, Y., Huang, L., Li, J., Zhao, H.: Latent consistency models: Synthesizing high-resolution images with few-step inference. arXiv preprint arXiv:2310.04378 (2023)
- McAllister, D., Ge, S., Huang, J.B., Jacobs, D.W., Efros, A.A., Holynski, A., Kanazawa, A.: Rethinking score distillation as a bridge between image distributions. arXiv preprint arXiv:2406.09417 (2024)
- Metzer, G., Richardson, E., Patashnik, O., Giryes, R., Cohen-Or, D.: Latent-NeRF for shape-guided generation of 3D shapes and textures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12663– 12673 (2023)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. Eur. Conf. Comput. Vis. (2021)
- Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. 2023 ieee. In: CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6038–6047 (2022)
- Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3D point clouds from complex prompts. arXiv preprint arXiv:2212.08751 (2022)
- Pham, H.A.: Reduction of function evaluation in differential evolution using nearest neighbor comparison. Vietnam Journal of Computer Science (2015)
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3D using 2d diffusion. In: Int. Conf. Learn. Represent. (2022)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: IEEE Conf. Comput. Vis. Pattern Recog. (2022)
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic textto-image diffusion models with deep language understanding. Adv. Neural Inform. Process. Syst. (2022)
- Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial diffusion distillation. arXiv preprint arXiv:2311.17042 (2023)
- 33. Seo, J., Jang, W., Kwak, M.S., Kim, H., Ko, J., Kim, J., Kim, J.H., Lee, J., Kim, S.: Let 2d diffusion model know 3D-consistency for robust text-to-3D generation. arXiv preprint arXiv:2303.07937 (2023)
- 34. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3D shape synthesis. Adv. Neural Inform. Process. Syst. (2021)
- Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3D generation. arXiv preprint arXiv:2308.16512 (2023)

17

- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- Song, Y., Dhariwal, P., Chen, M., Sutskever, I.: Consistency models. arXiv preprint arXiv:2303.01469 (2023)
- Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Scorebased generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456 (2020)
- Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3D content creation. arXiv preprint arXiv:2309.16653 (2023)
- 40. Tang, J., Wang, T., Zhang, B., Zhang, T., Yi, R., Ma, L., Chen, D.: Make-it-3D: High-fidelity 3D creation from a single image with diffusion prior. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
- Tsalicoglou, C., Manhardt, F., Tonioni, A., Niemeyer, M., Tombari, F.: Textmesh: Generation of realistic 3D meshes from text prompts. In: Int. Conf. on 3D Vision (3DV) (2024)
- Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3D generation. In: IEEE Conf. Comput. Vis. Pattern Recog. (2023)
- Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: Highfidelity and diverse text-to-3D generation with variational score distillation. Adv. Neural Inform. Process. Syst. (2024)
- Wu, Z., Zhou, P., Yi, X., Yuan, X., Zhang, H.: Consistent3D: Towards consistent high-fidelity text-to-3D generation with deterministic sampling prior. arXiv preprint arXiv:2401.09050 (2024)
- 45. Yang, X., Chen, Y., Chen, C., Zhang, C., Xu, Y., Yang, X., Liu, F., Lin, G.: Learn to optimize denoising scores for 3D generation: A unified and improved diffusion prior on NeRF and 3D gaussian splatting. arXiv preprint arXiv:2312.04820 (2023)
- 46. Yi, T., Fang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaussiandreamer: Fast generation from text to 3D gaussian splatting with point cloud priors. IEEE Conf. Comput. Vis. Pattern Recog. (2024)
- Yu, X., Guo, Y.C., Li, Y., Liang, D., Zhang, S.H., Qi, X.: Text-to-3D with classifier score distillation. arXiv preprint arXiv:2310.19415 (2023)
- Zhang, H., Chen, B., Yang, H., Qu, L., Wang, X., Chen, L., Long, C., Zhu, F., Du, D., Zheng, M.: Avatarverse: High-quality & stable 3D avatar creation from text and pose. In: AAAI (2024)
- Zhang, J., Zhang, X., Zhang, H., Liew, J.H., Zhang, C., Yang, Y., Feng, J.: Avatarstudio: High-fidelity and animatable 3D avatar creation from text. arXiv preprint arXiv:2311.17917 (2023)
- 50. Zhang, Q., Chen, Y.: Fast sampling of diffusion models with exponential integrator. arXiv preprint arXiv:2204.13902 (2022)
- Zheng, J., Hu, M., Fan, Z., Wang, C., Ding, C., Tao, D., Cham, T.J.: Trajectory consistency distillation. arXiv preprint arXiv:2402.19159 (2024)
- Zhou, Z., Tulsiani, S.: Sparsefusion: Distilling view-conditioned diffusion for 3D reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12588–12597 (2023)
- Zhu, J., Zhuang, P.: Hifa: High-fidelity text-to-3D with advanced diffusion guidance. arXiv preprint arXiv:2305.18766 (2023)