

# AnchorDS: Anchoring Dynamic Sources for Semantically Consistent Text-to-3D Generation

Jiayin Zhu<sup>1</sup>, Linlin Yang<sup>2</sup>, Yicong Li<sup>1\*</sup>, Angela Yao<sup>1</sup>

<sup>1</sup>National University of Singapore

<sup>2</sup>Communication University of China

zhujiayin@u.nus.edu, lyang@cuc.edu.cn, liyicong@u.nus.edu, ayao@comp.nus.edu.sg

## Abstract

Optimization-based text-to-3D methods distill guidance from 2D generative models via Score Distillation Sampling (SDS), but implicitly treat this guidance as static. This work shows that ignoring source dynamics yields inconsistent trajectories that suppress or merge semantic cues, leading to “semantic over-smoothing” artifacts. As such, we reformulate text-to-3D optimization as mapping a *dynamically evolving source* distribution to a fixed target distribution. We cast the problem into a dual-conditioned latent space, conditioned on both the text prompt and the intermediately rendered image. Given this joint setup, we observe that the image condition naturally anchors the current source distribution. Building on this insight, we introduce AnchorDS, an improved score distillation mechanism that provides state-anchored guidance with image conditions and stabilizes generation. We further penalize erroneous source estimates and design a lightweight filter strategy and fine-tuning strategy that refines the anchor with negligible overhead. AnchorDS produces finer-grained detail, more natural colours, and stronger semantic consistency, particularly for complex prompts, while maintaining efficiency. Extensive experiments show that our method surpasses previous methods in both quality and efficiency.

**Code** — <https://github.com/viridityzhu/AnchorDS>

**Extended version** — <https://arxiv.org/abs/2511.11692>

## 1 Introduction

With the growing demand for 3D content creation in gaming and virtual reality, text-to-3D generation has emerged as a significant research frontier. One prominent solution that builds on this trend is Score Distillation Sampling (SDS) (Poole et al. 2023). SDS can leverage powerful off-the-shelf 2D diffusion models to guide optimization-based text-to-3D generation without large-scale 3D datasets.

SDS-guided text-to-3D generation can be interpreted as an optimization process that gradually shifts the distribution of rendered images from the current 3D representation (*i.e.*, the source distribution) toward a distribution defined by a text-conditioned diffusion model (*i.e.*, the target distribution) (McAllister et al. 2024). Under this interpretation,

\*Corresponding author.

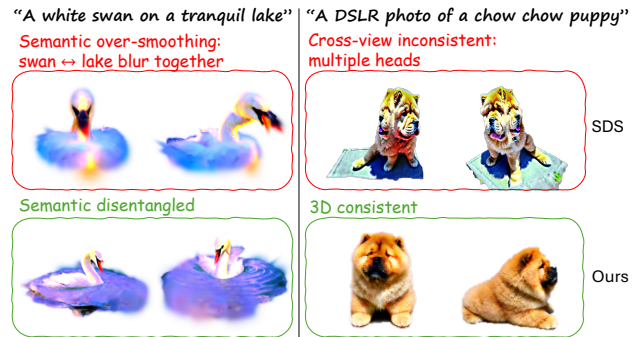


Figure 1: Comparison of Vanilla SDS vs. Ours. SDS suffers from semantic over-smoothing (mixing swan and lake) and cross-view inconsistency (multiple heads). Ours achieves semantic disentanglement and 3D consistency across views.

follow-up works have proposed enhancements follow two strategies. The first enhances the target distribution estimation by incorporating additional conditions, such as multiview images (Shi et al. 2023; Li et al. 2025) or depth maps (Qiu et al. 2024). The second addresses source distribution mismatch by refining the guidance, using reference-based scoring (Hertz, Aberman, and Cohen-Or 2023), enhanced negative prompts (McAllister et al. 2024), or variational optimization (Wang et al. 2023). Despite these advancements, SDS-based 3D generation still struggles with consistently preserving structural semantics across iterative updates. This manifests as two critical artifacts: (1) semantic over-smoothing, where object-specific features degenerate into homogenized, semantically ambiguous representations; and (2) cross-view inconsistency, where geometry and appearance are incoherent across perspectives (Fig. 1).

At its core, SDS optimization is formulated as a distribution transformation: it progressively shifts the rendered appearance of the 3D asset—from a *source distribution* reflecting the current 3D state toward a *target distribution* defined by a text-conditioned diffusion model. However, while the target distribution remains fixed to the text prompt, the source distribution is approximated using a static unconditional prior throughout optimization. This critically overlooks the *inherently non-stationary nature of the source dis-*

*tribution*: as the 3D representation evolves through optimization steps, rendered images continuously update, dynamically altering the source distribution. Consequently, SDS effectively discards accumulated structural information at each step by restarting from a semantics-agnostic prior. This fundamental mismatch between the static guidance mechanism and the evolving asset state destabilizes optimization, ultimately manifesting as the observed artifacts.

To address this issue, we propose to reframe score distillation as a dynamic editing process. Rather than treating generation as a one-shot projection from a static source, we view it as a progressive editing loop in which each optimization step refines the current 3D state based on the accumulated guidance from previous steps. Crucially, by explicitly recognizing and exploiting the evolving 3D state itself, our reformulation yields two advantages: 1) it supplies a stream of rich cues—geometry, colour, and semantics—that stabilise guidance and enforce cross-view consistency; and 2) feeding the state back into a pretrained conditional diffusion model enables dynamic, accurate, and lightweight source estimates without extra networks or handcrafted prompts.

Building on this perspective, we introduce AnchorDS, a dynamic form of SDS that anchors the source distribution using the rendered image at each optimization step. Specifically, we leverage a dual-conditioned diffusion model that incorporates both the text prompt and the intermediate image as guidance signals. Crucially, we observe that the diffusion model’s predicted noise conditioned on the current rendering naturally encodes structural and semantic cues from the image condition. This noise prediction inherently anchors the source distribution by correlating the gradient guidance with the evolving 3D state, mitigating distribution drift without explicit constraints. Notably, the image condition does not constrain the target output directly, but serves as a contextual anchor that steers the generation. This design is not sensitive to the selection of the dual-conditioned model. Our method is robust when utilizing various popular image-conditioned adapters, *e.g.* IP-Adapter (Ye et al. 2023) and ControlNet (Zhang, Rao, and Agrawala 2023).

Despite this flexibility, anchoring dynamic sources is still non-trivial because the implicit latent space of the diffusion model exhibits a distribution mismatch with rendered images. To bridge the gap, we incorporate two complementary components. First, we reconstruct a pseudo-source image at each step, offering a metric for evaluating the quality of the estimated source distribution. Second, we introduce two practical enhancements: a simple yet effective *Filtering* mechanism that discards unreliable source predictions, and a lightweight *Fine-tuning* strategy that better aligns the diffusion model with the domain of rendered images.

Our contributions are summarized as follows:

1. We reveal the evolving nature of the source distribution in SDS and identify it as the root cause of semantic over-smoothing and inconsistent optimization trajectories.
2. We propose AnchorDS, a novel score distillation framework that dynamically anchors the source estimation via a dual-conditioned diffusion model. We further introduce a filtering mechanism and a lightweight fine-tuning strat-

egy to better regularize the evolving source distribution.

3. Extensive experiments on T<sup>3</sup>Bench (He et al. 2023) and a representative suite of challenging prompts demonstrate that our method outperforms state-of-the-art (SoTA) SDS variants in both generation quality and efficiency.

## 2 Related Works

**Text-Guided 3D Generation.** DreamFusion demonstrates that a NeRF can be trained from scratch by following Score Distillation Sampling (SDS) gradients derived from 2D diffusions (Poole et al. 2023). Subsequent works keep the same optimization-from-noise paradigm while improving efficiency or fidelity (Lin et al. 2023; Tang et al. 2024b; Yi et al. 2024a; Li et al. 2024). Although these methods deliver the highest visual quality, they remain slow and are still vulnerable to view inconsistency and Janus problems. To eliminate per-scene optimization, another line of work trains generators that map text directly to 3D latents (Jun and Nichol 2023; Nichol et al. 2022; Siddiqui et al. 2024; Xiang et al. 2024). While generation is fast, training quality rely on large-scale 3D datasets (Deitke et al. 2022; Fu et al. 2021; Collins et al. 2022), and the outputs still trail optimization methods in geometric and color accuracy (He et al. 2023). Hybrid strategies therefore combine a feed-forward initialization with subsequent Gaussian/NeRF refinement (Liang et al. 2024; Yi et al. 2024b,a).

**Conditional and Controllable Text-to-3D.** Extending the diffusion prior with extra modalities improves controllability. ControlNet (Zhang, Rao, and Agrawala 2023) has been adopted for depth, normal, or multi-view constraints (Huang et al. 2024; Li et al. 2025). IP-Adapter (Ye et al. 2023) let users steer style or identity with a reference image and have been plugged into 3D pipelines (Zeng et al. 2023). All these works, however, treat the additional image as an external positive condition supplied a priori. We instead employs the intermediate rendering itself as a dynamic *source* anchor supplying self-consistent guidance.

**Refining Score Distillation Sampling.** Recent work revisits SDS from theoretical and practical perspectives. (Yu et al. 2023; Tang et al. 2024a; Katzir et al. 2024) separate mode-seeking and variance terms to stabilize optimization. (Liang et al. 2024; Lukoianov et al. 2024) avoid first-order errors through DDIM sampling and inversion. Mismatch remedies attempt to align the diffusion prior with the 3D asset. DDS pairs each original image with a reference prompt (Hertz, Aberman, and Cohen-Or 2023), SDS-Bridge introduces a handcrafted prompt describing the poor 3D state (McAllister et al. 2024), and ProlificDreamer trains a LoRA branch to approximate the particle distribution (Wang et al. 2023). These solutions improve stability yet rely on static prompts, specific references, or auxiliary networks, which introduce new bias and overhead. Instead, we remove this dependency by directly feeding the current rendering into the diffusion prior, yielding faithful and bias-free guidance efficiently.

### 3 Analysis on the Issue of Source Distribution Estimation

We begin by revisiting the formulation of text-to-3D generation via Score Distillation Sampling (SDS) (Poole et al. 2023), and highlight its key limitation - a lack of awareness of the rendered appearance from the current 3D state. We then draw connections between SDS and 2D editing paradigms, and reinterpret SDS as a dynamic editing process that conditions on the evolving source.

#### 3.1 Preliminaries

**Score Distillation Sampling (SDS).** SDS leverages pre-trained 2D diffusion models to optimize the 3D generation. Specifically, it applies the denoising process of a 2D diffusion model to a rendered image from the 3D model, through which it distills a prior on the generated 3D output. Given a sampled noise  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  and a latent representation  $z$  of the rendered image, the noisy latent at timestep  $t$  is given by

$$z_t = \sqrt{\bar{\alpha}_t} z + \sqrt{1 - \bar{\alpha}_t} \epsilon. \quad (1)$$

The corresponding noisy image is then used to compute the SDS gradient  $\nabla_{\Theta} \mathcal{L}_{SDS}(\phi, z)$ , where the predicted noise is assumed to approximate the score function. Concretely, following the score matching view, the diffusion model  $\phi$ 's noise prediction  $\hat{\epsilon}_{\phi}$  is defined as:

$$\hat{\epsilon}_{\phi} = -\sigma_t \nabla_{z_t} \log p(z_t; t, c), \quad \text{with } \sigma_t = \sqrt{1 - \bar{\alpha}_t}, \quad (2)$$

where  $c$  denotes conditioning information (typically a text prompt  $y$ , or a combination  $c = \{y, I\}$  when an image condition  $I$  is included). This formulation directs the optimization toward regions of high density in the conditional distribution  $p(z_t; t, c)$ . Applied to a 3D model parameterized by  $\Theta$ , the gradient is given as:

$$\nabla_{\Theta} \mathcal{L}_{SDS}(\phi, z) = w(t) \left( \hat{\epsilon}_{\phi}^{CFG}(z_t; t, c) - \epsilon \right) \frac{\partial z_t}{\partial \Theta}, \quad (3)$$

where  $\hat{\epsilon}_{\phi}^{CFG}(z_t; t, c)$  is the noise estimate (Eq. 2) under Classifier-Free Guidance (CFG).

**Classifier-Free Guidance (CFG).** CFG (Ho and Salimans 2022) balances conditional and unconditional predictions through adjustable weights. For a single-condition (text-only) setup, the CFG prediction is:

$$\hat{\epsilon}_{\phi}^{CFG}(z_t, t, y) = (1 + \omega) \hat{\epsilon}_{\phi}(z_t, t, y) - \omega \hat{\epsilon}_{\phi}(z_t, t, \emptyset), \quad (4)$$

where  $\omega$  controls the strength of the guidance.

#### 3.2 The Issue of Static Source Estimation

We first analyze why static source modeling in SDS causes artifacts (Fig. 1). Substituting Eq. 4 into Eq. 3 leads to two key terms  $m_1$  and  $m_2$ :

$$\begin{aligned} \hat{\epsilon}_{\phi}^{CFG}(z_t; t, c) - \epsilon &= (\omega - 1) \cdot \underbrace{(\hat{\epsilon}_{\phi}(z_t, t, c) - \hat{\epsilon}_{\phi}(z_t, t, \emptyset))}_{m_1} \\ &\quad + \underbrace{\hat{\epsilon}_{\phi}(z_t, t, \emptyset) - \epsilon}_{m_2}. \end{aligned} \quad (5)$$

A sufficiently large scaling of  $\omega$ , e.g.,  $\omega = 100$  will cause  $m_1$  to dominate the guidance, while  $m_2$  reduces variance of




Figure 2: Static vs. Dynamic Pseudo-Editing. Static source estimation leverages an unconditional prior, misaligned with the actual 3D state. Instead, dynamic pseudo-sources reflect evolving 3D renderings, ensuring faithful guidance.

$\hat{\epsilon}_{\phi}^{CFG}$  (McAllister et al. 2024; Tang et al. 2024a). Interpreting  $m_1$  as a score function:

$$m_1 = -\sigma_t \left( \nabla_{z_t} \log p(z_t; t, c) - \nabla_{z_t} \log p(z_t; t, \emptyset) \right), \quad (6)$$

and we see SDS pushes samples toward the higher-density regions of the conditional distribution  $p(z_t; t, c)$  while moving away from the unconditional prior  $p(z_t; t, \emptyset)$ .

This mirrors 2D image editing (e.g., DDIB (Su et al. 2023), SDEdit (Meng et al. 2022)), where edits aim to find the optimal mapping from a source image to a target image as two distributions. However, SDS approximates the source distribution—which should reflect the current 3D state—using the unconditional prior  $p(z_t; t, \emptyset)$ . The limitation becomes clear when rearranging Eq. 6 into a pseudo-editing formulation. Specifically, we invert Eq. 1 to define the pseudo-target and pseudo-source latent reconstructions:

$$\begin{aligned} \hat{z}_{t \rightarrow 0}^{\text{target}} &= \frac{1}{\sqrt{\bar{\alpha}_t}} \left( z_t - \sqrt{1 - \bar{\alpha}_t} \hat{\epsilon}_{\phi}(z_t, t, y) \right), \\ \hat{z}_{t \rightarrow 0}^{\text{source}} &= \frac{1}{\sqrt{\bar{\alpha}_t}} \left( z_t - \sqrt{1 - \bar{\alpha}_t} \hat{\epsilon}_{\phi}(z_t, t, \emptyset) \right). \end{aligned} \quad (7)$$

Then, the core update term simplifies elegantly as:

$$m_1 = \eta \left( \hat{z}_{t \rightarrow 0}^{\text{target}} - \hat{z}_{t \rightarrow 0}^{\text{source}} \right), \quad \text{where } \eta = \frac{\sqrt{\bar{\alpha}_t}}{\sqrt{1 - \bar{\alpha}_t}}. \quad (8)$$

Examining  $\hat{z}_{t \rightarrow 0}^{\text{source}}$ , we reveal that it loses critical information about the original rendering: (1)  $z_t$  is a noise-corrupted version that discards semantic content, while (2)  $\hat{\epsilon}_{\phi}(z_t, t, \emptyset) = -\sigma_t \nabla_{z_t} \log p(z_t; t, \emptyset)$  pushes toward an unconditional image prior that averages over diverse natural images. Crucially, neither term encodes the evolving 3D asset's current state, forcing  $\hat{z}_{t \rightarrow 0}^{\text{source}}$  to inherit this static, averaged characteristic rather than reflecting the dynamic source distribution. As a result, it fails to capture the semantics of intermediate renderings. The resulting update directions become inconsistent with the 3D asset's actual appearance, causing over-smoothing and cross-view inconsistency in Fig. 1.

Rather than static, we argue that the source should evolve with the 3D model (Fig. 2). We therefore reinterpret the text-to-3D optimization as a *dynamic edit*: at each step, we refine the 3D asset by (1) preserving favorable attributes (e.g., structural semantics consistent with current state) via faithful source modeling, and (2) correcting undesirable features (e.g., deviations from target distribution) through prompt alignment. This resolves the inconsistency in Fig. 1 by anchoring updates to the actual state rather than a static prior.

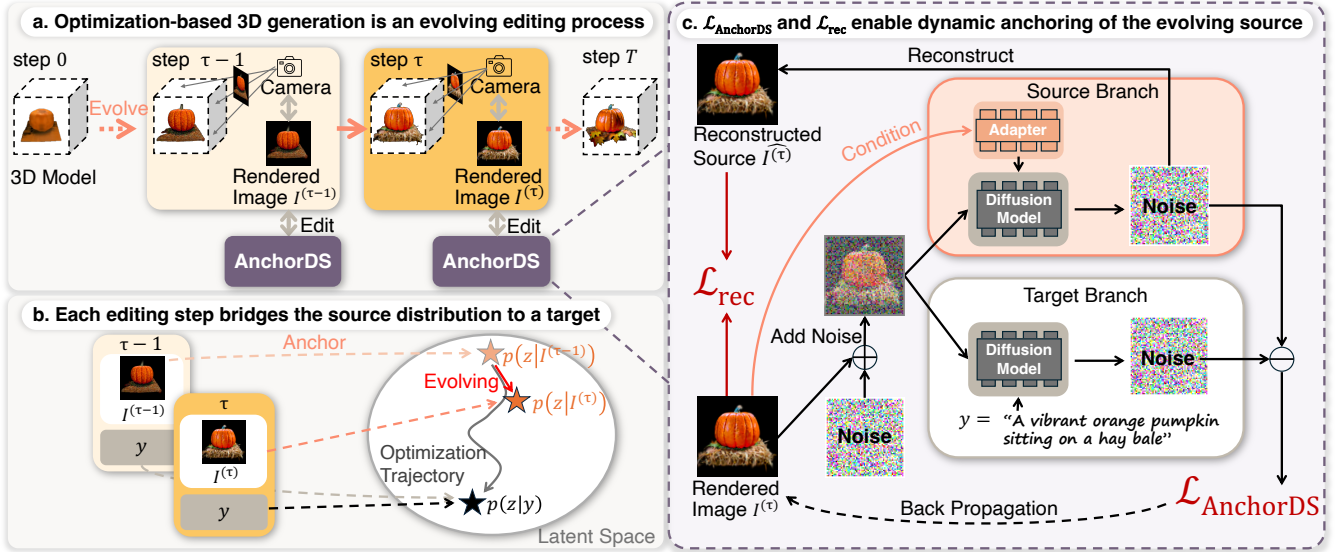


Figure 3: Overview of Our AnchorDS for Text-to-3D Generation. (a) Optimization-based 3D generation as an evolving editing process where each step refines the 3D model guided by AnchorDS. (b) Each step bridges source to target distributions through dynamic anchoring in latent space. (c) Technical framework showing how  $\mathcal{L}_{AnchorDS}$  and  $\mathcal{L}_{rec}$  enable dynamic source anchoring, with source and target branches processing the evolving 3D content through pretrained diffusion models.

## 4 Method

### 4.1 Dynamic Evolution of the Source Distribution

We formalize a critical insight regarding SDS-based text-to-3D generation: the source image distribution  $P_\theta(x)$  evolves dynamically throughout optimization rather than remaining static, as visualized in Fig. 3b. Mathematically, we reframe SDS optimization as a progressive transformation process:

$$P_\theta^{(0)}(x) \xrightarrow{\nabla_{\Theta} \mathcal{L}_{SDS}} P_\theta^{(1)}(x) \rightarrow \dots \rightarrow P_\theta^{(T)}(x), \quad (9)$$

where  $P_\theta^{(\tau)}(x)$  denotes the time-varying source distribution at optimization step  $\tau$ . Initially,  $P_\theta^{(0)}(x)$  resembles a diffuse unconditional prior when the 3D model  $\Theta$  is randomly initialized. At each iteration, the gradient update  $\nabla_{\Theta} \mathcal{L}_{SDS}$  transports probability mass such that:

$$P_\theta^{(\tau+1)}(x) = \mathcal{T} \left( P_\theta^{(\tau)}(x), \nabla_{\Theta} \mathcal{L}_{SDS}(z, y, \hat{\epsilon}_t) \right), \quad (10)$$

where  $\mathcal{T}(\cdot)$  represents the distributional shift operator. This process progressively morphs  $P_\theta^{(\tau)}(x)$  toward the target distribution  $P_{\text{target}}(x|y)$  defined by the text prompt  $y$ . The continuous distributional evolution stems from the mass transport interpretation of score-based dynamics (De Bortoli et al. 2021), establishing this formulation’s generality across score-distillation-based 3D generation methods.

Next, we introduce our AnchorDS to ensure accurate source distribution estimation at every optimization step.

### 4.2 Score Distillation via Dynamic Source Anchoring

AnchorDS dynamically anchors the score guidance at the evolving 3D state through image conditioning. Formally, at

optimization step  $\tau$  with rendered view  $I^{(\tau)}$ , we compute the guidance gradient as:

$$g_t^{(\tau)} = \hat{\epsilon}_\phi(z_t; t, y) - \hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)}), \quad (11)$$

where  $\hat{\epsilon}_\phi(z_t; t, y)$  targets the text-conditioned distribution, while  $\hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)})$  anchors the current source distribution  $P_\theta^{(\tau)}(x)$ . This differential formulation directs updates from the current state toward the target distribution.

Essentially, incorporating an image condition  $I^{(\tau)}$  recasts the problem into a dual-conditioned latent space. This preserves the text-conditioned target’s effectiveness while leveraging image conditions to anchor the source distribution. Specifically,  $I^{(\tau)}$  anchors the current source by providing structural grounding without content constraints, enabling contextual editing rather than output restriction. It yields two principal advantages: (1) Conditioning with  $I^{(\tau)}$  at continuous  $\tau$  maintains alignment between guidance and 3D state, mitigating drift and oversmoothing. (2) Source anchoring requires only a single additional U-Net forward pass per iteration, processed in parallel with the original pass, thus maintaining identical runtime to standard SDS.

AnchorDS bridges 2D editing principles with 3D generation through a key insight: pretrained diffusion models inherently possess image inversion capabilities within their conditional architecture. While 2D editing requires explicit inversion techniques, AnchorDS elegantly leverages this intrinsic property—directly utilizing the model’s natural ability to map images to latent distributions, achieving precise source anchoring without auxiliary inversion costs.

### 4.3 Source Anchoring via Image Conditioning

The effectiveness of AnchorDS depends critically on an accurate estimation of the current source distribution via image-conditioned diffusion models. We denote by  $I^{(\tau)} = R(\mathcal{X}_{3D}^{(\tau)})$  the rendered image of the 3D representation  $\mathcal{X}_{3D}^{(\tau)}$ . **Choice of Image Condition.** In general, one could apply a preprocessing function  $E(\cdot)$  to obtain a conditioning signal  $\bar{I}^{(\tau)} = E(I^{(\tau)})$  (e.g, converting  $I^{(\tau)}$  to a normal map, depth map, etc.) The conditioning signal must retain essential structural and semantic information from the 3D rendering while eliminating irrelevant noise. We initially consider that the identity mapping  $E(I^{(\tau)}) = I^{(\tau)}$  is particularly effective as it preserves maximal information about the current state. This aligns with evidence (Kadosh et al. 2025) that image-conditioned diffusion models learn invertible mappings between images and noise—crucial for source estimation. Meanwhile, we find that our method remains compatible with alternative signals (e.g, normal maps) provided that they sufficiently capture the 3D content’s core attributes.

**Pseudo-Source Reconstruction.** With the image-conditioned model in place, we can obtain an explicit estimate of the current source image distribution at any diffusion step. Given a noisy latent  $z_t$  (at noise level  $t$ ) that was produced from the current image  $I^{(\tau)}$ , the model’s image conditioned prediction  $\hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)})$  allows us to reconstruct a pseudo-reconstructed source image:

$$\hat{z}_{t \rightarrow 0}^{\text{anchored}, (\tau)} = \frac{1}{\sqrt{\alpha_t}} \left( z_t - \sqrt{1 - \alpha_t} \hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)}) \right), \quad (12)$$

which is the model’s one-step estimate of the denoised latent at timestep  $t$  – essentially a guess of the original image  $I^{(\tau)}$  (latent) before noise was added. With an image decoder  $\varepsilon(\cdot)$ , we can now explicitly evaluate reconstruction accuracy as:

$$\mathcal{L}_{\text{rec}} = \|\varepsilon(\hat{z}_{t \rightarrow 0}^{\text{anchored}, (\tau)}) - I^{(\tau)}\|_2^2. \quad (13)$$

This reconstruction not only provides a direct metric for source estimation quality but also enables two complementary mechanisms: (1) filtering out unstable predictions, and (2) fine-tuning the image adapter to better align with rendered images—both crucial for stable and accurate 3D optimization. We implement two strategies discussed below.

**Filtering.** A straightforward way to employ the objective (Eq. 13) is to apply a threshold-based filter to exclude unreliable source estimations:

$$\mathcal{M}_{\text{Filter}} = \begin{cases} 1 & \text{if } \mathcal{L}_{\text{rec}} < \gamma \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

$$\mathcal{L}_{\text{AnchorDS}} = \mathcal{M}_{\text{Filter}} \cdot \mathcal{L}_{\text{AnchorDS}}.$$

Here,  $\mathcal{L}_{\text{rec}}$  acts as a stabilizer, filtering out any spurious deviations in the anchored source prediction. Empirically, this translates to improved stability (no sudden jarring updates) and better preservation of existing content.

**Finetuning Image Adapter.** Another approach is to apply  $\mathcal{L}_{\text{rec}}$  to fine-tune the image adapter. Although pretrained 2D models are powerful enough to model the real-world image distribution, there remains a gap between the real-world and the synthesized rendered image distribution. Intuitively, we

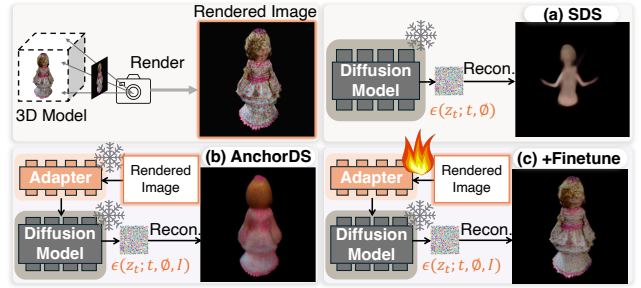


Figure 4: Source Estimation Strategies. (a) Vanilla SDS poorly reconstructs the source. (b) AnchorDS, conditioned on current rendering, recovers geometry and color faithfully. (c) Our fine-tuning accurately matches the source image.

aim to let 2D models “see” the real data. We address this through lightweight fine-tuning of the image adapter using  $\mathcal{L}_{\text{rec}}$ . This fine-tuning is minimal—only unfreezing one single layer of the image adapter is sufficient (increasing optimization time from  $\sim 25$  minutes to  $\sim 30$  minutes using 3DGS pipeline (Yi et al. 2024a)), while significantly improving source estimation accuracy (see Fig. 4c).

Minimizing  $\mathcal{L}_{\text{rec}}$  forces  $\hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)})$  to become consistent with the current image  $I^{(\tau)}$ . Note  $\hat{\epsilon}_\phi(z_t; t, \emptyset, I^{(\tau)})$  is *not* trained to match sampled noise; Instead, its role is to “invert” the current image into latent space. Eq. 12 simply uses this predicted noise to retrieve the model’s internal guess of the clean image  $I^{(\tau)}$ , and  $\mathcal{L}_{\text{rec}}$  ties that guess further. With this enhanced source estimate term, guidance  $g_t^{(\tau)}$  (Eq. 11) accurately reflects the difference between distributions of images that ensembles  $I^{(\tau)}$  and those that fulfill  $y$ .

### 4.4 Implementation

**Choice of Diffusion Model.** To incorporate  $\bar{I}^{(\tau)}$  as a condition, we leverage existing pre-trained dual-conditional diffusion models. A straightforward option is an image-to-image diffusion model such as IP2P (Brooks, Holynski, and Efros 2023). However, IP2P is fine-tuned for image editing and we observed it gives inconsistent guidance under the high CFG weights required for SDS; in practice, using IP2P led to unstable and desaturated results for text-to-3D, especially at large CFG scales. ControlNet (Zhang, Rao, and Agrawala 2023) processes derived maps (e.g, normal maps, sketches) but lacks direct training on raw images. Conversely, IP-Adapter (Ye et al. 2023) conditions Stable Diffusion (1.5) on unaltered images via an auxiliary latent, preserving the model’s expressive power without constraining image content. Crucially, our framework generalizes across adapters. We adopt IP-Adapter for primary experiments due to its direct image conditioning, while adopting ControlNet shows comparative results (detailed in Sec. 5).

**Pipeline.** Given a text prompt  $y$  and an optional initial 3D representation (which could be random or from a rough generator), our pipeline operates iteratively as illustrated in Fig. 3c. At each optimization step  $\tau$ , we render the current

Method	Base Model	CLIP ↑	Q1 ↓	Q2 ↓	Q3 ↓
VSD	SD 2.1	0.352	1.84	1.85	1.79
Ours (ControlNet)		<b>0.369</b>	<b>1.16</b>	<b>1.15</b>	<b>1.21</b>
VSD	SD 1.5	0.281	1.99	2.00	2.08
SDS-Bridge		0.233	2.38	2.35	2.29
Ours (IP-Adapter)		<b>0.334</b>	<b>1.63</b>	<b>1.66</b>	<b>1.63</b>

Q1: Which one has the best 3D consistency?

Q2: Which one shows accurately what the text describes?

Q3: Which one looks most realistic and natural?

Table 1: Quantitative Comparison. Q1–Q3 report averaged ranking of each method.

3D model from a random viewpoint to obtain image  $I^{(\tau)}$ , then encode it into the diffusion latent space and add noise at a random timestep  $t$  to produce  $z_t$ . We then query the diffusion model twice: once with text condition  $y$  to obtain the target prediction  $\hat{e}_\phi(z_t; t, y)$ , and once with an empty text and  $I^{(\tau)}$  to obtain the source prediction  $\hat{e}_\phi(z_t; t, \emptyset, I^{(\tau)})$ . The obtained AnchorDS guidance (Eq. 11) is then backpropagated through the rendering pipeline to update the 3D parameters. For the fine-tuning variant, we periodically update the unfrozen layer of the image adapter using  $\mathcal{L}_{rec}$ .

## 5 Experiments

**Experimental Setup.** We evaluate AnchorDS against vanilla SDS (Poole et al. 2023), SDS-Bridge (McAllister et al. 2024), and VSD (Wang et al. 2023) under both 3DGS (Yi et al. 2024a) and NeRF pipelines.

**Evaluation Metrics.** We follow (McAllister et al. 2024; Lee, Sohn, and Shin 2024; Dong et al. 2024) and use CLIP similarity (Radford et al. 2021) and T<sup>3</sup>Bench (He et al. 2023) for text–image alignment and visual quality, focusing on quality since text alignment is preserved by design.

### 5.1 Comparison with State-of-the-Art

To systematically evaluate SDS’s limitations, we curate 50 complex prompts from previous works (He et al. 2023; Poole et al. 2023), covering three challenging categories: fine-grained details, rich/mixed semantics, and multiple-object compositions. This evaluation scale aligns with established practices (McAllister et al. 2024; Liang et al. 2024; Lee, Sohn, and Shin 2024; Dong et al. 2024). We adopt the NeRF-based pipeline (Wang et al. 2023) across all experiments. While VSD demonstrates sensitivity to base models and typically performs reasonably only on Stable Diffusion (SD) 2.1, we evaluate on both SD 1.5 and SD 2.1. Our method (AnchorDS with Finetuning) uses IP-Adapter on SD 1.5 and ControlNet on SD 2.1 as image conditioners.

**Human Evaluation.** We conduct comprehensive human preference studies on MTurk, which comprises 20 batches covering our 50 complex prompts evaluated by 912 unique participants. For each comparison, evaluators rank all methods according to three questions: 3D consistency, text alignment, and visual quality, respectively.

**Results.** Table 1 shows our method consistently outperforms all methods in CLIP similarity and human preference evaluations. Qualitative comparisons in Figs. 5, 6, and 7 show

Method	All↑	Single↑	Surr↑	Multi↑
SDS (DreamFusion)	20.5	24.9	19.3	17.3
SDS (GaussianDreamer)	29.7	42.3	26.1	20.6
AnchorDS (IP-Adapter)	30.7	43.0	24.8	24.5
+ Filter	32.8	44.1	27.9	<b>26.5</b>
+ Finetune	<b>33.3</b>	<u>45.3</u>	<u>29.0</u>	<u>25.7</u>
AnchorDS (ControlNet)	30.8	43.9	27.2	21.3
+ Filter	<u>33.2</u>	<b>46.1</b>	<b>29.4</b>	24.0
+ Finetune	32.9	45.0	28.6	25.2

Table 2: Quantitative Comparison with Baseline on T<sup>3</sup>Bench Quality Metric. AnchorDS incorporates image conditioning via IP-Adapter or ControlNet. Filter and Finetune are alternative strategies for source distribution estimation.

our consistent generation of semantically faithful and structurally accurate results, while SDS-Bridge produces artifacts introduced by biased negative prompts, and VSD fails to accurately model the image distributions. Please refer to the Appendix for theoretical discussions.

### 5.2 Baseline Comparison

**SDS Baselines.** We comprehensively evaluate on T<sup>3</sup>Bench (He et al. 2023), a comprehensive benchmark containing 300 prompts across single object generation, object with surrounding generation, and multiple object generation categories. We compare against both DreamFusion-SDS and GaussianDreamer-SDS baselines. To ensure fair evaluation, we maintain all pipeline components constant except for the score guidance, using the state-of-the-art GaussianDreamer (Yi et al. 2024a) framework as our baseline. All methods are built upon SD 1.5.

**Comparison with SDS.** Quantitative results in Table 2 demonstrate that we consistently surpass SDS in all prompt categories. This validates that AnchorDS’s 3D state awareness ensures coherent semantic paths, preventing attribute mixing. Qualitative results are in the extended version.

### 5.3 Ablation Studies

Quantitative ablation results for our key components are shown in Table 2. AnchorDS consistently outperforms vanilla SDS through source anchoring via image conditioning. Our Filtering and Finetuning strategies provide complementary approaches for enhancing source estimation accuracy, with Finetuning achieving optimal performance. Please refer to Appendix for more ablations across prior models.

### 5.4 Method Generalizability

Experiments validate robust generalizability across 3DGS and NeRF. Please find additional results in the Appendix.

## 6 Conclusion

We demonstrate that the source distribution dynamically evolves during text-to-3D optimization—a fundamental property that has been largely overlooked by existing methods. While prior efforts have focused on reducing trajectory estimation error or improving the guidance prior, they

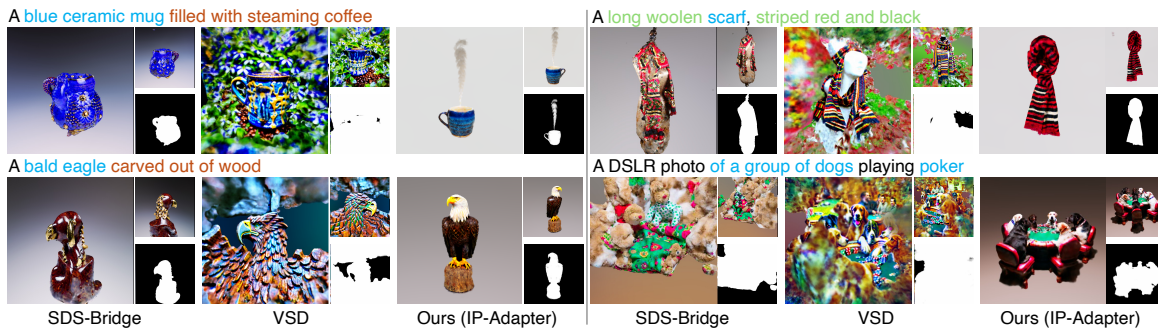


Figure 5: Qualitative Comparison on SD 1.5 on complex text prompts, including (1) **objects** with **rich or mixed semantics**; (2) **multiple** distinct **objects**; (3) **objects** with **fine-grained details**. SDS-Bridge suffers from prompt-induced texture bias and semantic inaccuracies, while VSD produces incoherent structures and exaggerated colours; our method delivers semantically faithful and structurally accurate results.



Figure 6: Qualitative Comparison on SD 2.1. Our method surpasses VSD, consistently capturing complex semantics (even challenging elements like barns and brick walls) while ensuring superior 3D structural accuracy.

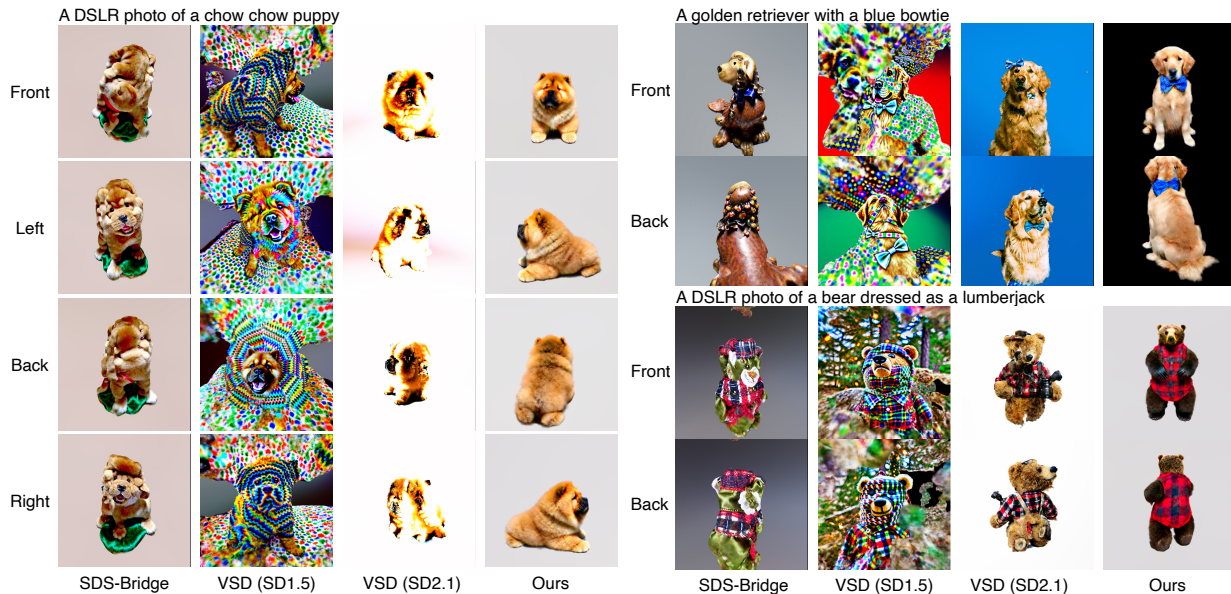


Figure 7: Multi-view Comparison. Ours achieves superior quality and multi-view consistency without common 3D consistency techniques (*e.g.* Perp-Neg (Armandpour et al. 2023)). By contrast, VSD, despite conditioning on view direction, still exhibits severe Janus artifacts, while SDS-Bridge fails to preserve consistent object semantics.

continue to exhibit critical issues such as semantic over-smoothing and multi-view inconsistency, due to inaccurate modeling of the evolving source. To address this, we introduce AnchorDS, which anchors the dynamic source distri-

bution by casting the problem into a dual-conditioned latent space and conditioning on rendering images. Experimental results confirm our method effectively mitigates these issues and achieves superior fidelity and visual quality.

## Acknowledgements

This research / project is supported by the Ministry of Education, Singapore, under the Academic Research Fund Tier 1 (FY2022).

## References

- Armandpour, M.; Sadeghian, A.; Zheng, H.; Sadeghian, A.; and Zhou, M. 2023. Re-imagine the negative prompt algorithm: Transform 2d diffusion into 3d, alleviate janus problem and beyond. *arXiv preprint arXiv:2304.04968*.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instructpix2pix: Learning to follow image editing instructions. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 18392–18402.
- Collins, J.; Goel, S.; Deng, K.; Luthra, A.; Xu, L.; Gundogdu, E.; Zhang, X.; Vicente, T. F. Y.; Dideriksen, T.; Arora, H.; Guillaumin, M.; and Malik, J. 2022. ABO: Dataset and Benchmarks for Real-World 3D Object Understanding. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 21094–21104. IEEE.
- De Bortoli, V.; Thornton, J.; Heng, J.; and Doucet, A. 2021. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in neural information processing systems*, 34: 17695–17709.
- Deitke, M.; Schwenk, D.; Salvador, J.; Weihs, L.; Michel, O.; VanderBilt, E.; Schmidt, L.; Ehsani, K.; Kembhavi, A.; and Farhadi, A. 2022. Objaverse: A Universe of Annotated 3D Objects. *arXiv preprint arXiv:2212.08051*.
- Dong, W.; Yang, B.; Ma, L.; Liu, X.; Cui, L.; Bao, H.; Ma, Y.; and Cui, Z. 2024. Coin3d: Controllable and interactive 3d assets generation with proxy-guided conditioning. In *ACM SIGGRAPH Conf. Proc.*, 1–10.
- Fu, H.; Jia, R.; Gao, L.; Gong, M.; Zhao, B.; Maybank, S.; and Tao, D. 2021. 3D-FUTURE: 3D Furniture Shape with TextURE. *Int. J. Comput. Vis.*, 129(12): 3313–3337.
- He, Y.; Bai, Y.; Lin, M.; Zhao, W.; Hu, Y.; Sheng, J.; Yi, R.; Li, J.; and Liu, Y.-J. 2023. T3Bench: Benchmarking Current Progress in Text-to-3D Generation. *arXiv preprint arXiv:2310.02977*.
- Hertz, A.; Aberman, K.; and Cohen-Or, D. 2023. Delta Denoising Score. In *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2328–2337.
- Ho, J.; and Salimans, T. 2022. Classifier-Free Diffusion Guidance. *arXiv preprint arXiv:2207.12598*.
- Huang, T.; Zeng, Y.; Zhang, Z.; Xu, W.; Xu, H.; Xu, S.; Lau, R. W.; and Zuo, W. 2024. Dreamcontrol: Control-based text-to-3d generation with 3d self-prior. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 5364–5373.
- Jun, H.; and Nichol, A. 2023. Shap-e: Generating conditional 3d implicit functions. *arXiv preprint arXiv:2305.02463*.
- Kadosh, E.; Goren, N.; Patashnik, O.; Garibi, D.; and Cohen-Or, D. 2025. Tight Inversion: Image-Conditioned Inversion for Real Image Editing. *arXiv preprint arXiv:2502.20376*.
- Katzir, O.; Patashnik, O.; Cohen-Or, D.; and Lischinski, D. 2024. Noise-free Score Distillation. In *Proc. Int. Conf. Learn. Represent.*
- Lee, K.; Sohn, K.; and Shin, J. 2024. DreamFlow: High-quality text-to-3D generation by Approximating Probability Flow. In *Proc. Int. Conf. Learn. Represent.*
- Li, W.; Chen, R.; Chen, X.; and Tan, P. 2024. Sweet-Dreamer: Aligning Geometric Priors in 2D diffusion for Consistent Text-to-3D. In *Proc. Int. Conf. Learn. Represent.*
- Li, Z.; Chen, Y.; Zhao, L.; and Liu, P. 2025. Controllable Text-to-3D Generation via Surface-Aligned Gaussian Splatting. In *Proc. Int. Conf. 3D Vis.*
- Liang, Y.; Yang, X.; Lin, J.; Li, H.; Xu, X.; and Chen, Y. 2024. Luciddreamer: Towards high-fidelity text-to-3d generation via interval score matching. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 6517–6526.
- Lin, C.-H.; Gao, J.; Tang, L.; Takikawa, T.; Zeng, X.; Huang, X.; Kreis, K.; Fidler, S.; Liu, M.-Y.; and Lin, T.-Y. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 300–309.
- Lukoianov, A.; S’aez de Oc’ariz Borde, H.; Greenewald, K.; Guizilini, V.; Bagautdinov, T.; Sitzmann, V.; and Solomon, J. M. 2024. Score distillation via reparametrized ddim. *Adv. Neural Inf. Process. Syst.*, 37: 26011–26044.
- McAllister, D.; Ge, S.; Huang, J.-B.; Jacobs, D. W.; Efros, A. A.; Holynski, A.; and Kanazawa, A. 2024. Rethinking Score Distillation as a Bridge Between Image Distributions. In *Proc. Int. Conf. Learn. Represent.*
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2022. SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations. In *Proc. Int. Conf. Learn. Represent.*
- Nichol, A.; Jun, H.; Dhariwal, P.; Mishkin, P.; and Chen, M. 2022. Point-E: A System for Generating 3D Point Clouds from Complex Prompts. *ArXiv*, abs/2212.08751.
- Poole, B.; Jain, A.; Barron, J. T.; and Mildenhall, B. 2023. DreamFusion: Text-to-3D using 2D Diffusion. In *Proc. Int. Conf. Learn. Represent.*
- Qiu, L.; Chen, G.; Gu, X.; Zuo, Q.; Xu, M.; Wu, Y.; Yuan, W.; Dong, Z.; Bo, L.; and Han, X. 2024. Richdreamer: A generalizable normal-depth diffusion model for detail richness in text-to-3d. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 9914–9925.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proc. Int. Conf. Mach. Learn.*, 8748–8763. PMLR.
- Shi, Y.; Wang, P.; Ye, J.; Long, M.; Li, K.; and Yang, X. 2023. Mvdream: Multi-view diffusion for 3d generation. *arXiv preprint arXiv:2308.16512*.
- Siddiqui, Y.; Monnier, T.; Kokkinos, F.; Kariya, M.; Kleiman, Y.; Garreau, E.; Gafni, O.; Neverova, N.; Vedaldi, A.; Shapovalov, R.; and Novotny, D. 2024. Meta 3D Asset-Gen: Text-to-Mesh Generation with High-Quality Geometry, Texture, and PBR Materials. In Globerson, A.; Mackey,

L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Adv. Neural Inf. Process. Syst.*, volume 37, 9532–9564. Curran Associates, Inc.

Su, X.; Song, J.; Meng, C.; and Ermon, S. 2023. Dual Diffusion Implicit Bridges for Image-to-Image Translation. In *Proc. Int. Conf. Learn. Represent.*

Tang, B.; Wang, J.; Wu, Z.; and Zhang, L. 2024a. Stable Score Distillation for High-Quality 3D Generation. *arXiv preprint arXiv:2312.09305*, (arXiv:2312.09305).

Tang, J.; Ren, J.; Zhou, H.; Liu, Z.; and Zeng, G. 2024b. DreamGaussian: Generative Gaussian Splatting for Efficient 3D Content Creation. In *Proc. Int. Conf. Learn. Represent.*

Wang, Z.; Lu, C.; Wang, Y.; Bao, F.; Li, C.; Su, H.; and Zhu, J. 2023. ProlificDreamer: High-Fidelity and Diverse Text-to-3D Generation with Variational Score Distillation. In Oh, A.; Neumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Adv. Neural Inf. Process. Syst.*, volume 36, 8406–8441. Curran Associates, Inc.

Xiang, J.; Lv, Z.; Xu, S.; Deng, Y.; Wang, R.; Zhang, B.; Chen, D.; Tong, X.; and Yang, J. 2024. Structured 3D Latents for Scalable and Versatile 3D Generation. *arXiv preprint arXiv:2412.01506*.

Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.

Yi, T.; Fang, J.; Wang, J.; Wu, G.; Xie, L.; Zhang, X.; Liu, W.; Tian, Q.; and Wang, X. 2024a. Gaussiandreamer: Fast generation from text to 3d gaussians by bridging 2d and 3d diffusion models. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 6796–6807.

Yi, T.; Fang, J.; Zhou, Z.; Wang, J.; Wu, G.; Xie, L.; Zhang, X.; Liu, W.; Wang, X.; and Tian, Q. 2024b. GaussianDreamerPro: Text to Manipulable 3D Gaussians with Highly Enhanced Quality. *arXiv:2406.18462*.

Yu, X.; Guo, Y.-C.; Li, Y.; Liang, D.; Zhang, S.-H.; and Qi, X. 2023. Text-to-3D with Classifier Score Distillation. *arXiv.org*.

Zeng, B.-W.; Li, S.; Feng, Y.; Yang, L.; Li, H.; Gao, S.; Liu, J.; He, C.; Zhang, W.; Liu, J.; Zhang, B.; and Yan, S. 2023. IPDreamer: Appearance-Controllable 3D Object Generation with Complex Image Prompts. *arXiv preprint arXiv:2310.05375*.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *IEEE International Conference on Computer Vision (ICCV)*.