

Optimized View and Geometry Distillation from Multi-view Diffuser

Youjia Zhang¹ Zikai Song¹ Junqing Yu¹ Yawei Luo² Wei Yang^{†1}

¹ Huazhong University of Science and Technology

² Zhejiang University

<https://youjiazhang.github.io/USD/>

Abstract. Generating multi-view images from a single input view using image-conditioned diffusion models is a recent advancement and has shown considerable potential. However, issues such as the lack of consistency in synthesized views and over-smoothing in extracted geometry persist. Previous methods integrate multi-view consistency modules or impose additional supervisory to enhance view consistency while compromising on the flexibility of camera positioning and limiting the versatility of view synthesis. In this study, we consider the radiance field optimized during geometry extraction as a more rigid consistency prior, compared to volume and ray aggregation used in previous works. We further identify and rectify a critical bias in the traditional radiance field optimization process through score distillation from a multi-view diffuser. We introduce an **Unbiased Score Distillation (USD)** that utilizes unconditioned noises from a 2D diffusion model, greatly refining the radiance field fidelity. We leverage the rendered views from the optimized radiance field as the basis and develop a two-step specialization process of a 2D diffusion model, which is adept at conducting object-specific denoising and generating high-quality multi-view images. Finally, we recover faithful geometry and texture directly from the refined multi-view images. Empirical evaluations demonstrate that our optimized geometry and view distillation technique generates comparable results to the state-of-the-art models trained on extensive datasets, all while maintaining freedom in camera positioning.

1 Introduction

Traditionally, the process of generating a three-dimensional model from a singular image necessitates extensive and meticulous efforts by highly skilled artists. However, recent advancements in neural networks, particularly through the adaptation of 2D diffusion models for 3D synthesis, have rendered the conversion of a single image into a 3D object feasible. The early breakthrough comes from the text to 3D domain, where DreamFusion [30] and Score Jacobian Chaining (SJC) [46] proposes a Score Distilling Sampling (SDS) strategy to distill the scores learned by 2D diffusion models from large-scale images to optimize a Neural Radiance Field (NeRF) [27], circumventing the need for 3D data. Successive approaches further improve the quality and diversity of generated geometries from textural prompts [2, 19, 48]. Particularly, RealFusion [26] migrates the



Fig. 1: In this study, our proposed methodology advances the extraction of geometries and views from a multi-view diffusion model, specifically Zero-1-to-3 [21]. This approach effectively addresses and mitigates the significant bias observed in previous geometry extraction techniques, concurrently enhancing the quality of generated views. Our technique produces multi-view images and geometries that are comparable, sometimes superior particularly for irregular camera poses, when benchmarked against the SOTA methodologies such as SyncDreamer [23] and Wonder3D [24], without training on large-scale data.

scheme to generate plausible 3D reconstruction matches to a single input image via textual inversion adapted supervision.

More relevantly, 3DiM [50] and MVDream [38] develop a pose-conditional image-to-image diffusion model, which generates the novel view at a target pose from a source view. Zero-1-to-3 [21] adopts a similar framework and learns control of viewpoints through a synthetic dataset and demonstrates zero-shot generalization to in-the-wild images. Though Zero-1-to-3 demonstrates plausible novel views, they are not multi-view consistent and the geometry distilled from SDS tends to be oversmoothed. To enhance the multi-view consistency, SyncDreamer [23] devises a volume-encoded multi-view noise predictor to share information across different views. Wonder3D [24] predicts the multi-view color images along with their normal maps from a cross-domain diffusion model. Though enhancing the multi-view consistency of image generation, these methods com-




Fig. 2: The unconditional noise predicted by Zero-1-to-3 [21] model tends to be biased. As a demonstration, we use the ‘Mario’ image as a toy example and add various levels of noise to the image (larger \mathbf{T} means more noise has been added). We use the predicted unconditional noise to recover the original image from noisy input and find the results of Zero-1-to-3 deviate from the input image greatly even for very small amount of noise. The right sub-figure shows the averaged difference between the predicted noise and the added noise.

promise the flexibility of camera positioning and only allow synthesis for a limited number of views.

In this study, we observe that the predicted unconditional noise from the multi-view diffuser, i.e., the Zero-1-to-3 model, appears to be biased. That is, even if we only add very low-level noise to a normal image and use the unconditional noise predicted by a Zero-1-to-3 model for denoising, the result still tends to deviate greatly from the original image. We analyze and rectify the critical bias by using an unconditioned noise from a pre-trained 2D diffusion model and greatly refining the geometry fidelity. Moreover, previous approaches use either 3D volume [23] or ray aggregation [45] to share information across views. We consider the radiance field as the consistency prior, and encourage the generated multi-view images to be consistent with the NeRF renderings. We develop a two-step specialization process of a 2D diffusion model, which is adept at conducting target-specific denoising and generating high-quality multi-view images from NeRF renderings. We then further use the refined views to generate the geometry and texture from NeuS [47], and in the meanwhile enforce input view consistency using view score distillation. Our approach generates comparable-quality of multi-view images and geometry to the SOTA approaches, including SyncDreamer and Wonder3D, without enforcing any restriction on camera poses. Consequently, we posit that our approach offers superior adaptability and effectiveness in addressing the challenges associated with generating consistent and high-quality multi-view imagery and geometry, affirming its substantial potential for widespread application in relevant fields.

2 Related Work

2D Diffusion Models for 3D Generation. Recent advancements in 2D diffusion models [4, 13, 34] and large-scale visual language models, notably the CLIP

model [33], have catalyzed new approaches for generating 3D assets. Pioneering efforts such as DreamFusion [30] and SJC [46] have developed methods for transforming 2D text into images, subsequently facilitating the generation of 3D shapes from text. This approach has inspired a range of subsequent studies that adopt a shape-by-shape optimization scheme. Additionally, the integration of 2D diffusion models with robust vision language models, especially CLIP [33], has emerged as a significant exploration in the generation of 3D assets [54]. The typical methodology involves optimizing a 3D representation, such as NeRF, mesh, or SDF, and then utilizing neural rendering to generate 2D images from various viewpoints. These images are processed through 2D diffusion models or the CLIP model to calculate SDS losses, which guide the optimization of the 3D shape. Building on the foundations laid by DreamFusion and SJC, numerous works have enhanced text-to-3D distillation methods in various aspects. Notably, Magic3D [19] develops a two-stage coarse-to-fine optimization framework for high-resolution 3D content generation, and ProlificDreamer [48] proposes a Variational Score Distillation (VSD) for generating highly detailed geometry. However, challenges such as low efficiency and the multi-face Janus problem, where optimized geometry tends to produce multiple faces due to the lack of explicit 3D supervision, remain prevalent. Furthermore, some works [32] have applied this distillation pipeline in single-view reconstruction tasks. While these methods have achieved impressive results, they often require extensive time for textual inversion and NeRF optimization, without always guaranteeing satisfactory outcomes. In contrast with the 2D diffusion to 3D extension, which is ignorant to multi-view consistency, our method focuses on the multi-view diffuser technique, which predicts the noises for novel views, and inherently avoids the Janus problem.

Multi-view Diffusion Models. In light of the complexities involved in ensuring the integrity of generated 3D content, recent efforts [1, 6, 9, 18, 22, 40, 42, 43, 45, 50, 52, 56, 57, 60] have explored the feasibility of directly generating novel views from a single image input. Notably, the Scene Representation Transformer [37] extends the vision transformer to image sets, enabling global information integration for 3D reasoning. Similarly, 3DiM [50] develops a pose-conditional image-to-image diffusion model, translating a single input view into consistent and sharp completions across multiple views. A seminal work in this area, Zero-1-to-3 [21], utilizes a similar network structure, trained on a large-scale synthetic 3D dataset, demonstrating notable generalizability. Recent advancements, such as One-2-3-45 [53], have leveraged the generalizable neural reconstruction method SparseNeuS [25] to directly produce 3D geometry from images generated by Zero-1-to-3. While this approach is more efficient and alleviates the Janus (multi-head) problem, it tends to produce lower-quality results with reduced geometric detail. In a different vein, SyncDreamer [23] concentrates on object reconstruction, generating images in a single reverse process and utilizing attention to synchronize states among views. This contrasts with Viewset Diffusion [39], which requires predicting a radiance field. SyncDreamer solely relies on attention for synchronization, fixing the viewpoints of generated views

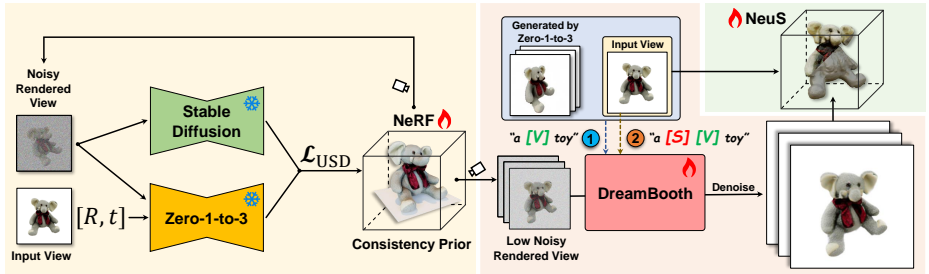


Fig. 3: The overall pipeline of our approach. We first use our Unbiased Score Distillation to extract an optimized underlying radiance field. And then we use the NeRF [27] as our consistency prior, i.e., the generated views should be consistent with the NeRF renderings. We propose a two-stage specialization scheme to obtain a specified DreamBooth [35] specifically for the target. We then denoise the NeRF renderings to obtain high-quality views and subsequently use NeuS [10, 47] technique to recover the geometry. Our optimized scheme generates comparable, sometimes better particularly for irregular camera poses, results to the SOTA works without training on large-scale data.

to enhance training convergence. A significant trend in recent research has been the design of various functional attention layers. Consistent123 [51] employs a shared self-attention layer, where all views query the same key and value from the input view. ConsistNet [55] introduces two sub-modules: a view aggregation module and a ray aggregation module, to extract features consistent across multiple views. MVDream [38] utilizes 3D self-attention. Wonder3D [24] goes a step further by not only outputting multi-view images but also inputting the normal map for each perspective, using cross-domain attention to maintain the consistency of the 3D structure. Both SyncDreamer and Wonder3D, resulting from fixed view outputs, exhibit sensitivity to the camera angle of the input image. In contrast, our method can reconstruct accurate 3D structures from inputs captured at varying camera angles. While existing approaches have focused on developing novel modules for enforcing multi-view consistency, our work demonstrates that an optimized distillation strategy can yield views and geometries comparable to models trained on large-scale datasets. This insight may inspire further exploration in improved strategies for geometry and view extraction.

3 Method

Our objective is to synthesize consistent multi-view images and high-quality geometric representations from a single input image. Notable prior work, such as the Zero-1-to-3, has demonstrated impressive results by utilizing an image and camera pose-conditioned diffusion model. However, this approach encounters limitations, particularly in terms of generating inconsistent multi-view images and a tendency for over-smoothing in the geometric output. The state-of-the-art research, including SyncDreamer [23] and Wonder3D [24], addresses these challenges by incorporating additional modules for consistency or employing normal supervision. However, this often compromises the flexibility of positioning




Fig. 4: The qualitative comparisons with baseline models on multi-view color images. Our approach generates consistent multi-view images while preserving the image details.

the target camera at will. Contrastingly, our work adopts a distinct methodology, showcasing that comparable quality in views and geometry can be achieved through a meticulously crafted distillation strategy. Central to our approach is the insight that the unconditional noise predictions from Zero-1-to-3 are inherently biased. We propose the utilization of unconditional noise from the Stable Diffusion [41] model to rectify this issue, as elaborated in Sec. 3.1. Our method, termed Unbiased Score Distillation (USD), significantly enhances the quality of the radiance field relative to previously used SDS/SJC methods. Furthermore, we employ the optimized NeRF [27] as a consistency prior, in contrast to previous implicit constraints such as 3D volume or ray aggregation. We ensure that the generated views and geometry align coherently with the distilled NeRF to achieve consistency. We posit that specializing a diffusion model to denoise the target object is crucial. To this end, we implement the DreamBooth [35] technique and engage in a two-stage fine-tuning process, detailed in Sec. 3.2. The first stage involves using multi-view images from Zero-1-to-3 as positive samples, contrasting them with text-prompt generated images as negative samples for image style learning. In the subsequent stage, the input image serves as the positive sample, with all Zero-1-to-3 generated images as negatives, focusing on learning finer details. Subsequently, we introduce a low level of noise into the NeRF renderings and employ the fine-tuned diffusion model for denoising. The final step involves applying the NeuS technique for mesh reconstruction, while in the meanwhile using View Score Distillation to ensure input view consistency. Fig. 3 illustrates the complete pipeline of our methodology.

3.1 Unbiased Score Distillation

In this section, we highlight the significant bias in Zero-1-to-3’s unconditional noise due to insufficient unconditional training and object-level dataset bias, affecting geometry quality in SDS-based distillation, both theoretically and empirically. We then propose a rectification method.

Bias in Unconditional Noise. As a multi-view diffuser, Zero-1-to-3 predicts noises of a target image latent \mathbf{z} given two conditions: the input image c_I and

relative camera pose c_P , i.e., it learns the probability distribution $P(\mathbf{z}|c_I, c_P)$ for image latent \mathbf{z} . We apply *Bayes' Rule* to decompose the conditional probability:

$$P(\mathbf{z}|c_I, c_P) = \frac{P(\mathbf{z}, c_I, c_P)}{P(c_I, c_P)} = \frac{P(c_I|c_P, \mathbf{z})P(c_P|\mathbf{z})P(\mathbf{z})}{P(c_I, c_P)} \quad (1)$$

Diffusion models estimate the score [15] of the data distribution, i.e., the derivative of the log probability, giving us the following expression:

$$\begin{aligned} \nabla_{\mathbf{z}} \log(P(\mathbf{z}|c_I, c_P)) &= \nabla_{\mathbf{z}} \log(P(c_I|c_P, \mathbf{z})) + \\ &\quad \nabla_{\mathbf{z}} \log(P(c_P|\mathbf{z})) + \nabla_{\mathbf{z}} \log(P(\mathbf{z})) \end{aligned} \quad (2)$$

This leads to the score estimate with classifier-free guidance (CFG [14]):

$$\begin{aligned} \epsilon_{\phi}^{\text{CFG}}(\mathbf{z}_t, c_I, c_P) &= \alpha_1 [\epsilon_{\phi}(\mathbf{z}_t, c_I, c_P) - \epsilon_{\phi}(\mathbf{z}_t, c_P, \emptyset)] + \\ &\quad \alpha_2 [\epsilon_{\phi}(\mathbf{z}_t, c_P, \emptyset) - \epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset)] + \\ &\quad \epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset) \end{aligned} \quad (3)$$

where ϵ_{ϕ} is a neural network for predicting noises given conditions, α_1 and α_2 are guidance scales that enable separately trading off the strength of conditions c_I and c_P separately. To ensure accurate prediction of $\epsilon_{\phi}(\mathbf{z}_t, c_P, \emptyset)$ and $\epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset)$, one needs to randomly drop c_I and c_P during training ϵ_{ϕ} . However, we observed that Zero-1-to-3 did not truly follow this procedure to supervise unconditional noise. As shown in Fig. 5, Zero-1-to-3 only randomly dropped image conditions c_I while keeping c_P untouched. During inference, it replaces the tensor after fully connected (FC) layer $f(\cdot)$ with zeros (not exactly equal to setting both c_I and c_P to zeros as $f(0) \neq 0$), leads to a bias. Moreover, the dataset used for fine-tuning is majorly object-centric, which may also introduce additional domain bias. Fully addressing this bias problem requires re-training the multi-view diffuser on a broader and more balanced dataset, which requires tremendous effort. Here we propose an effortless fix in the following.

Rectification. To alleviate the inaccurate unconditional noise problem, we first set $\alpha_1 = \alpha_2 = \omega$ to eliminate $\epsilon_{\phi}(\mathbf{z}_t, c_P, \emptyset)$. We then and yield a special case of $\epsilon_{\phi}^{\text{CFG}}(\mathbf{z}_t, c_I, c_P)$ as:

$$\epsilon_{\phi}^{\text{CFG}}(\mathbf{z}_t, c_I, c_P) = \omega [\epsilon_{\phi}(\mathbf{z}_t, c_I, c_P) - \epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset)] + \epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset) \quad (4)$$

To demonstrate our setting of α_1 and α_2 is valid, We provide an empirical validation as in Fig. 6. Further, we consider $\epsilon_{\phi}(\mathbf{z}_t, \emptyset, \emptyset)$ predicts noises from only the noisy latent and is equivalent to the unconditional noise $\epsilon_{\psi}(\mathbf{z}_t, \emptyset)$ of Stable Diffusion (SD) as it uses the same variational autoencoders (VAE [17]) as Zero-1-to-3. As such, the bias has been well rectified. To verify the effect of our rectification, Table 1 shows the denoising effect using various unconditional noise settings: (1) use SD unconditional noise $\epsilon_{\psi}(\mathbf{z}_t, \emptyset)$. (2) use Zero-1-to-3 noise $\epsilon_{\phi}(\mathbf{z}_t, f(0, c_P))$ with $f(0, c_P)$ as a condition (same to the training process as in Fig 5). (3) use Zero-1-to-3 noise $\epsilon_{\phi}(\mathbf{z}_t, 0)$ with 0 as a condition (referring to the




Fig. 5: Training and inference process of Zero-1-to-3 [21]. Training for predicting unconditional noise involves setting the c_I conditions to 0 at regular intervals.

Table 1: Quantitative evaluation of different unconditional noises. We randomly samples 1,000 different images from the GSO dataset [7], RTMV [44], and Objaverse [5], add a certain level of noise w.r.t. \mathbf{T} , where larger \mathbf{T} means more noise is added, and compare the denoising effect under different conditions of Stable Diffusion and Zero-1-to-3. The top three for each metric are highlighted in red, orange, and yellow respectively.

Noise level	$\mathbf{T}=50$			$\mathbf{T}=100$			$\mathbf{T}=200$			$\mathbf{T}=300$		
Metrics	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
$\epsilon_\psi(\mathbf{z}_t, \emptyset)$	40.26	0.992	0.008	36.94	0.988	0.014	33.99	0.981	0.023	31.98	0.974	0.027
$\epsilon_\phi(\mathbf{z}_t, f(0, c_P))$	36.17	0.981	0.015	34.85	0.979	0.024	33.07	0.975	0.031	30.44	0.969	0.046
$\epsilon_\phi(\mathbf{z}_t, 0)$	36.13	0.979	0.015	34.70	0.970	0.026	32.76	0.973	0.033	29.94	0.963	0.055

inference process, as shown in Fig 5). The SD noise generates the best result, and setting (2) is slightly better than (3) as it follows the training set up.

Unbiased Score Distillation. One major application of a multi-view diffuser is to distill 3D content, represented as NeRF parameter θ , using the SDS loss for optimization:

$$\nabla_{\theta} \mathcal{L}_{\text{SDS}} = \mathbb{E}_{t, c_P, \epsilon} \left[w(t) \left(\epsilon_{\phi}^{\text{CFG}}(\mathbf{z}_t, c_I, c_P) - \epsilon \right) \frac{\partial \mathbf{z}_t}{\partial \theta} \right] \quad (5)$$

where $w(t)$ is a weighting function, ϵ is standard Gaussian noise, and \mathbf{z}_t refers to the noisy latent as $\mathbf{z}_t = \sqrt{\alpha_t} \mathbf{z} + \sqrt{1 - \alpha_t} \epsilon$, with α_t being the noise scheduler. We rewrite the noise difference in Formula 5 by adding an additional weighting factor λ to $[\epsilon_\psi(\mathbf{z}_t, \emptyset) - \epsilon]$:

$$\epsilon_{\phi}^{\text{CFG}}(\mathbf{z}_t, c_I, c_P) - \epsilon = \omega [\epsilon_{\phi}(\mathbf{z}_t, c_I, c_P) - \epsilon_{\psi}(\mathbf{z}_t, \emptyset)] + \lambda [\epsilon_{\psi}(\mathbf{z}_t, \emptyset) - \epsilon] \quad (6)$$

where setting $\lambda = 1$ we get Formula 5. Inspired by DDS [12] and CSD [58], we found that setting $\lambda = 0$ can significantly improve the details of the 3D details generated using SDS, as shown in Fig. 11. We obtain our Unbiased Score Distillation (USD) as:

$$\nabla_{\theta} \mathcal{L}_{\text{USD}} = \mathbb{E}_{t, c_P, \epsilon} \left[w(t) [\omega (\epsilon_{\phi}(\mathbf{z}_t, c_I, c_P) - \epsilon_{\psi}(\mathbf{z}_t, \emptyset))] \frac{\partial \mathbf{z}_t}{\partial \theta} \right] \quad (7)$$

Our USD generates much better and consistent 3D than the SDS/SJC method used in Zero-1-to-3, as shown in Fig. 7.




Fig. 6: We compare the effect of the relationship between α_1 and α_2 on SDS [30], which can achieve the best when $\alpha_1 = \alpha_2$.




Fig. 7: We compare our geometry and texture distillation method based USD with Magic123 [31] and Zero-1-to-3 [21] (using SJC). Though the results of Magic123 tend to be smooth and clear, it heavily relies on the estimated depth from the input view, which may lead to incorrect estimations when the depth estimator fails. Zero-1-to-3 uses SJC for geometry and texture reconstruction, its formation is biased and hence lead to bad estimations. Ours, however, generate consistent and faithful estimations.

3.2 Consistent View and Geometry Distillation

Though with USD, we can obtain optimized NeRF for view synthesis, the generated images tend to be blurry, and direct extract geometry from the density field of NeRF produces noises on the mesh. Since our ultimate goal is to extract high-quality geometry and consistent multi-views, we propose to utilize the generated NeRF as the view consistency prior, i.e., the final high-quality view images should be consistent with the NeRF rendering. Thus, our problem transformed into a denoising problem.

Two-stage Specified Diffusion. We leverage the recent advance in diffusion model specialization, i.e., the DreamBooth [35], to fine-tune a 2D diffusion model for the specific target in the input image. We observe that the novel view images generated by Zero-1-to-3, though not multi-view consistent, tend to have the same style as the input image. We design a two-stage tuning method to gradually let the diffusion model learn the object details. In the first stage, we use the multi-view images generated by Zero-1-to-3 as positive samples, contrasting them with text-prompt generated images as negative samples for learning the visual style of the target. In the subsequent stage, the input image serves as the positive sample, with all Zero-1-to-3 generated images as negatives, focusing on learning finer details. During optimization, we use a unique identifier **[V]** to capture the visual style of the target. In the second stage, we set only the input image as the positive sample and the images generated by Zero-1-to-3 as negative ones, and use the additional identifier **[S]** to capture the identity of the specific target.

Geometry and Texture Distillation with Input View Supervision. With the specialization diffusion model, we add a small noise, Stable Diffusion sched-



Fig. 8: SOTA methods, like SyncDreamer [23] and Wonder3D [24] impose limitations on the viewing angles of the input image. For instance, these methods proficiently manage near-horizontal viewing angles but often falter with upward views. In contrast, our approach consistently generates reasonable structures across a broader range of input view poses.

uler $t = 200$, to the NeRF render images and conduct the denoising process. Then, we use the NeuS technique to reconstruct the geometry from the high-quality and clear images.

We observe that the input view is rarely used for optimization of the geometry and texture in multi-view diffusers, despite the fact that the input image is most faithful to the target object. To exploit the input view information, we further develop a reference view score distillation during the NeuS reconstruction process. Specifically, we consider the image render from NeuS at the input image viewpoint as $\mathbf{z}_t^* = \mathcal{R}(\Theta, p^*)$, where \mathcal{R} is the rendering function from NeuS model defined by Θ , p^* is the camera pose of the input image, can be set to a particular relative pose. We define the following reference view distillation loss as:

$$\mathcal{L}_{RV} = \mathbb{E}_{t, \epsilon} [w(t) \|\epsilon_\psi(\mathbf{z}_t^*, \varnothing) - \epsilon_\psi(\mathbf{y}_t, \varnothing)\|_2^2] \quad (8)$$

where \mathbf{y}_t is input view image. We add this reference view distillation loss \mathcal{L}_{RV} to original photometric loss with optimized images for optimizing the NeuS parameter Θ . This scheme is a better supervision strategy than directly applying MSE loss on the input view.

4 Experiments

We conduct extensive experiments, both qualitatively and quantitatively, to demonstrate the effectiveness of our method.

4.1 Implementation details

Optimizing the pipeline. We use exactly the same set of hyperparameters for all experiments and do not perform any per-object hyperparameter optimization. We use the implicit-volume implementation in threestudio [11] as our 3D representation (NeRF), which includes a multi-resolution hash-grid and a MLP




Fig. 9: Qualitative comparisons of our method with baseline approaches, namely Wonder3D [24], SyncDreamer [23], and One-2-3-45 [8], on the GSO dataset [7], focusing on the quality of the reconstructed textured meshes.

to predict density and RGB. The NeRF is optimized for 10,000 steps with an Adam optimizer at a learning rate of 0.01, weight decay of 0.05, and betas of (0.9, 0.95). For USD, the maximum and minimum time steps are decreased from 0.98 to 0.5 and 0.02, respectively, over the first 5,000 steps. We adopt the Stable Diffusion [41] model of V2.1. The classifier-free guidance (CFG) scale of the USD is set to 7.5 following [48]. The DreamBooth backbone is implemented by Stable Diffusion V2.1. In the first stage, we use Stable Diffusion to generate 200 images as negative samples for training. The USD process takes about 1.5 hours on a NVIDIA Tesla V100 (32GB) GPU. For DreamBooth fine-tuning, we train the model around for 600 steps with a learning-rate as $2e-6$, weight decay as 0.01 and a batch size of 2.

Camera setting. Since the reference image is unposed, we assume its camera parameters [8] are as follows. We set the field of view (FOV) of the camera is 40° , and the radial distance is 1.5 meters. Note this camera setting works for images subject to the front-view assumption. For images taken deviating from the front view, a manual change of polar angle or a camera estimation is required.

4.2 Evaluation Protocol

Evaluation Datasets. Following prior research [21, 23, 24], we adopt the Google Scanned Object dataset [7] for our evaluation, which includes a wide variety of common everyday objects. Our evaluation dataset matches that of SyncDreamer [23], consisting of 30 objects that span from everyday items to animals. For each object in the evaluation set, we render an image with a size of 256×256 and use it as the input. Additionally, to assess the generalization ability of our model, we include images with diverse styles collected from the website in Zero-1-to-3, SyncDreamer and Wonder3D.

Metrics. To evaluate the quality of single view reconstruction, we used two commonly used metrics: the chamfer distance (CD) between the ground truth shape and the reconstructed shape, and the volume IoU. Due to different methods using different normative systems, before calculating these two metrics, we

Table 2: Quantitative comparison with baseline methods. We report Chamfer Distance and Volume IoU on the GSO [7] dataset. The original implementation of SyncDreamer uses vanilla NeuS [47] for extracting 3D meshes, while Wonder3D uses Instant NGP(iNGP) adapted NeuS [10]. We report results using both techniques for better demonstration.

Method	Chamfer Dist.↓	Volume IoU↑
Realfusion [26]	0.0819	0.2741
Magic123 [31]	0.0516	0.4528
One-2-3-45 [20]	0.0629	0.4086
Point-E [29]	0.0426	0.2875
Shap-E [16]	0.0436	0.3584
Zero-1-to-3 [21]	0.0339	0.5035
SyncDreamer [23] (NeuS)	0.0261	0.5421
Wonder3D [24] (iNGP+NeuS)	0.0199	0.6244
Ours (NeuS)	0.0240	0.5688
Ours (iNGP+NeuS)	0.0177	0.6330

Table 3: The quantitative comparison in novel view synthesis. We report PSNR, SSIM [49], LPIPS [59] on the GSO [7] dataset.

Method	PSNR↑	SSIM↑	LPIPS↓
Realfusion [26]	15.26	0.722	0.283
Zero-1-to-3 [21]	18.93	0.779	0.166
SyncDreamer [23]	20.05	0.798	0.146
Wonder3D [24]	26.07	0.924	0.065
Ours	25.38	0.927	0.049

first align the generated shapes with the basic fact shapes. Moreover, we adopt the metrics PSNR, SSIM [49] and LPIPS [59] for evaluating the generated color images.

4.3 Ablation Study

We validate our design choices by ablating 4 major model variants, that are without Unbiased Score Distillation, without reference view distillation loss, without the two-stage DreamBooth(DB). The texture color of the 3D models generated by without USD exhibit bias, and the shapes are not smooth. Not using reference view supervision will result in the inability to recover the same texture details as the input image, and not using DB will result in blurring of texture details.

4.4 Different Viewing Angle Comparisons

We found that SyncDreamer [23] and Wonder3D [24] are very sensitive to viewing angles. If a relatively high viewing angle is input, SyncDreamer and Wonder3D will predict incorrect multi-view images, as shown in Fig. 8.

4.5 Single View Reconstruction

We evaluate the quality of the reconstructed geometry of different methods. The quantitative results are summarized in Table 2, and the qualitative comparisons are presented in Fig. 9. The quality of Wonder3D [24] shape reconstruction depends on the perspective of the input view, such as ‘Train’ shown in Fig. 9, where Wonder3D generated incorrect prediction results. The shape generated by SyncDreamer [23] undergoes deformation due to the camera pose in the input view. One-2-3-45 [20] attempts to reconstruct meshes from the multiview-inconsistent outputs of Zero-1-to-3 [21]. While it can capture coarse geometries, it loses important details in the process. In contrast, our method can achieve good reconstruction quality and texture in terms of geometric structure and texture. In




Fig. 10: Ablation study on the effect of each components, including the USD 3.1, the reference view score distillation 3.2, the DreamBooth 3.2 refinement of our method.

the paper, all the surface extraction demonstrated by our method is built on the Instant NGP(iNGP) [28] based SDF reconstruction method [10], and we use Blender Cycles [3] to render the results.

4.6 Novel View Synthesis

We evaluate the quality of novel view synthesis for different methods. The quantitative results are presented in Table 3, and the qualitative results can be found in Fig. 4. Zero-1-to-3 [21] produces visually reasonable images, but they lack multi-view consistency since it operates on each view independently (we don’t show the results of Zero-1-to-3). Although SyncDreamer [23] introduces a volume attention scheme to enhance the consistency of multi-view images, their model is sensitive to the elevation degrees of the input images and tends to produce unreasonable results. Wonder3D [24] ensures 3D consistency by generating normal maps, but may result in incorrect results for some input viewpoints.

4.7 Text-to-image-to-3D

As a case study, we combine text-to-image models, i.e., the Stable Diffusion or Imagen [36] to generate 3D models from text. We show some examples in Fig. 12. Compared to DreamFusion [30] and ProlificDreamer [48], our method shows no multi-face Janus problem and conforms to the text faithfully.

5 Conclusions

In this work, we introduce an optimized approach for distilling geometry and views from a multi-view diffuser, with a specific focus on the Zero-1-to-3 model. We observed that the direct application of the SDS/SJC technique to Zero-1-to-3 is often suboptimal, primarily due to bias issues inherent in unconditional noise. To address this, we propose an Unbiased Score Distillation (USD) strategy by leveraging unconditioned noises from a 2D diffusion model to effectively enhance the optimized radiance field. Moreover, we developed a two-stage DreamBooth

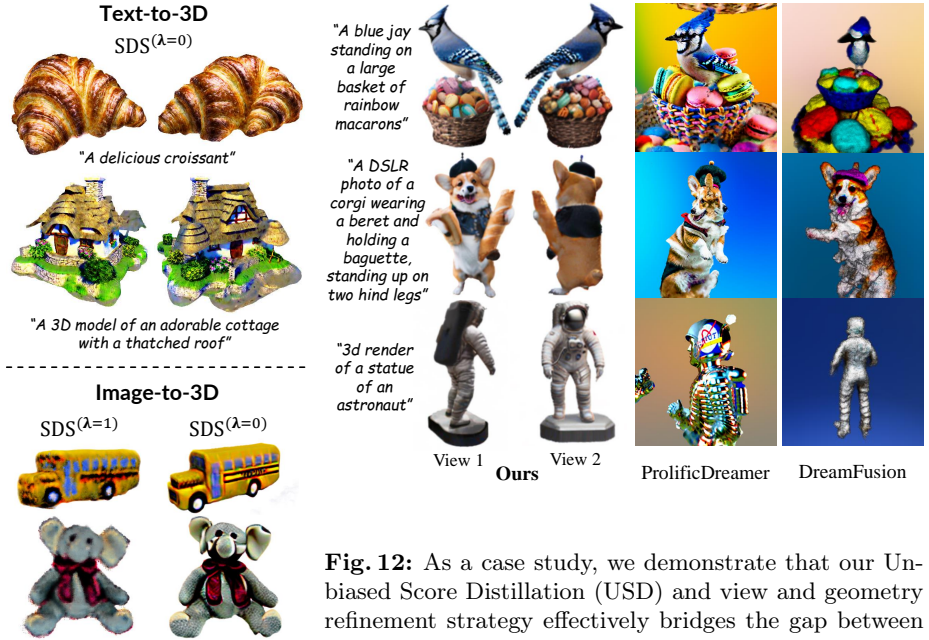


Fig. 11: Experiments on both text-to-3D and image-to-3D tasks demonstrate that SDS can achieve more detailed results when $\lambda = 0$.

Fig. 12: As a case study, we demonstrate that our Un-biased Score Distillation (USD) and view and geometry refinement strategy effectively bridges the gap between multi-view diffuser with text prompts. We use the Stable Diffusion model to generate an image from a text prompt and use our scheme to generate geometry from this image. The result shows a significant advantage of our method in mitigating the multi-face Janus problem (middle row).

refinement process to improve the rendering of views. This process ensures consistency across multiple perspectives while simultaneously enhancing image quality. While we have identified and addressed the bias issue in the Zero-1-to-3 model, the underlying causes remain to be fully understood. Future research will delve into the theoretical aspects of this bias problem. Additionally, we aim to explore the potential applications of USD in other fields, such as image translation and view synthesis.



Fig. 13: More results. Our model is generally applicable to various types of 2D images.

References

1. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aittala, M., De Mello, S., Karras, T., Wetzstein, G.: Generative novel view synthesis with 3d-aware diffusion models. In: ICCV (2023) [4](#)
2. Chen, R., Chen, Y., Jiao, N., Jia, K.: Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. arXiv preprint arXiv:2303.13873 (2023) [1](#)
3. Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), <http://www.blender.org> [13](#)
4. Croitoru, F.A., Hondru, V., Ionescu, R.T., Shah, M.: Diffusion models in vision: A survey. T-PAMI (2023) [3](#)
5. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., Vanderbilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: CVPR (2023) [8](#)
6. Deng, C., Jiang, C., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D., et al.: Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In: CVPR (2023) [4](#)
7. Downs, L., Francis, A., Koenig, N., Kinman, B., Hickman, R., Reymann, K., McHugh, T.B., Vanhoucke, V.: Google scanned objects: A high-quality dataset of 3d scanned household items. In: ICRA (2022) [8](#), [11](#), [12](#)
8. Face, H.: One-2-3-45. <https://huggingface.co/spaces/One-2-3-45/One-2-3-45> (2023) [11](#)
9. Gu, J., Trevithick, A., Lin, K.E., Susskind, J.M., Theobalt, C., Liu, L., Ramamoorthi, R.: Nerfdiff: Single-image view synthesis with nerf-guided distillation from 3d-aware diffusion. In: ICML (2023) [4](#)
10. Guo, Y.C.: Instant neural surface reconstruction (2022), <https://github.com/bennyguo/instant-nsr-pl> [5](#), [12](#), [13](#)
11. Guo, Y.C., Liu, Y.T., Shao, R., Laforte, C., Voleti, V., Luo, G., Chen, C.H., Zou, Z.X., Wang, C., Cao, Y.P., Zhang, S.H.: threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio> (2023) [10](#)
12. Hertz, A., Aberman, K., Cohen-Or, D.: Delta denoising score. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2328–2337 (2023) [8](#)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020) [3](#)
14. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022) [7](#)
15. Hyvärinen, A., Dayan, P.: Estimation of non-normalized statistical models by score matching. Journal of Machine Learning Research **6**(4) (2005) [7](#)
16. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023) [12](#)
17. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) [7](#)
18. Lei, J., Tang, J., Jia, K.: Generative scene synthesis via incremental view inpainting using rgbd diffusion models. In: CVPR (2022) [4](#)
19. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: CVPR (2023) [1](#), [4](#)

20. Liu, M., Xu, C., Jin, H., Chen, L., Xu, Z., Su, H.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. arXiv preprint arXiv:2306.16928 (2023) [12](#)
21. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: ICCV (2023) [2](#), [3](#), [4](#), [8](#), [9](#), [11](#), [12](#), [13](#)
22. Liu, X., Kao, S.h., Chen, J., Tai, Y.W., Tang, C.K.: Deceptive-nerf: Enhancing nerf reconstruction using pseudo-observations from diffusion models. arXiv preprint arXiv:2305.15171 (2023) [4](#)
23. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023) [2](#), [3](#), [4](#), [5](#), [10](#), [11](#), [12](#), [13](#)
24. Long, X., Guo, Y.C., Lin, C., Liu, Y., Dou, Z., Liu, L., Ma, Y., Zhang, S.H., Habermann, M., Theobalt, C., et al.: Wonder3d: Single image to 3d using cross-domain diffusion. arXiv preprint arXiv:2310.15008 (2023) [2](#), [5](#), [10](#), [11](#), [12](#), [13](#)
25. Long, X., Lin, C., Wang, P., Komura, T., Wang, W.: Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In: European Conference on Computer Vision. pp. 210–227. Springer (2022) [4](#)
26. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Vedaldi, A.: Realfusion: 360deg reconstruction of any object from a single image. In: CVPR (2023) [1](#), [12](#)
27. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020) [1](#), [5](#), [6](#)
28. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) **41**(4), 1–15 (2022) [13](#)
29. Nichol, A., Jun, H., Dhariwal, P., Mishkin, P., Chen, M.: Point-e: A system for generating 3d point clouds from complex prompts. arXiv preprint arXiv:2212.08751 (2022) [12](#)
30. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: ICLR (2023) [1](#), [4](#), [9](#), [13](#)
31. Qian, G., Mai, J., Hamdi, A., Ren, J., Siarohin, A., Li, B., Lee, H.Y., Skorokhodov, I., Wonka, P., Tulyakov, S., et al.: Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. arXiv preprint arXiv:2306.17843 (2023) [9](#), [12](#)
32. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (2021), <http://proceedings.mlr.press/v139/radford21a.html> [4](#)
33. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021) [4](#)
34. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) [3](#)
35. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22500–22510 (2023) [5](#), [6](#), [9](#)

36. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E.L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al.: Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS (2022)* **13**
37. Sajjadi, M.S., Meyer, H., Pot, E., Bergmann, U., Greff, K., Radwan, N., Vora, S., Lučić, M., Duckworth, D., Dosovitskiy, A., et al.: Scene representation transformer: Geometry-free novel view synthesis through set-latent scene representations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6229–6238 (2022) **4**
38. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. *CoRR abs/2308.16512* (2023). <https://doi.org/10.48550/ARXIV.2308.16512>, <https://doi.org/10.48550/arXiv.2308.16512> **2, 5**
39. Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Viewset diffusion: (0-)image-conditioned 3d generative models from 2d data. *CoRR abs/2306.07881* (2023). <https://doi.org/10.48550/ARXIV.2306.07881>, <https://doi.org/10.48550/arXiv.2306.07881> **4**
40. Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Viewset diffusion:(0-) image-conditioned 3d generative models from 2d data. *arXiv preprint arXiv:2306.07881* (2023) **4**
41. Takagi, Y., Nishimoto, S.: High-resolution image reconstruction with latent diffusion models from human brain activity. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14453–14463 (2023) **6, 11**
42. Tang, S., Zhang, F., Chen, J., Wang, P., Furukawa, Y.: Mvdifusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint arXiv:2307.01097* (2023) **4**
43. Tewari, A., Yin, T., Cazenavette, G., Rezhikov, S., Tenenbaum, J.B., Durand, F., Freeman, W.T., Sitzmann, V.: Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *arXiv preprint arXiv:2306.11719* (2023) **4**
44. Tremblay, J., Meshry, M., Evans, A., Kautz, J., Keller, A., Khamis, S., Müller, T., Loop, C., Morrical, N., Nagano, K., et al.: Rtmv: A ray-traced multi-view synthetic dataset for novel view synthesis. *arXiv preprint arXiv:2205.07058* (2022) **8**
45. Tseng, H.Y., Li, Q., Kim, C., Alsisan, S., Huang, J.B., Kopf, J.: Consistent view synthesis with pose-guided diffusion models. In: *CVPR (2023)* **3, 4**
46. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: *CVPR (2023)* **1, 4**
47. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: *NeurIPS (2021)* **3, 5, 12**
48. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *arXiv preprint arXiv:2305.16213* (2023) **1, 4, 11, 13**
49. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *TIP (2004)* **12**
50. Watson, D., Chan, W., Martin-Brualla, R., Ho, J., Tagliasacchi, A., Norouzi, M.: Novel view synthesis with diffusion models. *arXiv preprint arXiv:2210.04628* (2022) **2, 4**
51. Weng, H., Yang, T., Wang, J., Li, Y., Zhang, T., Chen, C., Zhang, L.: Consistent123: Improve consistency for one image to 3d object synthesis. *arXiv preprint arXiv:2310.08092* (2023) **5**

52. Xiang, J., Yang, J., Huang, B., Tong, X.: 3d-aware image generation using 2d diffusion models. arXiv preprint arXiv:2303.17905 (2023) 4
53. Xu, D., Jiang, Y., Wang, P., Fan, Z., Wang, Y., Wang, Z.: Neurallift-360: Lifting an in-the-wild 2d photo to A 3d object with 360° views. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023. pp. 4479–4489. IEEE (2023). <https://doi.org/10.1109/CVPR52729.2023.00435>, <https://doi.org/10.1109/CVPR52729.2023.00435> 4
54. Xu, J., Wang, X., Cheng, W., Cao, Y.P., Shan, Y., Qie, X., Gao, S.: Dream3d: Zero-shot text-to-3d synthesis using 3d shape prior and text-to-image diffusion models. In: CVPR (2023) 4
55. Yang, J., Cheng, Z., Duan, Y., Ji, P., Li, H.: Consistnet: Enforcing 3d consistency for multi-view images diffusion. arXiv preprint arXiv:2310.10343 (2023) 5
56. Yoo, P., Guo, J., Matsuo, Y., Gu, S.S.: Dreamsparse: Escaping from plato’s cave with 2d frozen diffusion model given sparse views. CoRR (2023) 4
57. Yu, J.J., Forghani, F., Derpanis, K.G., Brubaker, M.A.: Long-term photometric consistent novel view synthesis with diffusion models. In: ICCV (2023) 4
58. Yu, X., Guo, Y.C., Li, Y., Liang, D., Zhang, S.H., Qi, X.: Text-to-3d with classifier score distillation. arXiv preprint arXiv:2310.19415 (2023) 8
59. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) 12
60. Zhou, Z., Tulsiani, S.: Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In: CVPR (2023) 4