

PlacidDreamer: Advancing Harmony in Text-to-3D Generation

Anonymous Authors

ABSTRACT

Recently, text-to-3D generation has attracted significant attention, resulting in notable performance enhancements. Previous methods utilize end-to-end 3D generation models for initializing 3D Gaussians, and multi-view diffusion models to enforce multi-view consistency. Moreover, they employ text-to-image diffusion models to refine details with score distillation algorithms. However, these methods exhibit two limitations. Firstly, they encounter conflicts in generation directions since different models aim to produce diverse 3D assets. Secondly, the issue of over-saturation in score distillation has not been thoroughly investigated and solved. To address these limitations, we propose PlacidDreamer, a text-to-3D framework that harmonizes initialization, multi-view generation, and text-conditioned generation with a single multi-view diffusion model, while simultaneously employing a novel score distillation algorithm to achieve balanced saturation. To unify the generation direction, we introduce the Latent-Plane module, a training-friendly plug-in extension that enables multi-view diffusion models to provide fast geometry reconstruction for initialization and enhanced multi-view images to personalize the text-to-image diffusion model. To address the over-saturation problem, we propose to view score distillation as a multi-objective optimization problem and introduce the Balanced Score Distillation algorithm, which offers a Pareto Optimal solution that achieves both rich details and balanced saturation. Extensive experiments validate the outstanding capabilities of our PlacidDreamer. The code will be available on GitHub.

CCS CONCEPTS

• Computing methodologies → Computer vision problems.

KEYWORDS

3D Generation, text-to-3D, score distillation

1 INTRODUCTION

The task of generating 3D assets from text, known as text-to-3D, has garnered significant attention for its potential to simplify 3D creation, a process once requiring specialized knowledge. Due to the relative scarcity and lower quality of 3D data compared with 2D data, one promising approach is to adapt pre-trained 2D models for 3D generation. An optimization-based approach leveraging score distillation algorithms, which distill generative capability from pre-trained 2D diffusion models to guide subsequent 3D generations, has emerged as a dominant approach in this field.

Unpublished working draft. Not for distribution.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted by ACM, provided that the copies are not made for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM MM, 2024, Melbourne, Australia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnn>

Since the introduction of the first score distillation algorithm, Score Distillation Sampling (SDS) [37], subsequent works have significantly advanced this optimization-based approach, enhancing both generation quality and speed. One significant factor influencing generation quality is the multi-face problem. To address this problem, Magic123 [38], DreamCraft3D [47], Consistent123 [27], and EfficientDreamer [61] incorporate multi-view diffusion models to enhance multi-view consistency. More recently, the introduction of 3D Gaussian Splatting [22] has further optimized the pipeline with convenient initialization, faster rendering, and training speed. LucidDreamer [25], GaussianDreamer [58], and GSGEN [6] propose to leverage end-to-end 3D generation models [20, 35] to provide a robust 3D Gaussian initialization, thereby enhancing overall quality.

Despite significant progress in the aforementioned methods, there are still two main limitations that require attention:

- **Conflicting Optimization Directions.** The integration of multiple generative models within a single pipeline can lead to contradictory optimization directions. For instance, the score distillation guidance derived from multi-view diffusion models may be at odds with that from text-to-image diffusion models [38], necessitating the development of specialized balancing strategies [27] or additional refinement [47] for optimal performance. Moreover, employing distinct generative models for supervision across different stages [6, 58] can cause the generative model in later stages to ignore the outputs of its predecessors and independently generate new results based on its intrinsic data interpretations.
- **Over-Saturation in Score Distillation.** The problem of over-saturation within score distillation algorithms remains insufficiently explored and unresolved. This issue manifests as a discrepancy in color distribution between the 3D content created through score distillation techniques and the 2D images generated by diffusion inference processes. Although certain methodologies yield results with more appropriate levels of saturation, they impose a significant computational load. Examples include LoRA finetuning [53], adversarial training [5], or supervision in multiple spaces with iterative sampling [62]. Thus, a comprehensive understanding of the over-saturation problem and a fast yet effective score distillation algorithm is needed.

To overcome these challenges, we propose PlacidDreamer, a framework that harmonizes initialization, multi-view generation, and text-conditioned generation with a single multi-view diffusion model, while simultaneously employing a novel score distillation algorithm to achieve balanced saturation. More specifically, to address the first limitation, we newly devised a Latent-Plane module. This module enhances the multi-view diffusion model by enabling fast geometry reconstruction and improving its capabilities in generating multi-view images. The reconstructed geometry is then utilized to initialize 3D Gaussian points, and the improved multi-view images are used to personalize the text-to-image diffusion model with directional prompts. This coordinated approach aligns the generation directions within the pipeline with the outcomes from the multi-view diffusion model, promoting convergence and



A golden retriever with a blue bowtie.



A bright red fire hydrant.



A leather-bound book with gold details.



A pirate flag with skull and crossbones.



A shiny emerald green beetle.



An ivory candlestick holder.



A worn-out leather briefcase.



An intricate ceramic vase with peonies painted on it.



A wet, vibrant beach ball.



A blooming potted orchid with purple flowers.

Figure 1: 3D generations of PlacidDreamer. More results are provided in the Supplementary Material.

significantly enhancing the quality of the generated content. Additionally, the Latent-Plane module is designed to be training-efficient and is adaptable to a variety of multi-view diffusion models, accommodating different viewpoint configurations. To address the second limitation, we delve into the causes of over-saturation. We decompose the score distillation equation into two primary components: classifier guidance and smoothing guidance. Prior algorithms have suffered from an imbalance where classifier guidance overwhelmingly dominates smoothing guidance, leading to prevalent over-saturation. Our analysis reveals that in over 30% of instances, the optimization directions of these two guidances form obtuse angles, resulting in negative optimization along certain directions with a fixed Classifier-Free Guidance (CFG) parameter, making it challenging to maintain control over balance after multiple optimization steps. To rectify this imbalance, we propose to treat score distillation as a multi-objective optimization problem and introduce

a Balanced Score Distillation (BSD) algorithm. This algorithm incorporates a multi-objective optimization solution, Multiple-Gradient Descent Algorithm (MGDA), which dynamically adjusts the optimization directions to converge at Pareto Optimal points, where the generated results exhibit both rich details and balanced saturation.

We evaluate PlacidDreamer both qualitatively and quantitatively to demonstrate its effectiveness. Extensive experiments indicate the superiority of our method over previous methods. Quantitative evaluations conducted on the T3Bench [10] benchmark reveal that PlacidDreamer consistently outperforms baseline methods by a margin of at least ten points in both generation quality and alignment metrics. Furthermore, we conduct ablation studies to respectively evaluate the contribution of each proposed module to the overall quality. To further highlight the BSD algorithm's effectiveness, we replaced the score distillation algorithms in various open-source text-to-3D frameworks with BSD, while keeping other

components constant. Results show that BSD consistently enhances these frameworks' performance.

Our contributions can be summarized as follows:

- We introduce PlacidDreamer, a novel framework designed for high-fidelity text-to-3D generation. PlacidDreamer advances a more harmonious generation process through two novel approaches: the Latent-Plane module, which enhances multi-view diffusion models, and the Balanced Score Distillation algorithm, which optimizes detail richness and saturation control.
- We identify the causes of the over-saturation problem in score distillation algorithms and propose to view score distillation as a multi-objective optimization problem, with the goal of optimizing towards Pareto Optimal points to stabilize the outcomes of generation.
- Extensive experiments, including both quantitative metrics and qualitative assessments, demonstrate that PlacidDreamer significantly outperforms existing state-of-the-art methods.

2 RELATED WORKS

Text-to-3D Generation. DreamFields [18] initially employs a pre-trained model, CLIP [40], to guide the optimization of NeRF [34]. Trying to leverage the generative nature of the diffusion model, Dreamfusion [37] introduces SDS, a loss associated with a score function derived from distilling 2D diffusion models. Subsequent works significantly improve SDS-based methods, including those using multi-view consistent models [24, 30, 44, 61], enhancing pipeline structures [4, 26, 32, 33, 43, 54, 58], introducing extra generation priors [2, 15, 17, 19, 48, 56, 60], or exploring timestep scheduling [16]. Recently, some fast-forward reconstruction models [13, 28, 30, 31, 50] have emerged for faster 3D generation.

Score Distillation. Given the heavy reliance on score distillation methods in these studies, addressing associated issues with SDS is crucial. While many works [11, 21, 48, 51, 52, 59] focus on addressing over-smoothing, fewer tackle over-saturation problems. ProlificDreamer [53] introduces VSD, leveraging LoRA finetuning for 3D distribution modeling, which can also alleviate over-saturation. IT3D [5] handles over-saturation by training a discriminator distinguishing 3D assets from text-to-2D images. Additionally, HiFA [62] presents an iterative score distillation process for a more accurate sampling direction. It's worth noting that these methods introduce a significant computational burden, slowing down overall speed.

3 PRELIMINARIES

3.1 Diffusion Models

For discrete-time diffusion models [9, 12, 36, 45], given a data distribution $q(\mathbf{x}) = q_0(\mathbf{x}_0)$, we construct a forward process with a series of distributions $q_t(\mathbf{x}_t) = \mathcal{N}(\alpha_t \mathbf{x}_0, (1 - \alpha_t^2)\mathbf{I})$ with decreasing $\{\alpha_t | t \in [0, T], t \in \mathbb{Z}\}$, where $\alpha_0 = 1$ and $\alpha_T \approx 0$. In the generation process, we start from $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and generate a sample of the previous timestep iteratively with a denoising network $\epsilon_\phi(\mathbf{x}_t, t)$ trained by minimizing the prediction of added noise, which is given by

$$\mathbb{E}_{\mathbf{x}_0 \sim q_0(\mathbf{x}_0), t, \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} w(t) \|\epsilon_\phi(\alpha_t \mathbf{x}_0 + \sigma_t \epsilon, t) - \epsilon\|_2^2, \quad (1)$$

where $w(t)$ is to balance losses between different timesteps, t is uniformly selected from 0 to T , and $\sigma_t = \sqrt{1 - \alpha_t^2}$.

3.2 Score Distillation Sampling (SDS)

Score distillation is an optimization-based generation method that distills knowledge from pre-trained 2D diffusion models to guide other generations. The first score distillation method, SDS [37] is denoted as

$$\nabla_\theta \mathcal{L}_{\text{SDS}}(\phi, \mathbf{x} = g(\theta)) \triangleq \mathbb{E}_{t, \epsilon} [\omega(t) (\hat{\epsilon}_\phi(\mathbf{x}_t; y, t) - \epsilon) \frac{\partial \mathbf{x}}{\partial \theta}], \quad (2)$$

where $\mathbf{x}_t = \alpha_t \mathbf{x} + \sigma_t \epsilon$ and other symbols are defined the same as in Equation (1). Intuitively, this loss perturbs \mathbf{x} with a random amount of noise corresponding to the timestep t , and estimates an update direction that follows the score function of the diffusion model to move to a higher density region.

3.3 Multiple-Gradient Descent Algorithm

MGDA is a gradient-based algorithm used to solve multi-objective optimization problems. Multi-objective optimization refers to the task of finding a Pareto Optimal solution under multiple optimization criteria, where optimizing one objective does not deteriorate the solution of another objective during the optimization process. In the context of solving multi-objective optimization tasks, MGDA determines a descent direction that is common to all criteria. For instance, when applied to a binary-objective optimization problem optimizing $\mathcal{L}^1(x)$ and $\mathcal{L}^2(x)$, MGDA will identify the direction orthogonal to $\nabla_x \mathcal{L}^1(x) - \nabla_x \mathcal{L}^2(x)$. The algorithm has been demonstrated to converge to a Pareto Optimal point, where there is no available optimization that can improve one objective without worsening another. Further elaboration on MGDA can be found in the work of Désidéri [8].

4 METHODS

4.1 Pipeline

Given a text prompt, we initially generate a reference image using Stable Diffusion [41] or MVDream [44]. Unlike previous methods that transform text-to-3D generation into single-view reconstruction, we do not aim to precisely reconstruct the reference object. After background removal, the reference image is fed into the multi-view diffusion model to generate multi-view images. To optimize the initialization of 3D Gaussian points [22] and ensure its compatibility with the multi-view images, we introduce Latent-Plane. This module, as detailed in Section 4.2, seamlessly integrates into the latent layers of any multi-view diffusion model, enabling fast reconstruction of a volume density field and enhancement of the generated images within 40 seconds. Moreover, it is training-friendly (only requiring $4 \times A6000$ GPUs for 18 hours), as it directly leverages 3D spatial knowledge from pre-trained multi-view diffusion model. Subsequently, the volume density field is used to initialize the 3D Gaussian points on the object surface, while the generated multi-view images are used to fine-tune Stable Diffusion with LoRA [14] technique. During this fine-tuning process, directional prompts such as "left view" are incorporated into the original text, enhancing the diffusion model's awareness of 3D space. Finally, we supervise the splatted images using the BSD guidance from the fine-tuned

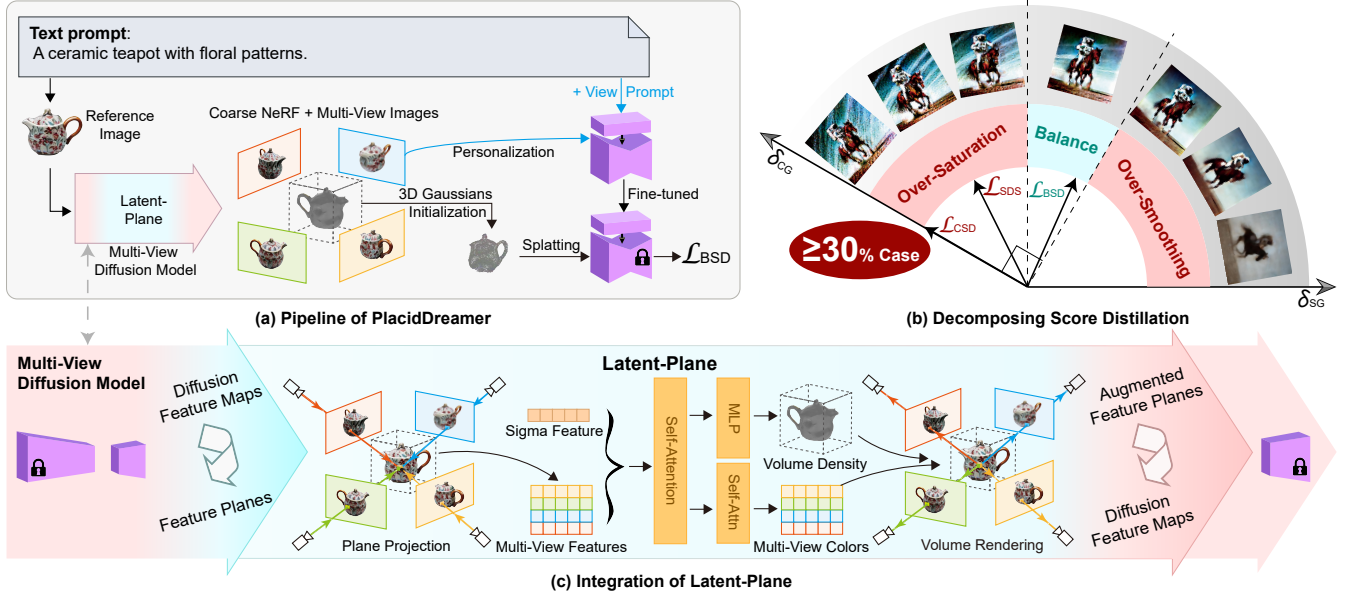


Figure 2: (a) The pipeline of PlacidDreamer. (b) Score distillation can be decomposed into two directions: classifier guidance δ_{CG} and smoothing guidance δ_{SG} . CSD [59] only utilizes classifier guidance. In more than 30% of cases, the angle between these two guidance vectors is obtuse. In such scenarios, using a fixed CFG parameter in SDS may result in negative optimization in the δ_{SG} direction, leading to over-saturation. However, BSD algorithm ensures that each optimization step is non-negative in both directions. (c) The integration of the Latent-Plane module with multi-view diffusion models.

diffusion model. As elaborated in Section 4.3, we decompose score distillation and reveal a conflict between optimization directs, leading to color discrepancies. We suggest treating score distillation as a multi-objective optimization problem and we introduce the BSD algorithm, achieving rich details and reasonable colors.

4.2 Latent Plane

The proposed method, Latent-Plane, serves as a plug-in module compatible with various multi-view diffusion models. These diffusion models typically take an input image I_{π_0} of an object from a specific viewpoint π_0 and generate corresponding images I_{π_i} from other viewpoints $\pi_i, i = 1, 2, \dots, N$. Inspired by ConsisNet [57], which reinforces feature patches by utilizing coordinate relationships derived from back-projection, we hypothesize that pre-trained features within multi-view diffusion models possess sufficient knowledge of the 3D space to conduct direct reconstruction. Leveraging this idea, we treat latent feature maps akin to the feature plane in the Tri-Plane [3] model, allowing us to reconstruct a volume density field. Furthermore, the reconstruction results can be utilized to enhance the latent feature map through volume rendering, thereby improving the accuracy of multi-view predictions.

4.2.1 Multi-View Feature Gathering. Most multi-view diffusion models currently utilize the Unet [42] architecture for their denoisers. To illustrate our strategy of selecting the appropriate latent layer for inserting the Latent-Plane module, we consider the Unet architecture as an exemplar. Within the Unet architecture, assuming there are L decoder blocks generating L feature maps denoted as $F_j \in \mathbb{R}^{H \times W \times D}$ with j indexing these maps and D representing the feature dimension. We choose the feature map with the highest

resolution and the smallest index, denoted as F_k , as it encompasses the deepest and most intricate features extracted by the model, containing comprehensive geometric information.

At each timestep t , during the reverse diffusion process, we obtain N feature maps $F_k^{(i)}, i = 1, 2, \dots, N$ from N viewpoints. For every point \mathbf{x} in the 3D space, we project it onto each feature map $F_k^{(i)}$ and derive its corresponding feature $f_k^{(i)}(\mathbf{x})$ through tri-linear interpolation, denote as,

$$f^{(i)}(\mathbf{x}) = \text{Interp_2D}(\text{Proj}(\mathbf{x}, F_k^{(i)}), F_k^{(i)}), \quad (3)$$

where Interp_2D represents the 2D interpolation function, and Proj denotes the function for projecting 3D spatial points onto 2D planes. We then apply an additional linear layer to extract low-dimensional features from the extracted features. To better represent the features of spatial point \mathbf{x} , we concatenate each feature with its respective camera embeddings $e_{\text{cam}}^{(i)}$ and coordinate embeddings $e_{\text{pos}}(\mathbf{x})$.

$$e^{(i)}(\mathbf{x}) = \text{Concat}[\text{Linear}(f^{(i)}(\mathbf{x})), e_{\text{cam}}^{(i)}, e_{\text{pos}}(\mathbf{x})]. \quad (4)$$

4.2.2 Plane-Based NeRF. So far, \mathbf{x} has N features gathered from N viewpoints, serving as N tokens, which will be mutually enhanced through the Attention layers. To obtain the volume density $\sigma(\mathbf{x})$ for the spatial point, before feeding into the attention layer, we additionally add a token τ_σ corresponding to obtaining sigma. τ_σ is obtained by concatenating a trainable embedding e_σ with embeddings of timestep t , followed by a linear layer.

$$\tau_\sigma = \text{Linear}(\text{Concat}[e_\sigma, e_t]), \quad (5)$$

$$\tau(\mathbf{x}) = \text{Concat}[\tau_\sigma, e^{(1)}(\mathbf{x}), e^{(2)}(\mathbf{x}), \dots, e^{(N)}(\mathbf{x})], \quad (6)$$

$$\tau'(\mathbf{x}) = \text{MultiHeadSelfAttn}(\tau(\mathbf{x})), \quad (7)$$

where `MultiHeadSelfAttn` denotes the Multi-Head Self Attention layers. After the multi-view augmentation, there are a total of $N+1$ tokens in $\tau'(\mathbf{x})$. For the sigma token $\tau'_\sigma(\mathbf{x})$, we use an MLP to obtain its volume density value $\sigma(\mathbf{x})$. For the multi-view enhanced feature tokens $\tau'(\mathbf{x})$, we conduct volume rendering to gather features enhanced on the object surface.

$$\sigma(\mathbf{x}) = \text{MLP}(\tau'_\sigma(\mathbf{x})). \quad (8)$$

$$\mathbf{F}'_k^{(i)} = \text{Volume_Rendering}(\sigma(\mathbf{x}), \tau'^{(i)}(\mathbf{x})) \quad (9)$$

We directly add the feature map to the original feature map. The enhanced feature map continues to function in the decoder of the diffusion model, ultimately producing the model's predicted noise ϵ , which is then integrated into diffusion training and inference.

4.2.3 Training. The training comprises two stages. In the first stage, we independently train the volume density generation module, enabling the selection of arbitrary camera viewpoints, rather than being constrained to predefined fixed camera viewpoints. Subsequently, we can calculate the depth of each spatial point relative to the selected camera. Following volume rendering, we supervise the generated occupancy map $M^{(i)}$ and depth maps $D^{(i)}$, which are denoted as

$$\mathcal{L}_\sigma = \mathcal{L}_{\text{BCE}}(M^{(i)}, M_{\text{gt}}^{(i)}) + \lambda(1 - \rho(D^{(i)}, D_{\text{gt}}^{(i)})), \quad (10)$$

where, \mathcal{L}_{BCE} denotes the Binary Cross Entropy loss, ρ represents the Pearson correlation coefficient, and λ controls the balance between the two losses. For training the feature map augmentation, we use the standard diffusion loss described in Equation (1).

4.3 Balanced Score Distillation

In this chapter, our exploration of score distillation is conducted through 2D generation experiments, where we directly optimize the pixel values of a 2D image. The use of 2D experiments allows the exclusion of various factors unrelated to score distillation algorithms, such as camera viewpoints and 3D representation selection, thereby highlighting the effectiveness of score distillation methods. Furthermore, the outcomes of 2D experiments can effectively represent the effects observed in 3D experiments [11, 53]. Our focus on the experimental results is specifically directed toward the richness of details and color saturation at two distinct time points: the optimal state of image quality (referred to as the convergence state) and the period of sustained training after reaching the convergence state (referred to as the over-trained state). In 3D generation, the lack of an early-stop mechanism frequently results in the widespread occurrence of over-training.

4.3.1 Decomposing Score Distillation. We adopt a modeling approach different from SDS [37] to obtain our decomposition. Score distillation is an optimization-based generation method aimed at distilling knowledge from pre-trained 2D diffusion models to guide other generations. In order to utilize 2D diffusion models for guidance, it is necessary to establish a connection between the 2D image distribution p_0 modeled by diffusion models and the 3D representation distribution q . We make the assumption that the probability

density distribution $q(\theta|y)$ for the parameters θ of the 3D representation, conditioned on the text prompt y , is proportional to the product of the conditional probability densities of its rendered images \mathbf{x}_0^π under various viewpoint π s, denoted as,

$$q(\theta|y) \propto \prod_{\pi} p_0(\mathbf{x}_0^\pi(\theta)|y). \quad (11)$$

We adopt the average negative logarithm of the probability,

$$\mathcal{L}_{\text{SD}} = -\frac{1}{N} \log q(\theta|y), \quad (12)$$

as the loss for neural network training, where N is the number of π s. Consequently, as the loss decreases, we obtain a sample as described in the text prompt. We can get the gradient on θ as,

$$\nabla_{\theta} \mathcal{L}_{\text{SD}} = -\mathbb{E}_{\pi} [\nabla_{\theta} \log p_0(\mathbf{x}_0^\pi(\theta)|y)]. \quad (13)$$

However, p_0 represents the unknown distribution of general 2D images, and high-density regions of p_0 are sparsely populated [46]. Given that current diffusion models are designed to model the score function of the distribution p_t of noisy images with noise of level t , we instead leverage it to generate an optimized gradient for the noisy images \mathbf{x}_t that is obtained by adding noise of level t to \mathbf{x}_0^π . Subsequently, this gradient can be back-propagated to the image \mathbf{x}_0^π and further transmitted to the 3D parameters θ , denoted as,

$$\nabla_{\theta} \mathcal{L}_{\text{SD}} = -\mathbb{E}_{\pi, t} [\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t(\theta)|y) \frac{\partial \mathbf{x}_t}{\partial \mathbf{x}_0^\pi} \frac{\partial \mathbf{x}_0^\pi}{\partial \theta}]. \quad (14)$$

Please note that we observe a distinction between our formula and Equation (2) from SDS [37], particularly in the absence of the final term $-\epsilon$ in our formulation. The first term $\nabla_{\mathbf{x}_t} \log p_0(\mathbf{x}_t(\theta)|y)$ is a 2D score function modeled by diffusion models. In 2D diffusion generation, the score function of conditional distribution can be decomposed into the classifier guidance and the score function of an unconditional distribution:

$$\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|y) = \nabla_{\mathbf{x}_t} \log p_t(y|\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t). \quad (15)$$

So, we naturally decompose the score distillation process into two functional terms.

$$\delta_{\text{CG}} = -\nabla_{\mathbf{x}_t} \log p_t(y|\mathbf{x}_t) = (\epsilon(\mathbf{x}_t, t, y) - \epsilon(\mathbf{x}_t, t, \emptyset))/\sigma_t \quad (16)$$

$$\delta_{\text{SG}} = -\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t) = \epsilon(\mathbf{x}_t, t, \emptyset)/\sigma_t \quad (17)$$

$$-\nabla_{\mathbf{x}_t} \log p_t(\mathbf{x}_t|y) = u \cdot \delta_{\text{CG}} + v \cdot \delta_{\text{SG}}, \quad (18)$$

where u and v control the ratio between the two terms. The first term is the classifier guidance term (δ_{CG}). When only utilizing classifier guidance, referred to as CSD [59], images exhibit fine details but suffer from artifacts and over-saturation, as shown in the upper right corner of Figure 2. Therefore, the introduction of the second term aims to alleviate the artifacts caused by δ_{CG} and guide \mathbf{x}_0 towards the higher density region in the distribution of 2D images. As depicted in Figure 2, starting from \mathcal{L}_{CSD} , as u/v decreases, the optimization directions move towards the "balance" area, the generation results become smoother. Hence, we refer to it as smoothing guidance (δ_{SG}). However, when δ_{CG} starts to dominate the proportion, the image tends to become overly smoothed.




Figure 3: 2D generation results of score distillation algorithms, annotated with computational costs. "Forward" represents undergoing one forward process of the diffusion model. The results of our BSD closely resemble the text-to-image ground truth. BSD converges at the Pareto Optimal points, ensuring that its results maintain balanced saturation during over-training.

4.3.2 Balanced Score Distillation. Based on our analysis, the parameters of CFG (u/v) in score distillation algorithms play a critical role in regulating the balance between the two guidance terms during each optimization step. However, we have observed that their effectiveness becomes highly unstable after numerous optimization steps, particularly highlighted by the failure of most score distillation methods to preserve image properties during the over-training phase. We conduct experiments of previous score distillation algorithms including SDS [37], VSD [53], ISM [25], CSD [59], SSD [48]. We find that most of them exhibit over-saturation during over-trained steps, as depicted in Figure 3.

To address this issue, we investigated the distribution characteristics of δ_{CG} and δ_{SG} . We observed that over 30% of the steps in the iterative optimization process exhibit a negative dot product between δ_{CG} and δ_{SG} , indicating obtuse angles between their optimization directions in the high-dimensional space, as shown in Figure 2. In such instances, simply combining the two directions based on a fixed ratio may lead to the final optimization direction projecting negatively onto one of the term's optimization directions. Consequently, the control of balance becomes ineffective due to the presence of negative optimization, ultimately resulting in over-saturation or over-smoothing.

Therefore, we propose to consider score distillation as a multi-objective optimization task, where the optimization objectives are $\mathcal{L}^1 = -\lambda \log(p_t(y|x_t))$ and $\mathcal{L}^2 = -\log(p_t(x_t))$, with λ representing the hyper-parameter that controls the ratio of two guidance terms. The solution of multi-objective optimization tasks will find the Pareto optimal point, where no action can improve one objective without deteriorating another. If generation reaches the

Pareto Optimal points, it will stabilize at this optimal point, and the generated content will exhibit rich details while maintaining balanced saturation. In this scenario, the hyper-parameter λ genuinely controls the final inclination towards both optimization directions, rather than attempting to control the balance at each step.

We propose the Balanced Score Distillation, which employs MGDA [8]. Suppose the optimization combination of $\nabla_{x_t} \mathcal{L}^1$ and $\nabla_{x_t} \mathcal{L}^2$ is

$$\nabla_{x_t} \log \tilde{p}_t(x_t|c) = \alpha \nabla_{x_t} \mathcal{L}^1 + (1 - \alpha) \nabla_{x_t} \mathcal{L}^2. \quad (19)$$

Following Désidéri [8], α is the solution of

$$\min_{\alpha, 1-\alpha \geq 0} \left\{ \left\| \alpha \nabla_{x_t} \mathcal{L}^1 + (1 - \alpha) \nabla_{x_t} \mathcal{L}^2 \right\|_2^2 \right\}. \quad (20)$$

For binary situation like in Equation (20), we have closed form solution for α , which is $\alpha = \min(\max(0, \hat{\alpha}), 1)$, where

$$\hat{\alpha} = \frac{(\nabla_{x_t} \mathcal{L}^2 - \nabla_{x_t} \mathcal{L}^1)^T \nabla_{x_t} \mathcal{L}^2}{\|\nabla_{x_t} \mathcal{L}^2 - \nabla_{x_t} \mathcal{L}^1\|_2^2}. \quad (21)$$

Thus, the final formula for BSD is:

$$\nabla_{\theta} \mathcal{L}_{SD} = \mathbb{E}_{\pi, t} [\omega(t) (\alpha \lambda \delta_{CG} + (1 - \alpha) \delta_{SG}) \frac{\partial x_0^T}{\partial \theta}], \quad (22)$$

$$\alpha = \min \left[\max \left[0, \frac{(\delta_{SG} - \lambda \delta_{CG})^T \delta_{SG}}{\|\delta_{SG} - \lambda \delta_{CG}\|_2^2} \right], 1 \right]. \quad (23)$$

As illustrated in Figure 3, images generated by the BSD algorithm approximate the saturation of 2D images, showcasing smooth color transitions and detailed contours. In states of over-training, BSD maintains a harmonized saturation level, unlike other methods that suffer from loss of details due to over-saturation. A comparable



Figure 4: Qualitative comparison with baseline methods. More comparisons with recent methods utilizing NeRF [34] as 3D representations are provided in the Supplementary Material.

method to BSD is VSD [53], which exhibits fewer instances of over-saturation and finer details compared to other baseline methods. However, as annotated in the figure, VSD requires significantly longer computational time and higher GPU utilization than BSD.

4.3.3 Relationship with Previous Methods. We provide an analysis of the relationship between our decomposition and score distillation algorithm with previous works in the Supplementary Material.

5 EXPERIMENTS

Table 1: Results of quantitative comparison with baselines and ablation studies. The values in parentheses represent the direct scores given by the ImageReward [55] model.

	Quality	Alignment	Average
DreamFusion	24.9 (-1.255)	24.0	24.4
Magic3D	38.7 (-0.565)	35.3	37.0
LatentNeRF	34.2 (-0.790)	32.0	33.1
Fantasia3D	29.2 (-1.040)	23.5	26.4
SJC	26.3 (-1.185)	23.0	24.7
ProlificDreamer	51.1 (+0.055)	47.8	49.4
PlacidDreamer (Ours)	62.4 (+0.620)	59.8	61.1
w/o Latent-Plane	56.5 (+0.325)	54.6	55.6
w/o Initialization	57.3 (+0.365)	56.1	56.7
w/o Personalization	61.8 (+0.590)	59.5	60.7
w/o BSD	60.2 (+0.510)	57.4	58.8

5.1 Implementation Details

The Latent Plane model has $N = 8$ viewpoints. Given the resolution limitation of feature map $F^{(2)}$ to 32×32 , it becomes imperative to regard the pixel space as continuous, ensuring accurate calculation of boundary correspondence. Even a minor difference of 0.5 pixels along the boundary could lead to notable errors. For all the MLP described in Section 4, we use a single linear layer for computing efficiency. We train the Latent-Plane module with Zero-1-to-3 [29]

for 18 hours on $4 \times$ Nvidia RTX A6000 GPUs with a batch size of 32 and loss scale $\lambda = 0.05$ with Objaverse Dataset [7] rendered by SyncDreamer [30], filtered by RichDreamer [39] and LGM [49]. For more details, please refer to the Supplementary Materials.

5.2 Comparing with Baselines.

Quantitative comparison. T3Bench [10] serves as a benchmark for text-to-3D generation, standardizing camera poses for rendering images and utilizing pre-trained models [1, 23, 40, 55] to ensure a fair evaluation. The primary 3D representation employed in this benchmark is a textured mesh, which aids in the convergence of generated 3D samples, forming seamless surfaces. However, extracting meshes from 3D Gaussian representations poses challenges. Consequently, we employ 3D Gaussian splatting to generate RGB images directly. This choice challenges our generation’s quality by not eliminating noise in direct RGB image generation.

As PlacidDreamer focuses on the generation of individual objects, our evaluations are confined to the single-object part of T3Bench. Considering DreamFusion [37], Magic3D [26], Fantasia3D [4], ProlificDreamer [53], Latent-NeRF [33], and SJC [51] as baseline methods, all of which were tested by T3Bench authors, He et al. [10]. The outcomes reveal that PlacidDreamer significantly outperforms the baseline methods in both generative quality and accuracy corresponding to the provided text. Our average quality score surpasses the next-highest ProlificDreamer by 11 points. Since this score is a linear transformation of the ImageReward [55] score, we annotate the original ImageReward scores, providing a clearer perspective on the superiority of our generative quality.

Qualitative comparison. We also conduct qualitative comparisons on several baseline methods which also builds on 3D Gaussian Splatting [22], including DreamGaussian [50], GaussianDreamer [58], GSGEN [6], and LucidDreamer [25]. In Figure 4, we present the results of different methods generating the same prompts. DreamGaussian can generate a textured mesh in several minutes, thus, its generation quality is lower compared with other methods. GaussianDreamer and GSGEN leverages end-to-end 3D generation model

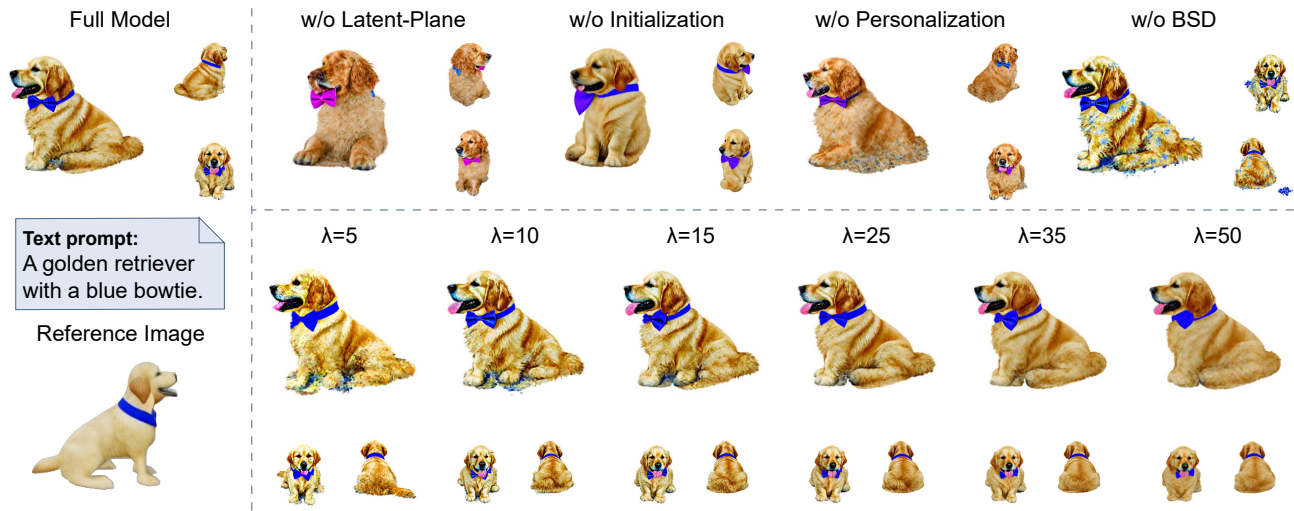


Figure 5: Results of ablation studies. In the first line, we evaluate PlacidDreamer by removing each component individually. In the second line, we investigate the impact of different λ values, validating that BSD enables stable control of the balance between color saturation and detail level.

[20, 35] for initialization with multi-view consistency. However, their final results still have multi-face problems, because the 3D model and Stable Diffusion [41] generate towards different directions. LucidDreamer’s generations feature bright colors according to the ISM algorithm and exhibit good geometric structure. However, with Latent-Plane enhancing texture and the BSD algorithm, PlacidDreamer achieves superior textures and richer details.

5.3 Ablation Studies.

We perform ablation studies, as illustrated in Figure 5, and carry out quantitative evaluations using T3Bench, as detailed in Table 1.

Latent-Plane. The “Latent-Plane” module in the pipeline serves the functions of initializing 3D Gaussian points and personalizing the diffusion model. We investigate the effects of removing both functionalities (w/o Latent-Plane), removing only initialization (w/o Initialization), and removing only fine-tuning (w/o Personalization) on the experimental results. Firstly, we remove the Latent-Plane module (w/o Latent-Plane), utilizing the Point-e [35] model as the initialization model for 3D Gaussian points, following LucidDreamer. Point-e struggles to comprehend the complex prompts and fails to generate a complete shape for the dog. Additionally, due to the color deviation introduced by Stable Diffusion during generation, the color of the dog’s tie is inaccurate. Subsequently, we only remove the initialization function of the Latent-Plane, retaining the fine-tuning part (w/o Initialization). It can be observed that the Janus problem in the dog’s head is resolved. Through extensive experimentation, we find that under the probabilistic supervision of score distillation, 3D Gaussian points do not undergo significant changes in position, making them highly sensitive to initialization. Next, we solely remove the fine-tuning part, keeping the initialization part intact (w/o Personalization). This results in a more complete structure for the dog’s bodies, but issues arise with the multi-face problem and tie color.

BSD. We replace BSD with SDS [37] and observe severe over-saturation and some blue artifacts. Compared with BSD, SDS makes

3D Gaussian points harder to converge. To further investigate BSD, we conduct ablation experiments on λ to validate its effectiveness in controlling the ratio of δ_{CG} and δ_{SG} . It can be observed that when $\lambda = 5$, the dog’s color is over-saturated. As λ increases to the range of 15-25, over-saturation alleviates without significant loss of details. Only when λ reaches the value of 35 does the loss of details become apparent. However, at this point, the dog’s features are still recognizable. It is worth noting that in the SDS algorithm, CFG can only be set to a large value to get recognizable outcomes.

These results prove that each proposed component significantly contributes to the overall effectiveness of the framework.

5.4 Experiments on BSD

BSD is a general score distillation method independent of the choice of NeRF, camera perspectives, diffusion models, etc. To test its distillation ability, we integrate BSD into various 3D generation pipelines, ensuring a fair evaluation of score distillation capabilities. This included exclusively substituting alternative score distillation algorithms, without introducing additional modifications to the pipeline. Our results demonstrate that BSD enhances the output of each pipeline. Further details and results are provided in the Supplementary Material.

6 CONCLUSION

We have focused on resolving conflicts in current text-to-3D approaches, including conflicts within a single model’s guidance and conflicts arising from guidance provided by different models. To address these conflicts, we introduced a novel framework PlacidDreamer with the newly designed Latent-Plane module and the Balanced Score Distillation algorithm to achieve mutual optimization. As a result, we have achieved a more harmonious and balanced text-to-3D generation process, leading to high-fidelity 3D outcomes. We hope our work will inspire further investigations into more harmonious methodologies for 3D generation.

REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. 2023. DreamAvatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv preprint arXiv:2304.00916* (2023).
- [3] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas J Guibas, Jonathan Tremblay, Sameh Khamis, et al. 2022. Efficient geometry-aware 3d generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 16123–16133.
- [4] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. 2023. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873* (2023).
- [5] Yiwen Chen, Chi Zhang, Xiaofeng Yang, Zhongang Cai, Gang Yu, Lei Yang, and Guosheng Lin. 2023. It3d: Improved text-to-3d generation with explicit view synthesis. *arXiv preprint arXiv:2308.11473* (2023).
- [6] Zilong Chen, Feng Wang, and Huaping Liu. 2023. Text-to-3d using gaussian splatting. *arXiv preprint arXiv:2309.16585* (2023).
- [7] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. 2023. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 13142–13153.
- [8] Jean-Antoine Désidéri. 2012. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *Comptes Rendus Mathématique* 350, 5-6 (2012), 313–318.
- [9] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. In *Proceedings of Advances in Neural Information Processing Systems*.
- [10] Yuze He, Yushi Bai, Matthieu Lin, Wang Zhao, Yubin Hu, Jenny Sheng, Ran Yi, Juanzi Li, and Yong-Jin Liu. 2023. T³Bench: Benchmarking Current Progress in Text-to-3D Generation. *arXiv:2310.02977* [cs.CV]
- [11] Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. 2023. Delta denoising score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2328–2337.
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. In *Proceedings of Advances in Neural Information Processing Systems*.
- [13] Yicong Hong, Kai Zhang, Juxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. 2023. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400* (2023).
- [14] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).
- [15] Shuo Huang, Zongxin Yang, Liangting Li, Yi Yang, and Jia Jia. 2023. AvatarFusion: Zero-shot Generation of Clothing-Decoupled 3D Avatars Using 2D Diffusion. In *Proceedings of the 31st ACM International Conference on Multimedia*. 5734–5745.
- [16] Yukun Huang, Jianan Wang, Yukai Shi, Xianbiao Qi, Zheng-Jun Zha, and Lei Zhang. 2023. DreamTime: An Improved Optimization Strategy for Text-to-3D Content Creation. *arXiv preprint arXiv:2306.12422* (2023).
- [17] Yukun Huang, Jianan Wang, Ailing Zeng, He Cao, Xianbiao Qi, Yukai Shi, Zheng-Jun Zha, and Lei Zhang. 2023. DreamWaltz: Make a Scene with Complex 3D Animatable Avatars. *arXiv preprint arXiv:2305.12529* (2023).
- [18] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. 2022. Zero-shot text-guided object generation with dream fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 867–876.
- [19] Ruixiang Jiang, Can Wang, Jingbo Zhang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. 2023. AvatarCraft: Transforming Text into Neural Human Avatars with Parameterized Shape and Pose Control. *arXiv preprint arXiv:2303.17606* (2023).
- [20] Heewoo Jun and Alex Nichol. 2023. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463* (2023).
- [21] Oren Katzir, Or Patashnik, Daniel Cohen-Or, and Dani Lischinski. 2023. Noise-free score distillation. *arXiv preprint arXiv:2310.17590* (2023).
- [22] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics* 42, 4 (2023).
- [23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597* (2023).
- [24] Weiyu Li, Rui Chen, Xuelin Chen, and Ping Tan. 2023. Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. *arXiv preprint arXiv:2310.02596* (2023).
- [25] Yixun Liang, Xin Yang, Jiantao Lin, Haodong Li, Xiaogang Xu, and Yingcong Chen. 2023. LucidDreamer: Towards High-Fidelity Text-to-3D Generation via Interval Score Matching. *arXiv:2311.11284* [cs.CV]
- [26] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiao-hui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 300–309.
- [27] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. 2023. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. *arXiv preprint arXiv:2309.17261* (2023).
- [28] Minghua Liu, Ruoxi Shi, Linghao Chen, Zhuoyang Zhang, Chao Xu, Xinyue Wei, Hansheng Chen, Chong Zeng, Jiayuan Gu, and Hao Su. 2023. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. *arXiv preprint arXiv:2311.07885* (2023).
- [29] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. 2023. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9298–9309.
- [30] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. 2023. SyncDreamer: Generating Multiview-consistent Images from a Single-view Image. *arXiv preprint arXiv:2309.03453* (2023).
- [31] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. 2023. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv preprint arXiv:2310.15008* (2023).
- [32] Jonathan Lorraine, Kevin Xie, Xiaohui Zeng, Chen-Hsuan Lin, Towaki Takikawa, Nicholas Sharp, Tsung-Yi Lin, Ming-Yu Liu, Sanja Fidler, and James Lucas. 2023. Att3d: Amortized text-to-3d object synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 17946–17956.
- [33] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. 2023. Latent-nerf for shape-guided generation of 3d shapes and textures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12663–12673.
- [34] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Commun. ACM* 65, 1 (2021), 99–106.
- [35] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. 2022. Point-e: A system for generating 3d point clouds from complex prompts. *arXiv preprint arXiv:2212.08751* (2022).
- [36] Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*. PMLR, 8162–8171.
- [37] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. 2022. DreamFusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988* (2022).
- [38] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. 2023. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843* (2023).
- [39] Lingteng Qiu, Guanying Chen, Xiaodong Gu, Qi zuo, Mutian Xu, Yushuang Wu, Weihao Yuan, Zilong Dong, Liefeng Bo, and Xiaoguang Han. 2023. RichDreamer: A Generalizable Normal-Depth Diffusion Model for Detail Richness in Text-to-3D. *arXiv preprint arXiv:2311.16918* (2023).
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*. PMLR, 8748–8763.
- [41] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10684–10695.
- [42] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. 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*. Springer, 234–241.
- [43] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. 2023. Zeronvs: Zero-shot 360-degree view synthesis from a single real image. *arXiv preprint arXiv:2310.17994* (2023).
- [44] Yichun Shi, Peng Wang, Jianglong Ye, Mai Long, Kejie Li, and Xiao Yang. 2023. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512* (2023).
- [45] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of International Conference on Machine Learning*.
- [46] Yang Song and Stefano Ermon. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* 32 (2019).
- [47] Jingxiang Sun, Bo Zhang, Ruizhi Shao, Lizhen Wang, Wen Liu, Zhenda Xie, and Yebin Liu. 2023. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818* (2023).

929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044

- 1045 [48] Boshi Tang, Jianan Wang, Zhiyong Wu, and Lei Zhang. 2023. Stable Score
1046 Distillation for High-Quality 3D Generation. *arXiv preprint arXiv:2312.09305*
1047 (2023). 1103
- 1048 [49] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and
1049 Ziwei Liu. 2024. LGM: Large Multi-View Gaussian Model for High-Resolution
1050 3D Content Creation. *arXiv preprint arXiv:2402.05054* (2024). 1104
- 1051 [50] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. 2023. Dream-
1052 gaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv*
1053 *preprint arXiv:2309.16653* (2023). 1105
- 1054 [51] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich.
1055 2023. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d
1056 generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
1057 *Pattern Recognition*. 12619–12629. 1106
- 1058 [52] Peihao Wang, Zhiwen Fan, Dejie Xu, Dilin Wang, Sreyas Mohan, Forrest Iandola,
1059 Rakesh Ranjan, Yilei Li, Qiang Liu, Zhangyang Wang, et al. 2023. SteinDreamer:
1060 Variance Reduction for Text-to-3D Score Distillation via Stein Identity. *arXiv*
1061 *preprint arXiv:2401.00604* (2023). 1107
- 1062 [53] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun
1063 Zhu. 2023. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation
1064 with Variational Score Distillation. *arXiv preprint arXiv:2305.16213* (2023). 1108
- 1065 [54] Jinbo Wu, Xiaobo Gao, Xing Liu, Zhengyang Shen, Chen Zhao, Haocheng Feng,
1066 Jingtuo Liu, and Errui Ding. 2024. Hd-fusion: Detailed text-to-3d generation
1067 leveraging multiple noise estimation. In *Proceedings of the IEEE/CVF Winter*
1068 *Conference on Applications of Computer Vision*. 3202–3211. 1109
- 1069 [55] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie
1070 Tang, and Yuxiao Dong. 2023. Imagereward: Learning and evaluating human
1071 preferences for text-to-image generation. *arXiv preprint arXiv:2304.05977* (2023). 1110
- 1072 [56] Yuanyou Xu, Zongxin Yang, and Yi Yang. 2023. SEEAAvatar: Photorealistic Text-
1073 to-3D Avatar Generation with Constrained Geometry and Appearance. *arXiv*
1074 *preprint arXiv:2312.08889* (2023). 1105
- 1075 [57] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. 2023. Consist-
1076 net: Enforcing 3d consistency for multi-view images diffusion. *arXiv preprint*
1077 *arXiv:2310.10343* (2023). 1106
- 1078 [58] Taoran Yi, Jiemin Fang, Guanjun Wu, Lingxi Xie, Xiaopeng Zhang, Wenyu Liu,
1079 Qi Tian, and Xinggang Wang. 2023. Gaussiandreamer: Fast generation from text
1080 to 3d gaussian splatting with point cloud priors. *arXiv preprint arXiv:2310.08529*
1081 (2023). 1107
- 1082 [59] Xin Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Song-Hai Zhang, and Xi-
1083 aojuan Qi. 2023. Text-to-3d with classifier score distillation. *arXiv preprint*
1084 *arXiv:2310.19415* (2023). 1108
- 1085 [60] Yifei Zeng, Yuanxun Lu, Xinya Ji, Yao Yao, Hao Zhu, and Xun Cao. 2023. Avatar-
1086 Booth: High-Quality and Customizable 3D Human Avatar Generation. *arXiv*
1087 *preprint arXiv:2306.09864* (2023). 1109
- 1088 [61] Minda Zhao, Chaoyi Zhao, Xinyue Liang, Lincheng Li, Zeng Zhao, Zhipeng Hu,
1089 Changjie Fan, and Xin Yu. 2023. EfficientDreamer: High-Fidelity and Robust 3D
1090 Creation via Orthogonal-view Diffusion Prior. *arXiv preprint arXiv:2308.13223*
1091 (2023). 1110
- 1092 [62] Joseph Zhu and Peiye Zhuang. 2023. HiFA: High-fidelity Text-to-3D with Ad-
1093 vanced Diffusion Guidance. *arXiv preprint arXiv:2305.18766* (2023). 1111
- 1094 1112
- 1095 1113
- 1096 1114
- 1097 1115
- 1098 1116
- 1099 1117
- 1100 1118
- 1101 1119
- 1102 1120
- 1121
- 1122
- 1123
- 1124
- 1125
- 1126
- 1127
- 1128
- 1129
- 1130
- 1131
- 1132
- 1133
- 1134
- 1135
- 1136
- 1137
- 1138
- 1139
- 1140
- 1141
- 1142
- 1143
- 1144
- 1145
- 1146
- 1147
- 1148
- 1149
- 1150
- 1151
- 1152
- 1153
- 1154
- 1155
- 1156
- 1157
- 1158
- 1159
- 1160