

NVS-Adapter: Plug-and-Play Novel View Synthesis from a Single Image

Yoonwoo Jeong^{*1}, Jinwoo Lee^{*2}, Chiheon Kim³,
Minsu Cho^{†1}, and Doyup Lee^{†3}

POSTECH¹ Cinamon² Runway³

Abstract. Recent advancements in Novel View Synthesis (NVS) from a single image have produced impressive results by leveraging the generation capabilities of pre-trained Text-to-Image (T2I) models. However, previous NVS approaches require extra optimization to use other plug-and-play image generation modules such as ControlNet and LoRA, as they fine-tune the T2I parameters. In this study, we propose an efficient plug-and-play adaptation module, NVS-Adapter, that is compatible with existing plug-and-play modules without extensive fine-tuning. We introduce target views and reference view alignment to improve the geometric consistency of multi-view predictions. Experimental results demonstrate the compatibility of our NVS-Adapter with existing plug-and-play modules. Moreover, our NVS-Adapter shows superior performance over state-of-the-art methods on NVS benchmarks although it does not fine-tune billions of parameters of the pre-trained T2I models. The code and data are publicly available at postech-cvlab.github.io/nvsadapter/¹

1 Introduction

The Novel View Synthesis (NVS) task, *i.e.*, generating novel views from a set of images, has drawn significant attention in 3D vision due to its wide applicability in AR, VR, and robotics. Among different setups for the task, NVS from a single image is particularly challenging due to limited information in generating unseen and occluded regions of objects and scenes. As generative models have shown remarkable progress in learning geometry and semantics of the real world, recent methods often leverage the models as a prior to reconstruct 3D geometry from a single image, *e.g.*, by training diffusion models [18, 24, 51] on NVS datasets to generate reliable multi-views [58, 66] or neural fields [2, 4, 16].

Despite their notable efficacy, these methods require categorical priors of 3D objects due to the limited scale of NVS datasets, resulting in a lack of generalizability to unseen visual objects. As a breakthrough, recent work [2, 4, 16] aims to improve generalizability by fine-tuning pre-trained Text-to-Image (T2I) models on large-scale 3D object datasets [8, 9]. However, since they fine-tune pre-trained

^{*} Equally contributed.

[†] Corresponding authors.

¹ This work was done during the summer internship program at KakaoBrain.

T2I models, they require additional fine-tuning [48] to adopt plug-and-play image generation modules such as ControlNet [64] and LoRA [19]. This fine-tuning process necessitates expensive datasets that provide both the new condition and multi-view images, limiting the usage in practice.

In this work, we propose an efficient framework for training a plug-and-play T2I adaptation module, *NVS-Adapter*, for NVS from a single image by freezing the parameters of the pre-trained T2I model. NVS-Adapter aligns the geometries of multi-views, resulting in consistent multi-view predictions that are crucial for reliable 3D generation. Moreover, our framework, which trains only additional cross-attentions, obviates the necessity of training billions of parameters, in contrast to previous approaches. In detail, our NVS-Adapter consists of two trainable components: *view-consistency cross-attention* and *global semantic conditioning*. The view-consistency cross-attention learns visual correspondences between views to align their local details while efficiently aggregating features from multi-views. Additionally, the global semantic conditioning facilitates the model to grasp global structure of visual objects.

Experimental results demonstrate that NVS-Adapter synthesizes geometrically consistent multi-views, achieving superior performance to previous approaches although we do not train billions of parameters of the T2I model. We further verify that NVS-Adapter is fully compatible with existing plug-and-play modules such as ControlNet [64] and LoRA [19], allowing controlled and customized multi-view generation without additional training.

Our contributions can be summarized as follows: 1) We propose a novel plug-and-play NVS-Adapter with two view-alignment cross-attentions. We additionally adopt global semantic conditioning to facilitate the model to grasp the global structure of visual objects. 2) We demonstrate that NVS-Adapter generates geometrically consistent multi-views from a single image. 3) We verify that our plug-and-play adaptation module is fully compatible with other plug-and-play T2I modules without additional training.

2 Related Work

Plug-and-Play Image Generation Modules Existing studies propose plug-and-play image generation modules for conditional [34, 64] and customized [19] image generation by preserving the parameters of pre-trained T2I models. T2I-Adapter [34] optimizes additional trainable projectors to handle auxiliary conditions. ControlNet [64] clones the T2I model and optimizes the cloned one as a condition handler. LoRA [19] efficiently optimizes residual parameters by rank decomposition matrices for customized image generation. Recently, Zero123++ [48] pioneers a controllable multi-view generation by fine-tuning a multi-view diffusion model with depth-conditioned ControlNet. However, it still requires an additional training process to leverage the depth-conditioned ControlNet.

Generative Models for 3D Objects Generative models have shown remarkable progress in synthesizing 3D objects. Geometry-aware generative mod-

els [3, 15, 39, 59] adopt adversarial training to implicitly learn 3D representation of faces [3] or scenes [39, 59] from 2D images. Recent studies train diffusion models [18, 24, 51] to directly generate 3D representations [2, 7, 13, 22, 23, 35, 37, 38, 40, 50, 63]. As another direction, several work trains diffusion models on multi-view datasets [5, 43, 62] to synthesize novel views via cross-attentions [58, 66] or implicit feature fields [4, 16, 61].

Text-to-3D Generation via 2D Diffusion Prior Recent work on T2I diffusion models enables generating 3D representations from a text prompt without 3D data. The seminal work [41] proposes score distillation sampling (SDS) that optimizes a 3D representation by supervising its rendered views with the scores of pre-trained T2I models. Along with the advances in neural fields [36, 47, 55, 56], the existing studies [6, 27, 54, 57] have consistently improved the performance of SDS. As another direction, low-rank adaptation [19], structural conditions [64], and view-dependent texts [41] are leveraged to maintain the multi-view consistency of generated 3D representations [32, 46, 57]. However, the generated 3D contents still lack geometric consistency despite their high fidelity since the training data of T2I models do not contain precise information of underlying 3D geometries.

Fine-tuning T2I Models for Novel View Synthesis Existing studies modify and fine-tune T2I models on multi-view datasets to generate high-quality and diverse images for novel view synthesis. Zero-1-to-3 [8, 29] concatenates a reference image into the input of U-Net after modifying text cross-attentions for image embeddings with relative camera extrinsics [26]. Despite fine-tuning on Objaverse [8, 9], generated views suffer from geometric inconsistency and inaccurate camera viewpoints since the model cannot explicitly learn visual correspondences between the reference and generated views [48]. ConsistNet [60] and SyncDreamer [30] are built on Zero-1-to-3 to synthesize multiple novel views with an unprojection operation. MVDream [49] leverages 3D self-attentions for 3D-aware multi-view generation. Wonder3D [31] fine-tunes the Stable Diffusion Image Variations Model [26] to generate multi-view normal maps alongside color images. Zero123++ [48] and One-2-3-45++ [28], which are concurrent to our work, adopt the reference attention technique [65] and fine-tune the original self-attentions for NVS from a single image. However, they are limited to render only fixed viewpoints. MVD-Fusion [20] jointly infers depth to produce multi-view consistent novel views. In detail, it generates multi-view RGB-D images and leverages the depth estimates to obtain reprojection-based conditioning to maintain multi-view consistency. EscherNet [25] learns implicit and generative 3D representations coupled with a specialised camera positional encoding to predict novel views. Despite their capability for high-quality generation, these models still require training billions of parameters, as they fine-tune the pa-

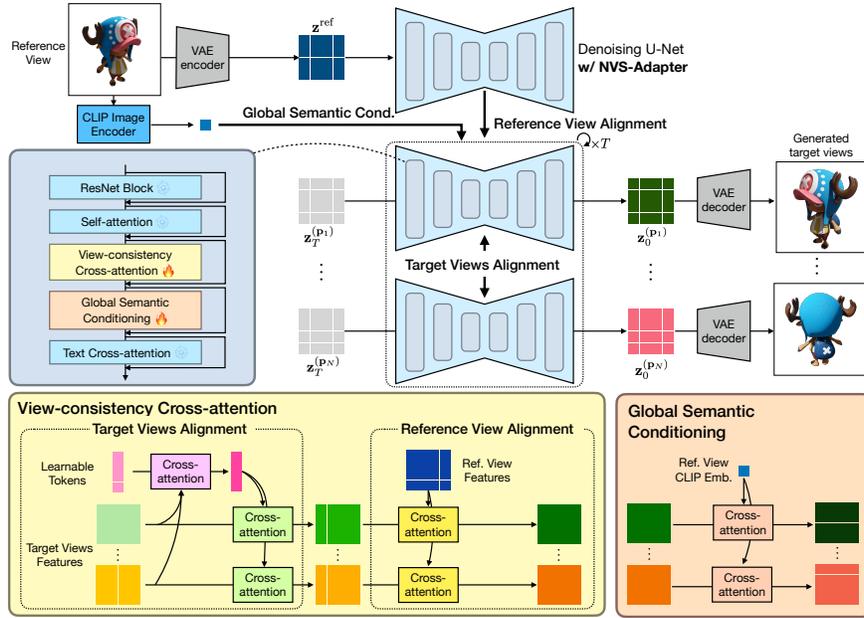


Fig. 1: Overview of our framework. NVS-Adapter synthesizes novel views by incorporating two components into each U-Net block of a pre-trained T2I model: *view-consistency cross-attention* which aligns features of each target view with the relevant features of other views, and *global semantic conditioning* which aligns the features of target views with the global semantic structure of the reference.

rameters of T2I models. In contrast, our framework enables efficient training by training only additional cross-attention modules.

3 Methods

Our framework synthesizes novel views from a single image for diverse 3D objects, while preserving the original parameters of T2I models. Figure 1 illustrates the overall architecture. We first introduce the preliminaries of diffusion models and novel view synthesis, and then formulate the task of novel view synthesis from a single image. Lastly, we explain how our NVS-Adapter synthesizes geometrically consistent multi-views of visual objects in a plug-and-play manner.

3.1 Preliminaries

Diffusion Models Diffusion models learn a data distribution $q(\mathbf{x})$ using gradually corrupted latents $\mathbf{x}_1, \dots, \mathbf{x}_T$, where $q(\mathbf{x}_t | \mathbf{x}_0) := \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t} \mathbf{x}_0, (1 - \alpha_t) \mathbf{I})$ and $\mathbf{x}_0 = \mathbf{x}$. Then, a diffusion model $p_\theta(\mathbf{x}_{0:T})$ aims to predict the reverse process of the Markov chain with Gaussian transitions to denoise the latents as

$p_\theta(\mathbf{x}_{0:T}) := p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)$, where $p(\mathbf{x}_T) := \mathcal{N}(\mathbf{x}_T|\mathbf{0}, \mathbf{I})$, $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_t; \mu_\theta(\mathbf{x}_t, t), \sigma_t^2 \mathbf{I})$, and σ_t is a constant at t . Given a constant α_t at timestep t , the predicted mean is defined as $\mu_\theta(\mathbf{x}_t, t) := (\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_\theta(\mathbf{x}_t, t)) / \sqrt{\alpha_t}$, and the noise predictor ϵ_θ minimizes the reweighted variational lower bound [18]:

$$L_{\text{simple}} = \mathbb{E}_{t, \mathbf{x}_0, \epsilon} [\|\epsilon - \epsilon_\theta(\mathbf{x}_t, t)\|^2], \quad (1)$$

where $\mathbf{x}_t = \sqrt{\alpha_t} \mathbf{x}_0 + \sqrt{1 - \alpha_t} \epsilon$, and $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Latent Diffusion Models Latent diffusion models (LDMs) [44] generate high-resolution images with their two-stage framework. In the first stage, an auto-encoder [14], which consists of an encoder \mathcal{E} and a decoder \mathcal{D} , learns to compress a high-resolution RGB image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ into a low-resolution feature map $\mathbf{z} = \mathcal{E}(\mathbf{x}) \in \mathbb{R}^{H/f \times W/f \times C}$ with the encoder and reconstruct the image from the latent as $\mathbf{x} \approx \mathcal{D}(\mathcal{E}(\mathbf{x}))$ with the decoder, where f is the downsampling factor and C is the dimension of latent features. In the second stage, the U-Net [10, 45] of diffusion models processes features on the latent space of \mathbf{z} instead of the image space of \mathbf{x} . Each U-Net block of LDMs, such as Stable Diffusion, comprises a ResNet block followed by a self-attention of image features and a cross-attention between image and text features.

Novel View Synthesis from a Single Image Given a reference view of a 3D object, we aim to generate rendered images from novel viewpoints. Let the reference RGB image $\mathbf{x}^{\text{ref}} \in \mathbb{R}^{H \times W \times 3}$, N relative camera viewpoints to a reference viewpoint $\mathbf{P} = \{\mathbf{p}_1, \dots, \mathbf{p}_N\}$, the rendered novel views $\mathbf{x}^{(\mathbf{P})} = \{\mathbf{x}^{(\mathbf{p}_1)}, \dots, \mathbf{x}^{(\mathbf{p}_N)}\}$ are given, where $\mathbf{p}_i = [R_i; T_i] \in \mathbb{R}^{3 \times 4}$ is i -th camera viewpoint with relative camera rotation $R_i \in \mathbb{R}^{3 \times 3}$ and translation $T_i \in \mathbb{R}^3$, and its rendered novel view $\mathbf{x}^{(\mathbf{p}_i)}$, respectively. Then, a generative model learns a conditional distribution $q(\mathbf{x}^{(\mathbf{P})}|\mathbf{x}^{\text{ref}}, \mathbf{P})$ to synthesize novel views $\mathbf{x}^{(\mathbf{P})}$ given a reference view \mathbf{x}^{ref} . For the brevity of notation, we denote $q(\mathbf{x}^{(\mathbf{P})}|\mathbf{x}^{\text{ref}}, \mathbf{P})$ by $q(\mathbf{x}^{(\mathbf{P})}|\mathbf{x}^{\text{ref}})$. Note that a powerful generative model such as a diffusion model is required, since the model has to extrapolate unseen or occluded parts of the 3D object.

3.2 LDMs for NVS from a Single Image

We formulate the task of synthesizing novel views from a single image based on LDMs, since we exploit the pre-trained LDM, Stable Diffusion. Specifically, a LDM learns a conditional distribution $q(\mathbf{z}^{(\mathbf{P})}|\mathbf{z}^{\text{ref}})$, where $\mathbf{z}^{\text{ref}} = \mathcal{E}(\mathbf{x}^{\text{ref}})$ and $\mathbf{z}^{(\mathbf{P})} = \{\mathbf{z}^{(\mathbf{p}_i)}\}_{i=1}^N$ denote latent features of a reference view and N target views, respectively. For simplicity, previous approaches [29, 58] assume that all the target views $\mathbf{z}^{(\mathbf{P})}$ are independently drawn from the reference image, *i.e.*, $p_\theta(\mathbf{z}_{t-1}^{(\mathbf{P})}|\mathbf{z}_t^{(\mathbf{P})}, \mathbf{z}^{\text{ref}}) = \prod_{i=1}^N p_\theta(\mathbf{z}_{t-1}^{(\mathbf{p}_i)}|\mathbf{z}_t^{(\mathbf{p}_i)}, \mathbf{z}^{\text{ref}})$, where noised latent features of the target views are $\mathbf{z}_t^{(\mathbf{p}_i)} = \sqrt{\alpha_t} \mathbf{z}_0^{(\mathbf{p}_i)} + \sqrt{1 - \alpha_t} \epsilon$ with $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Then, the approximate posterior at a specific timestep t is formulated as $q(\mathbf{z}_t^{(\mathbf{P})}|\mathbf{z}_0^{(\mathbf{P})}, \mathbf{z}^{\text{ref}}) =$

$\prod_{i=1}^N q(\mathbf{z}_t^{(\mathbf{P}^i)} | \mathbf{z}_0^{(\mathbf{P}^i)}, \mathbf{z}^{\text{ref}})$. However, the independence assumption leads to geometric inconsistency across predicted target views, since the generation process does not align the geometries among the target views. Thus, our denoising process aligns views to consider the geometric correspondences among the target views:

$$p_\theta(\mathbf{z}_{t-1}^{(\mathbf{P})} | \mathbf{z}_t^{(\mathbf{P})}, \mathbf{z}^{\text{ref}}) = \prod_{i=1}^N p_\theta(\mathbf{z}_{t-1}^{(\mathbf{P}^i)} | \mathbf{z}_t^{(\mathbf{P}^i)}, \mathbf{z}^{\text{ref}}), \quad (2)$$

where the latent features of each i -th target view $\mathbf{z}_{t-1}^{(\mathbf{P}^i)}$ is aligned with the reference view \mathbf{z}^{ref} and the other target views $\{\mathbf{z}_t^{(\mathbf{P}^j)}\}_{j \in \{1, \dots, N\} \setminus \{i\}}$. Note that we do not corrupt the reference view to preserve fine-grained details of visual objects.

3.3 NVS-Adapter

We propose a plug-and-play T2I adaptation module, *NVS*, to effectively synthesize consistent novel multi-views from a single image. We incorporate the adaptation modules into each U-Net block as illustrated in Figure 1. While the features of different views are separately processed, the two components, *view-consistency cross-attention* and *global semantic conditioning*, effectively align the intermediate features across views. The view-consistency cross-attention aligns the intermediate target features to have consistent local details by learning visual correspondences between different views. The global semantic conditioning provide the semantic information of visual objects in the reference image.

View-Consistency Cross-Attention Our view-consistency cross-attention sequentially conducts the *target views alignment* and the *reference view alignment*, which align a target view with the other target views and the reference view, respectively. Since every block in a U-Net share the same details, we describe the details in on one specific block without the loss of generality.

Target Views Alignment The ResNet and the self-attention block from the pre-trained T2I model separately extracts the features $\mathbf{f}^{(\mathbf{P})} = \{\mathbf{f}^{(\mathbf{P}^i)}\}_{i=1}^N$ of target views, where $\mathbf{f}^{(\mathbf{P}^i)} \in \mathbb{R}^{h \times w \times c}$ is the feature of the i -th target view with a $h \times w$ resolution and c channels. Inspired by the attention operation of Perceiver [21], we aggregate the features of the target views using L learnable tokens $\mathbf{q} \in \mathbb{R}^{L \times c}$ as

$$\bar{\mathbf{q}} := \text{MHA}(\text{Q} = \mathbf{q}, \text{KV} = \mathbf{f}^{(\mathbf{P})} + \boldsymbol{\gamma}^{(\mathbf{P})}), \quad (3)$$

where Q and KV denote the query and key-value of a multi-head attention (MHA) [53] operation, and $\boldsymbol{\gamma}^{(\mathbf{P})} = \{\boldsymbol{\gamma}^{(\mathbf{P}^i)}\}_{i=1}^N$ denotes the positional embeddings of the camera poses \mathbf{P} of the target views. Note that we use relative camera poses to the reference view. When we encode the positional embeddings $\boldsymbol{\gamma}^{(\mathbf{P}^i)} \in \mathbb{R}^{h \times w \times c}$ under the i -th view into sinusoidal embeddings [33], we represent ray positions with the Plücker coordinate [11] to effectively learn the geometric relation among rays. We remark that the parameters of learnable tokens \mathbf{q}

are shared across all target views. Subsequently, each target view leverages $\bar{\mathbf{q}}$ to aggregate the matched features from the other target views as

$$\mathbf{f}_{\text{TA}}^{(\mathbf{P})} := \mathbf{f}^{(\mathbf{P})} + \text{MHA}(\text{Q} = \mathbf{f}^{(\mathbf{P})} + \boldsymbol{\gamma}^{(\mathbf{P})}, \text{KV} = \bar{\mathbf{q}}), \quad (4)$$

where $\mathbf{f}_{\text{TA}}^{(\mathbf{P})}$ are target-views-aligned features of $\mathbf{f}^{(\mathbf{P})}$.

Reference View Alignment After the target views alignment in Eq. (4), we perform the reference view alignment to ensure the consistency of details between the reference view, which is the only condition for NVS from a single image, and the generated target views. In detail, the model forwards an additional MHA operation as

$$\mathbf{f}_{\text{RA}}^{(\mathbf{P})} := \mathbf{f}_{\text{TA}}^{(\mathbf{P})} + \text{MHA}(\text{Q} = \mathbf{f}_{\text{TA}}^{(\mathbf{P})} + \boldsymbol{\gamma}^{(\mathbf{P})}, \text{KV} = \mathbf{f}^{\text{ref}} + \boldsymbol{\gamma}^{\text{ref}}), \quad (5)$$

where \mathbf{f}^{ref} and $\boldsymbol{\gamma}^{\text{ref}}$ are the ResNet feature and the positional embedding of the reference view. Similarly, the reference view alignment is applied to the reference features, considering it as a self-attention operation:

$$\mathbf{f}_{\text{RA}}^{\text{ref}} := \mathbf{f}^{\text{ref}} + \text{MHA}(\text{Q} = \mathbf{f}_{\text{TA}}^{\text{ref}} + \boldsymbol{\gamma}^{\text{ref}}, \text{KV} = \mathbf{f}^{\text{ref}}). \quad (6)$$

The aligned features $\mathbf{f}_{\text{RA}}^{(\mathbf{P})}$ and $\mathbf{f}_{\text{RA}}^{\text{ref}}$ are used as an input to the next module, global semantic conditioning.

Global Semantic Conditioning Our NVS-Adapter also incorporates a cross-attention of global semantic conditioning [29] to inject the global semantic structure of visual objects in the reference image to different novel views. Note that NVS from a single image assumes a reference image is given without its text captions, unlike the conventional T2I setup, which takes the text as input. We extract CLIP [42] image embeddings of the reference view to conduct a cross-attention operation between the features of each view and the CLIP image embeddings, as shown in Figure 1.

4 Experiments

In this section, we conduct extensive experiments to demonstrate the effectiveness of our NVS-Adapter. In Section 4.1, we evaluate our framework on NVS benchmarks and also show qualitative 3D generation results on various scenarios such as image-to-3D (I23D) and T23D. Then, in Section 4.2, we show that our plug-and-play module is compatible with other plug-and-play T2I modules without extra optimization. Furthermore, in Section 4.3, we show our model generates high-fidelity radiance fields via Score Distillation Sampling (SDS). Lastly, in Section 4.4 and 4.5, we provide in-depth studies to analyze the efficacy of our NVS-Adapter to understand the geometry across different viewpoints of visual objects.

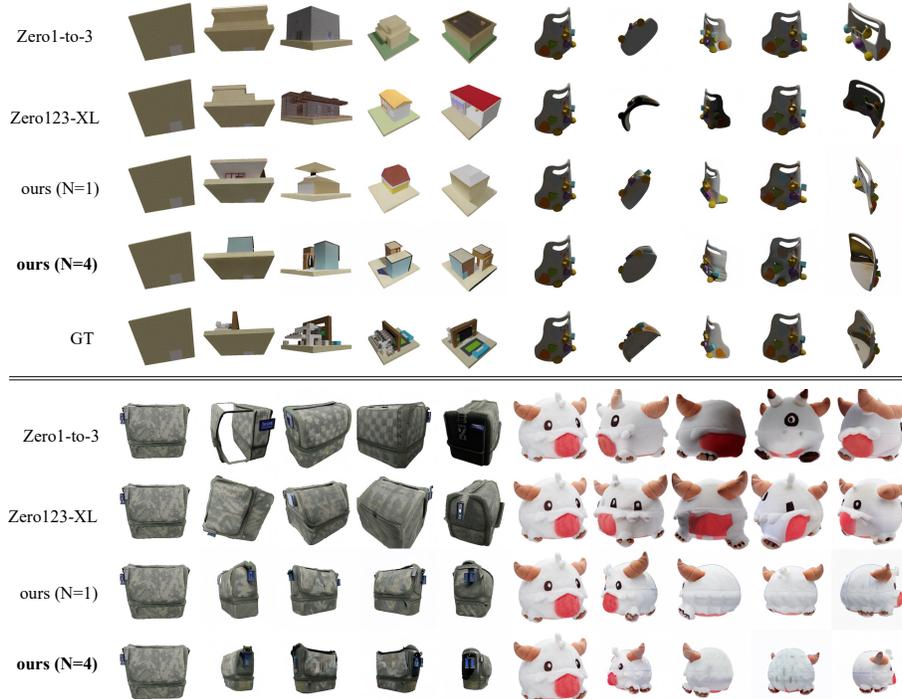


Fig. 2: Novel view synthesis examples by Zero-1-to-3 [29], Zero123-XL [8], and our NVS-Adapter with $N = 1$ and $N = 4$. Top images: NVS results conditioned on an image from Objaverse [9] and GSO [12] validation set. Bottom images: NVS results conditioned on a single image used in SyncDreamer [30]. The first column presents a reference image, and the rest of four columns are synthesized views by each model. Note that the "Bottom images" do not have the ground truth images for target viewpoints since they are not from multi-view datasets.

Implementation Details We implement our NVS-Adapter on Stable Diffusion 2.1-base [1] (SD2.1) and train our framework on Objaverse [9], which includes 800K 3D assets. To demonstrate the compatibility with ControlNet and LoRA models, we repeat this process on Stable Diffusion 1.5-base (SD1.5) instead of SD2.1, as SD1.5 provides a wider range of ControlNet [64] and LoRA [19] models and thereby allows comprehensive experimentation. For a fair comparison with Zero-1-to-3 [29], we use the same training dataset with the image resolution of 256×256 . Our NVS-Adapter uses $h \times w$ learnable tokens \mathbf{q} for target views alignment in Eq. (3) at every U-Net block, while aggregating $N = 4$ target views. Our model is trained for 200K iterations with the batch size of 256, using 16 NVIDIA A100 80GB GPUs for 5 days. For more details, please refer to the supplementary materials.

Table 1: Comparison of NVS scores on Objaverse [9] (left) and Google Scanned Objects [12] (right) with Zero-1-to-3 [29] and Zero123-XL [8]. We also report scores when our model is combined with ControlNet [64] with three conditions: depth, canny, and HED. Note that no additional fine-tuning is required to attach the ControlNet.

	Objaverse			Google Scanned Objects		
	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
Zero-1-to-3	19.52	0.8603	0.1251	18.36 \pm 0.02	0.8418 \pm 0.0005	0.1283 \pm 0.0006
Zero123-XL	17.71	0.8258	0.1552	18.46 \pm 0.11	0.8351 \pm 0.0015	0.1267 \pm 0.0009
ours (SD21)	19.58	0.8658	0.1135	18.80 \pm 0.07	0.8469 \pm 0.0010	0.1207 \pm 0.0010
ours (SD15)	17.05	0.8338	0.1480	17.87 \pm 0.14	0.8331 \pm 0.0004	0.1327 \pm 0.0007
w. canny	20.82	0.8788	0.0977	20.64 \pm 0.03	0.8624 \pm 0.0009	0.1030 \pm 0.0004
w. depth	20.99	0.8839	0.0900	20.20 \pm 0.03	0.8611 \pm 0.0006	0.1057 \pm 0.0002
w. HED	22.02	0.8982	0.0811	22.77 \pm 0.03	0.8820 \pm 0.0003	0.0912 \pm 0.0004

4.1 Novel View Synthesis from a Single Image

We evaluate our framework on the validation split of Objaverse [9] and a subset of Google Scanned Objects (GSO) [12], containing high-quality scanned 3D assets. We use 12 and 16 rendered views for each object following the rendering pipeline of Zero-1-to-3 [29] on Objaverse and GSO, respectively. We randomly partition 16 target views to have 4 target views for each batch and report the mean and standard deviation of results with three repetitions for GSO. Due to the huge computational budgets of the validation on Objaverse, we omit the repetition. Table 1 (left) shows the competitive results with previous methods of Zero-1-to-3 and Zero123-XL in terms of PSNR, SSIM, and LPIPS [52]. In Table 1 (right), our NVS-Adapter also outperforms Zero-1-to-3 and Zero123-XL on GSO, achieving higher generalization performance. Note that our NVS-Adapter shows notably better LPIPS score than previous approaches, indicating that our model fully exploits the generation capacity of the pre-trained T2I model. We also remark that Zero-1-to-3 and Zero123-XL fine-tune over one billion parameters, and Zero123-XL learns 12 times more 3D assets than our NVS-Adapter. These results indicate our plug-and-play NVS-Adapter can efficiently and effectively adapt a T2I model to synthesize novel views from a single image.

Figure 2 visualizes the synthesized novel views by our NVS-Adapter and previous methods [29]. Additionally, we visualize NVS results on images used in SyncDreamer [30] to show that our NVS-Adapter well generalizes on arbitrary images that are not from multi-view datasets. Zero-1-to-3 [29] and Zero123-XL [8, 29] commonly suffer from the inconsistency of generated views, since they cannot synthesize multi-views at once. Although NVS-Adapter with $N = 1$ lacks the consistency of generated views, it shows better correspondences between the reference and target views than Zero-1-to-3, since the view-consistency cross-attention layers learn the correspondences among view features. In contrast, our NVS-Adapter with $N = 4$ well aligns the features of multi-views and effectively generates consistent multi-views at once.

4.2 Compatibility with ControlNets and LoRAs

Since NVS-Adapter is fully compatible with existing plug-and-play modules, our framework can cover various scenarios on NVS without extra optimization. We demonstrate the effectiveness of our NVS-Adapter on two widely used scenarios, controllable generation with ControlNet [64] and customized generation with LoRA [19], tailored to NVS from a single image.

For controllable NVS with ControlNet, we exploit the three conditions of ControlNet, canny edge, depth, and HED, to guide the multi-view generation. Table 1 shows that our NVS-Adapter can exploit canny edge, depth, and HED conditions to significantly improve the NVS performance. Figure 4 (left) shows that the synthesized novel views are guided by each condition. For instance, the first two examples are generating novel views in unobserved regions. NVS-Adapters with ControlNet address the ambiguities by the guidance from conditions. In the last example, our NVS-Adapter without ControlNet suffers from intricate geometries of a molecular toy while the ControlNet-guided NVS-Adapters generate accurate results.

For customized NVS with LoRA, we update the pre-trained T2I model with three LoRA models -friedegg, gemstone, and blueresin- that are tailored to generate objects with unique textures. We use the same LoRA models for generating a reference image and multi-view images. We found that our NVS-Adapter with LoRA preserves the styles provided by the LoRA when generating multi-views as visualized in Figure 4 (right). According to the results, NVS-Adapter with LoRA avoids oversimplified appearance (up), captures complex details (middle), and includes physical phenomena like reflection (down). We remark that our framework does not require additional training for the compatibility with ControlNets and LoRAs, different from previous approaches [48]. For more details, please refer to the supplementary materials.

4.3 3D Reconstruction

We leverage our framework with NVS-Adapter to reconstruct 3D geometry from a single image, incorporating Score Distillation Sampling (SDS) [41]. We freeze the parameters of our framework and optimize InstantNGP [36] to generate 3D consistent representations. We use threestudio [17] implementations based on the setting of Zero-1-to-3 to compare our framework with other methods. We also employ opacity regularization to reduce noisy structures. In the coarse stage of training, we optimize the model for 5,000 iterations with batch size 16. Following the coarse stage, we further optimize the model for 3,000 iterations with a reduced batch size of 8. We sample a timestep t in the range $[0.02, 0.98]$ and linearly decay the maximum timestep t_{max} to 0.5 during the coarse stage [27]. We attach the implementation details to the supplementary materials.

Figure 3 compares the generated 3D representations via SDS of Zero-1-to-3, Zero123-XL, and our framework with NVS-Adapter. We first generate images from input texts as reference views to compare our framework with Dreamfusion [41]. Zero-1-to-3, Zero123-XL, and our framework generate 3D following

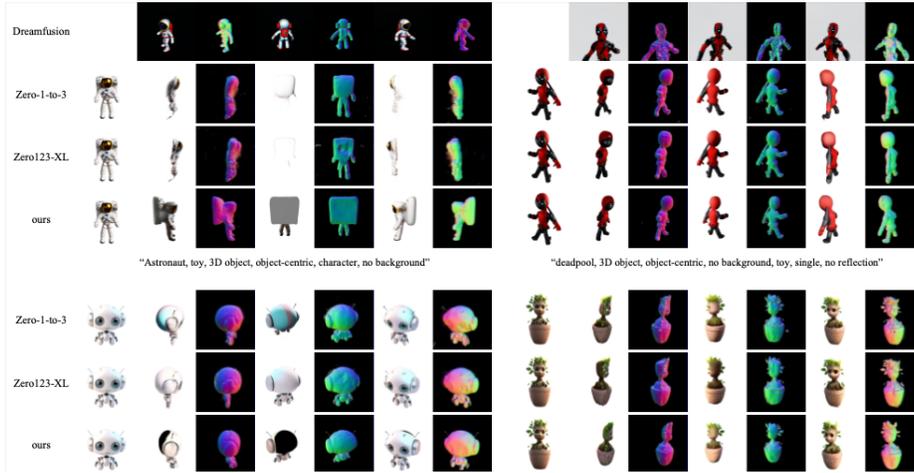


Fig. 3: 3D reconstruction examples via Score Distillation Sampling (SDS) [41] with baselines [8, 29] and our NVS-Adapter. Top images shows 3D reconstructions results conditioned on an image generated by T2I model, and bottom images shows results on an image in the wild.

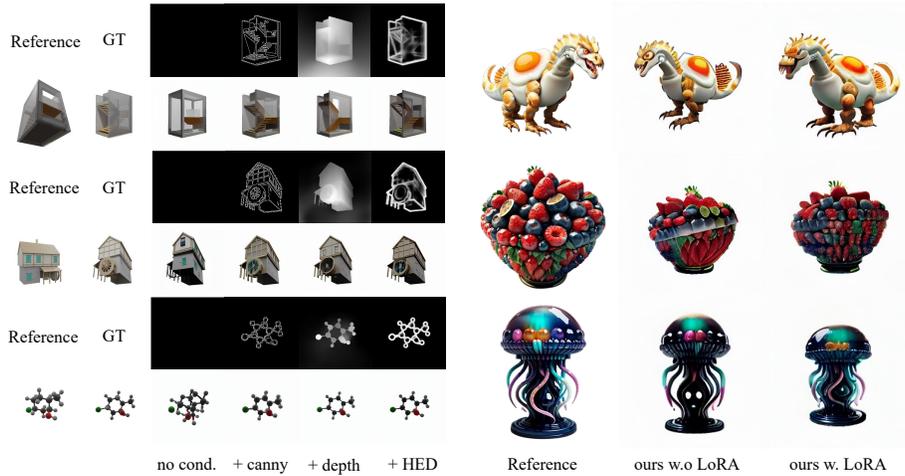


Fig. 4: (left) Examples of NVS results of our NVS-Adapter with and without ControlNet [64] variants. (right) Examples of NVS results of our NVS-Adapter with and without LoRA [19] models. We use the same LoRA models for generating a reference image and multi-view images.

the high fidelity of images generated from texts, unlike the oversimplified results of Dreamfusion. However, Zero-1-to-3 and Zero123-XL suffer from geometric inconsistency, but our framework shows high-quality results following the ref-

Table 2: The effects of the number of learnable tokens $L = h \times w \times c$ in target views alignment for NVS on GSO [12].

	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
$c = 0.5$	18.26 ± 0.05	0.8406 ± 0.0010	0.1277 ± 0.0007
$c = 1$	18.56 ± 0.09	0.8446 ± 0.0018	0.1243 ± 0.0015
$c = 2$	18.51 ± 0.10	0.8441 ± 0.0015	0.1251 ± 0.0008

Table 3: The effects of the number of target views N for NVS on GSO [12].

	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
$N = 1$	14.38 ± 0.02	0.8025 ± 0.0003	0.1812 ± 0.0013
$N = 2$	18.55 ± 0.01	0.8451 ± 0.0010	0.1221 ± 0.0003
$N = 4$	19.01 ± 0.03	0.8501 ± 0.0011	0.1190 ± 0.0013
$N = 6$	18.38 ± 0.02	0.8419 ± 0.0012	0.1254 ± 0.0009

Table 4: The effects of using global semantic conditioning (GSC) in a view-consistency cross-attention for NVS on GSO [12].

	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
ours	18.45 ± 0.05	0.8409 ± 0.0003	0.1240 ± 0.0002
w/o GSC	18.27 ± 0.05	0.8388 ± 0.0007	0.1254 ± 0.0004

Table 5: The effects of changing ray positional embeddings in the view-consistency cross-attention for NVS on GSO [12].

	PSNR(\uparrow)	SSIM(\uparrow)	LPIPS(\downarrow)
extrinsic	18.02 ± 0.02	0.8329 ± 0.0006	0.1298 ± 0.000
ray o, ray d	18.41 ± 0.03	0.8423 ± 0.0004	0.1254 ± 0.0004
Plücker	18.45 ± 0.05	0.8409 ± 0.0003	0.1240 ± 0.0002

erence images. Our framework also shows competitive results on image-to-3D, while generating geometrically consistent 3D representations compared to other methods.

4.4 Ablation Study

We conduct an extensive ablation study to evaluate the efficacy of our NVS-Adapter in synthesizing novel views. Due to the computational costs, we evaluate the trained models after 100K training iterations for the ablation study. For additional ablation studies, please refer to the supplementary materials.

The number of learnable tokens. We analyze the effects of the number of learnable tokens \mathbf{q} , which aggregates the features of target views in Eq. (3). With

a fixed number of target views $N = 4$, we vary the number of learnable tokens $L = h \cdot w \cdot c$ with $c = 0.5, 1, 2$ for a $h \times w$ resolution U-Net block. Table 2 shows that increasing the number of learnable tokens improves the results, since a large number of tokens can easily preserve the information of target views. However, NVS-Adapter with $N = 4$ achieves the best performance with $L = h \times w$ although it has fewer tokens than $2 \times h \times w$. Thus, we set the number of learnable tokens to be $h \times w$ for each U-Net block for the rest of experiments.

The number of target views. We analyze the effects of changing the number of target views N . We simultaneously generate all 16 views using our models trained with $N = 1, 2, 4, 6$ to decide the number of target views to use. Since our NVS-Adapter with $N = 1$ is not supervised to jointly generate multiple viewpoints, the model lacks geometric consistency when generating multiple views together. NVS-Adapter with more target views tend to show better performance, but expanding to $N = 6$ deteriorates the overall performance, since the capacity of $L = h \times w$ learnable tokens is limited to aggregate $N > 4$ target views. As shown in Figure 2, our NVS-Adapter with $N = 4$ consistently synthesizes multi-views in a single generation unlike our NVS-Adapter with $N = 1$, which shows different geometries for each predicted viewpoint. We set $N = 4$ for the rest of experiments since it achieves the best performance.

The effects of global semantic conditioning. We conduct an ablation study to validate the effectiveness of global semantic conditioning. Table 4 shows that leveraging CLIP image features of a reference view enhances the semantic consistency of synthesized target views. This indicates the novel views generated without global semantic conditioning suffer from lower quality due to the lack of semantic understanding of visual objects.

Ray positional embeddings. We examine the effects of ray positional embeddings on the view-consistency cross-attention. We change the Plücker coordinates, which represent the ray positions in view features, into the offset and direction vectors of rays. In addition, we also compare the results using the embedding of camera extrinsic parameters, as in Zero-1-to-3, instead of rays. Table 5 shows that utilizing each ray’s position improves the performance, since it allows our view-consistency cross-attention to leverage ray geometry for matching view correspondences.

4.5 Visualization Analysis of Cross-Attention Maps

We analyze the role of our view-consistency cross-attention to align the features of target and reference views. We visualize the mean of cross-attention maps in the last U-Net block over the attention heads.

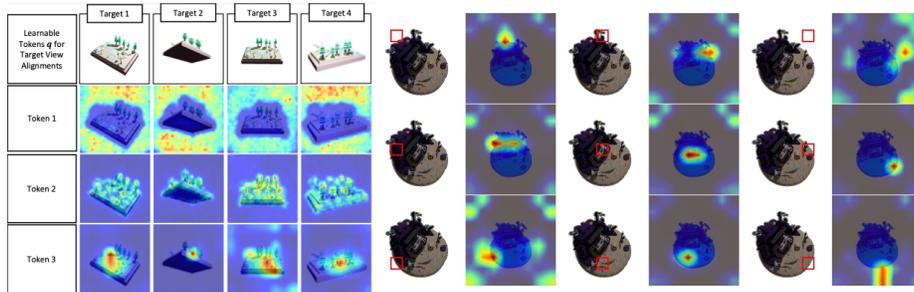


Fig. 5: Visualization of attention maps of target views alignments (left) and reference view alignments (right). For the target views alignments, each token aggregates different regions of target views. For the reference view alignments, the queries in the target view (red boxes) correspond to the highlighted regions in the reference view.

Target views alignment. Figure 5 (left) visualizes the cross-attention maps that use the learnable tokens as queries and the features of target views as keys. Each learnable token effectively aggregates and encodes the information of target views based on the semantic and positional features, reducing redundancy. As an example, token 1 (up) aggregates the white background of all target views, and token 2 (middle) attends to prominent regions in the foreground. Meanwhile, token 3 (down) attends to a specific region of the foreground, while aggregating the same positions across different viewpoints. Thus, our method aggregates the features of target views into a fixed number of tokens, enabling efficient synthesis of well-aligned multi-views.

Reference view alignment. Figure 5 (right) visualizes the cross-attention maps, considering a target and reference view as a query and a key, respectively. Note that the highlighted regions in attention maps correspond to the query regions (red box) in Figure 5. That is, our reference view alignment focuses on matching target and reference view correspondences, aggregating the local details from the reference view for consistent synthesis.

5 Conclusion

We have proposed a plug-and-play module, NVS-Adapter, for T2I models to synthesize geometrically consistent multi-views from a single image. Our NVS-Adapter employs view-consistency cross-attention for aligning target and reference views, and the embedding of a reference view is used for the global semantic condition of views. The experimental results show that our plug-and-play module is fully compatible with existing plug-and-play modules such as ControlNet [64] and LoRA [19]. Moreover, our NVS-Adapter achieves superior performance in NVS by learning geometric correspondences across views.

Acknowledgements

This work was supported by IITP grants (RS-2021-II212068: AI Innovation Hub (50%), RS-2022-II220113: Developing a Sustainable Collaborative Multi-modal Lifelong Learning Framework (45%), RS-2019-II191906: AI Graduate School Program at POSTECH (5%)) funded by the Korea government.

References

1. AI, S.: Stable diffusion version 2. <https://github.com/Stability-AI/stablediffusion> (2022)
2. Anciukevičius, T., Xu, Z., Fisher, M., Henderson, P., Bilen, H., Mitra, N.J., Guerrero, P.: Renderdiffusion: Image diffusion for 3d reconstruction, inpainting and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12608–12618 (2023)
3. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16123–16133 (2022)
4. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aittala, M., Mello, S.D., Karras, T., Wetzstein, G.: GeNVS: Generative novel view synthesis with 3D-aware diffusion models. In: arXiv (2023)
5. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
6. 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)
7. Cheng, Y.C., Lee, H.Y., Tulyakov, S., Schwing, A.G., Gui, L.Y.: Sdfusion: Multi-modal 3d shape completion, reconstruction, and generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4456–4465 (2023)
8. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., et al.: Objaverse-xl: A universe of 10m+ 3d objects. arXiv preprint arXiv:2307.05663 (2023)
9. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13142–13153 (2023)
10. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
11. Dorst, L., Fontijne, D., Mann, S.: *Geometric algebra for computer science: an object-oriented approach to geometry*. Elsevier (2010)
12. 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: 2022 International Conference on Robotics and Automation (ICRA). pp. 2553–2560. IEEE (2022)
13. Erkoç, Z., Ma, F., Shan, Q., Nießner, M., Dai, A.: Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. arXiv preprint arXiv:2303.17015 (2023)

14. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021)
15. Gao, J., Shen, T., Wang, Z., Chen, W., Yin, K., Li, D., Litany, O., Gojcic, Z., Fidler, S.: Get3d: A generative model of high quality 3d textured shapes learned from images. In: Advances In Neural Information Processing Systems (2022)
16. Gu, J., Trevithick, A., Lin, K.E., Susskind, J., Theobalt, C., Liu, L., Ramamoorthi, R.: Nerfdiff: Single-image view synthesis with nerf-guided distillation from 3d-aware diffusion. In: International Conference on Machine Learning (2023)
17. 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)
18. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. arXiv preprint arXiv:2006.11239 (2020)
19. Hu, E.J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: Lora: Low-rank adaptation of large language models. In: International Conference on Learning Representations (2021)
20. Hu, H., Zhou, Z., Jampani, V., Tulsiani, S.: Mvd-fusion: Single-view 3d via depth-consistent multi-view generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9698–9707 (2024)
21. Jaegle, A., Gimeno, F., Brock, A., Vinyals, O., Zisserman, A., Carreira, J.: Perceiver: General perception with iterative attention. In: International conference on machine learning. pp. 4651–4664. PMLR (2021)
22. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023)
23. Karnewar, A., Vedaldi, A., Novotny, D., Mitra, N.J.: Holodiffusion: Training a 3d diffusion model using 2d images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18423–18433 (2023)
24. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. *Advances in Neural Information Processing Systems* **35**, 26565–26577 (2022)
25. Kong, X., Liu, S., Lyu, X., Taher, M., Qi, X., Davison, A.J.: Eschnet: A generative model for scalable view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9503–9513 (2024)
26. Labs, L.: Stable diffusion image variations model card. <https://huggingface.co/lambdalabs/stable-diffusion-image-conditioned> (2022)
27. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 300–309 (2023)
28. Liu, M., Shi, R., Chen, L., Zhang, Z., Xu, C., Chen, H., Zeng, C., Gu, J., Su, H.: One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion (2023)
29. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. arXiv preprint arXiv:2303.11328 (2023)
30. 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)

31. 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)
32. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Vedaldi, A.: Realfusion: 360deg reconstruction of any object from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8446–8455 (2023)
33. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
34. Mou, C., Wang, X., Xie, L., Wu, Y., Zhang, J., Qi, Z., Shan, Y., Qie, X.: T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. arXiv preprint arXiv:2302.08453 (2023)
35. Müller, N., Siddiqui, Y., Porzi, L., Bulo, S.R., Kotschieder, P., Nießner, M.: Diffrrf: Rendering-guided 3d radiance field diffusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4328–4338 (2023)
36. 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)
37. Nam, G., Khlifi, M., Rodriguez, A., Tono, A., Zhou, L., Guerrero, P.: 3d-ldm: Neural implicit 3d shape generation with latent diffusion models. arXiv preprint arXiv:2212.00842 (2022)
38. 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)
39. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2021)
40. Ntavelis, E., Siarohin, A., Olszewski, K., Wang, C., Van Gool, L., Tulyakov, S.: Autodecoding latent 3d diffusion models. arXiv preprint arXiv:2307.05445 (2023)
41. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. In: The Eleventh International Conference on Learning Representations (2022)
42. 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. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/radford21a.html>
43. Reizenstein, J., Shapovalov, R., Henzler, P., Sbordone, L., Labatut, P., Novotny, D.: Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10901–10911 (2021)
44. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
45. Ronneberger, O., Fischer, P., Brox, T.: 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. pp. 234–241. Springer (2015)

46. Seo, J., Jang, W., Kwak, M.S., Ko, J., Kim, H., Kim, J., Kim, J.H., Lee, J., Kim, S.: Let 2d diffusion model know 3d-consistency for robust text-to-3d generation. arXiv preprint arXiv:2303.07937 (2023)
47. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. *Advances in Neural Information Processing Systems* **34**, 6087–6101 (2021)
48. Shi, R., Chen, H., Zhang, Z., Liu, M., Xu, C., Wei, X., Chen, L., Zeng, C., Su, H.: Zero123++: a single image to consistent multi-view diffusion base model. arXiv preprint arXiv:2310.15110 (2023)
49. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023)
50. Shue, J.R., Chan, E.R., Po, R., Ankner, Z., Wu, J., Wetzstein, G.: 3d neural field generation using triplane diffusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20875–20886 (2023)
51. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems* **32** (2019)
52. Talebi, H., Milanfar, P.: Learned perceptual image enhancement. In: *2018 IEEE international conference on computational photography (ICCP)*. pp. 1–13. IEEE (2018)
53. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)
54. Wang, H., Du, X., Li, J., Yeh, R.A., Shakhnarovich, G.: Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 12619–12629 (2023)
55. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems* **34**, 27171–27183 (2021)
56. Wang, Y., Han, Q., Habermann, M., Daniilidis, K., Theobalt, C., Liu, L.: Neus2: Fast learning of neural implicit surfaces for multi-view reconstruction. arXiv preprint arXiv:2212.05231 (2022)
57. 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)
58. Watson, D., Chan, W., Brualla, R.M., Ho, J., Tagliasacchi, A., Norouzi, M.: Novel view synthesis with diffusion models. In: *The Eleventh International Conference on Learning Representations (2023)*, <https://openreview.net/forum?id=HtoA0oT30jC>
59. Xue, Y., Li, Y., Singh, K.K., Lee, Y.J.: Giraffe hd: A high-resolution 3d-aware generative model. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18440–18449 (2022)
60. 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)
61. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4578–4587 (2021)
62. Yu, X., Xu, M., Zhang, Y., Liu, H., Ye, C., Wu, Y., Yan, Z., Zhu, C., Xiong, Z., Liang, T., et al.: Mvimngnet: A large-scale dataset of multi-view images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9150–9161 (2023)

63. Zeng, X., Vahdat, A., Williams, F., Gojcic, Z., Litany, O., Fidler, S., Kreis, K.: Lion: Latent point diffusion models for 3d shape generation. In: Advances in Neural Information Processing Systems (NeurIPS) (2022)
64. Zhang, L., Agrawala, M.: Adding conditional control to text-to-image diffusion models. arXiv preprint arXiv:2302.05543 (2023)
65. Zhang, L.: Reference-only control. <https://github.com/Mikubill/sd-webui-controlnet/discussions/1236> (2023)
66. Zhou, Z., Tulsiani, S.: Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12588–12597 (2023)