

# VaLID: Variable-Length Input Diffusion for Novel View Synthesis

Shijie Li<sup>1,2\*</sup> Farhad G. Zanjani<sup>2</sup> Haitam Ben Yahia<sup>2</sup> Yuki Asano<sup>2</sup> Jürgen Gall<sup>1</sup>  
 Amirhossein Habibian<sup>2</sup>

<sup>1</sup>University of Bonn <sup>2</sup>Qualcomm AI Research<sup>†</sup>

lsj@uni-bonn.de, gall@iai.uni-bonn.de, {fzanjani, hyahia, asano, ahabibia}@qti.qualcomm.com

## Abstract

*Novel View Synthesis (NVS), which tries to produce a realistic image at the target view given source view images and their corresponding poses, is a fundamental problem in 3D Vision. As this task is heavily under-constrained, some recent work, like Zero123 [18], tries to solve this problem with generative modeling, specifically using pre-trained diffusion models. Although this strategy generalizes well to new scenes, compared to neural radiance field-based methods, it offers low levels of flexibility. For example, it can only accept a single-view image as input, despite realistic applications often offering multiple input images. This is because the source-view images and corresponding poses are processed separately and injected into the model at different stages. Thus it is not trivial to generalize the model into multi-view source images, once they are available. To solve this issue, we try to process each pose image pair separately and then fuse them as a unified visual representation which will be injected into the model to guide image synthesis at the target-views. However, inconsistency and computation costs increase as the number of input source-view images increases. To solve these issues, the Multi-view Cross Former module is proposed which maps variable-length input data to fix-size output data. A two-stage training strategy is introduced to further improve the efficiency during training time. Qualitative and quantitative evaluation over multiple datasets demonstrates the effectiveness of the proposed method against previous approaches. The code will be released according to the acceptance.*

## 1. Introduction

Zero-shot novel view synthesis (NVS) has attracted more and more attention recently as the development of AR/VR applications [14] and content creation [25,46] rises.

\*Work was completed during an internship at Qualcomm AI Research.

<sup>†</sup>Qualcomm AI Research is an initiative of Qualcomm Technologies, Inc.



Figure 1. Compared to the previous method (Zero123 [18]), which can only receive a single image as input even when multiple images are available, the proposed architecture (VaLID) can accept variable input views in both training and inference. Thus it can learn a robust representation during training and can utilize more information in the inference time.

NVS tries to synthesize high-quality and visually consistent views given source view images and relative poses. Yet, because of the general under-constrainedness of reconstruction and due to occlusion, it is a highly challenging task that usually involves solving this large ambiguity.

Current solutions can be roughly partitioned into inference-time optimization-based methods and optimization-free methods. The former [22, 24] can produce high-quality outcomes by conducting time-consuming iterative optimization for each object and thus cannot generalize to other objects directly. As a comparison, the latter [18–21, 25, 29, 40, 40, 43, 45] models the whole task by a neural network. Although the generation is a single forward in the inference time which is highly efficient, the quality of generated images is not desirable. Compared to the above methods, Zero123 [18] shows a promising alternative solution. By utilizing the powerful diffusion model, it can produce high-quality images with good generalization ability and efficiency. However, it has been constrained to a single source-view image as input whereas, in the case of the availability of multi-view source images, it cannot be utilized in the process of NVS. The flexibility of a model to a variable number of source-view images is more desirable in real applications where more source images reduce ambiguity through multi-view fusion. Figure 3 (a) shows an overview of the Zero123 [18] architecture. The model adopts an appearance-pose disentanglement conditioning

strategy. The appearance information is injected at the input of U-Net whereas the corresponding relative pose information is injected in the attention modules in the U-Net. Due to separate conditioning of the diffusion process on appearance and pose embeddings, it is non-trivial to generalize the model to multi-view source images, once they are available.

To resolve this restriction and enable the model with the ability to receive variable size source-view images as input, an appearance-pose-entanglement conditioning strategy is proposed which is shown in Figure 3 (b). Based on this conditioning strategy, we propose a novel architecture named **Variable-Length Input Diffusion (VaLID)**. Compared to prior works including Zero123 [18], VaLID can receive variable size source-view images as input to perform NVS, during both training and inference time. This flexibility allows to handle single-view and multi-view input at the same time and to produce even higher fidelity results when more views are available. The proposed VaLID first transfers the source-view image(s) into rich visual embeddings through a Vision Transformer (ViT) encoder, pre-trained with the Masked Auto-Encoder (MAE) task [9]. The ViT encoder produces multiple spatial output tokens. Similar to Zero123 [18], the camera pose(s) are transferred into the pose embedding(s) through an MLP network. The image tokens are concatenated with their corresponding camera pose features to form the input for conditioning the diffusion network.

However, when multiple source-view images are available, there will be an inconsistency among tokens from different images. What is worse, the required computation resource will increase a lot due to the increased number of image tokens from multiple images. To make tokens from different views work harmoniously while efficient, we propose a new “Multi-view Cross Former” module. In this module, all the tokens will be fused and transferred to a fixed number of tokens independent of the number of input source-view images. Finally, these transferred tokens will be injected into the attention modules in the diffusion model to guide the target-view image synthesis.

Although the proposed architecture has the ability to receive multiple source-view images as input, this also requires more training resources. To alleviate this issue, an efficient two-stage training strategy is introduced where the first stage mainly focuses on learning NVS from a single source-view image whereas the second stage mainly solves inconsistency when multiple source-view images are available by only finetuning relevant modules. A token sampling strategy is applied at stage 2 to further constrain training cost.

In summary, our contributions can be summarised as:

- We introduce a diffusion-based NVS model to address variable-size multi-view image fusion, both in train-

ing and inference times. The proposed appearance pose entanglement conditioning strategy can outperform previous methods quantitatively and qualitatively even when only a single source-view image is used.

- We introduce Multi-view Cross Former to transfer the variable-size input tokens into a fixed-size representation by learning a set of learnable tokens. This improves the consistency and efficiency of conditioning to generate novel views while being agnostic to the number of input images.
- With the proposed two-stage training strategy, the training efficiency is improved. What’s more, the performance of single-view image-conditioned NVS is also improved by involving multiple views in the training.

## 2. Related work

Novel View Synthesis is a fast-moving research area with many concurrent works. We broadly categorize the existing and concurrent works into single and multi-view models and highlight the key differences to our work.

### 2.1. Single image novel-view synthesis

Several NVS methods based on single-view input image have been presented on NeRF-based [22] approaches to act on a single view [21, 32, 41, 42, 47] to generate novel views. While these methods produce impressive results, NeRFs famously suffer from requiring test time optimization and the single-view variants typically do not work as well as their multi-view counterparts [47]. To counter the lack of information from a single view, some include additional geometric priors as supervision, input or inductive bias, such as depth maps [8, 13, 36, 41, 42] point clouds [23] or epipolar attention [35]. Our work, in contrast, does not use any geometric representation except for the relative pose. Although these geometric priors work well to aid in reconstructing the observed angle, it does not circumvent that some angles are simply occluded or unobserved. Therefore a large body of work aims at filling these gaps with generative models. Specifically, diffusion models [11, 31] have recently become the primary method to fill in gaps in the observed views as in RealFusion [21], DreamFusion [25] and Zero123 [18]. In particular, Zero-1-to-3 [18] has shown that this is possible with a single view and a relative camera pose. They use CLIP [26] to extract appearance features, combined with a relative pose to generate novel views. Our work is closest to their setup, but we reduce high variance in unobserved regions by extending it to the multi-view setting for training and allowing for single or multi-view inference. This is the *key difference* between single-view methods and ours: we investigate and quantify the effect of increasing




Figure 2. Overview of the proposed architecture (VaLID ).

the number of views, either during training or inference, without any geometric priors. Our work is orthogonal to many of the aforementioned methods and could be further improved as a combination.

## 2.2. Multi-image novel view synthesis

In this section, we look in particular at multi-view works in sparse view regimes. Similar to the single-view setup, there are several geometry-aware methods [10, 12, 16, 32, 37, 44], which try to mitigate the number of views needed for novel view synthesis. Several papers focus on geometry-free NVS, such as ENR [5] that lifts the extracted features of a pair input images into 3D space and minimizes the cross-rendering loss. The introduced 3D lifting is computationally expensive and degrades the quality of rendered images. LFNs [30] samples a single point per ray cast and SRT [28] uses transformers to align multiple frames implicitly. Our work likewise is a geometry-free approach, but we additionally use generative modeling to synthesize novel views. EG3D [1] and 3DiM [39], similarly use generative modeling to align 3D features, using a GAN [6] or diffusion model respectively. 3DiM [39] is closest to our work with its diffusion model, but they only use multiple views during training time and generate poses in an auto-regressive manner at inference that suffers from accumulation of errors in generating and feeding in the additional views and being much slower than our work that consumes all the inputs in a single generation process.

MVDiffusion [33] uses text prompts, depth maps, pixel correspondences, multi-view, and diffusion models on scene reconstruction. Magic3D [15] trains a coarse NeRF with a diffusion model on top for high-resolution NVS based on a text prompt. SyncDreamer [19] does NVS from a single view, by lifting several noisy target images into a frustum and conditioning on this in a pre-trained Zero123. Several works try to enforce consistency across Zero-1-to-3 in different ways [17, 40, 43, 45] *i.e.* by projecting features to 3D [43]. The closest to us is Consistent123 [40], which also uses cross-attention among different views, but the main difference is that we look at single-view output rather than simultaneous multi-view output. Wonder3D [20] estimates

normal maps and textures and combines them using optimization of a neural implicit signed distance field (SDF) to amalgamate all 2D generated data.

Unlike the previously mentioned methods, the proposed VaLID does not require optimization during inference and can accommodate a varying number of source views, making it suitable for practical applications.

## 3. VaLID Novel View Synthesis

We aim for generating a realistic image  $y$  at the target-view, given any number of source-view images  $x = \{x_i\}_{i=1}^{N_v}$ , and corresponding camera poses relative to the target-view  $\pi = \{\pi_i\}_{i=1}^{N_v}$ , where  $N_v \geq 1$  is a variable number of input source-view images. VaLID involves two steps: *i)* Extracting features from the source-view images as described in Sec. 3.1. *ii)* Aggregating features from source-view images to generate the target-view as discussed in Sec. 3.2. A high-level overview of our pipeline is illustrated in Figure 2.

### 3.1. Source view tokenization

**Limitations of existing conditionings** The prior work Zero123 [18] is made of two conditioning mechanisms as illustrated in Fig 3 (a): *i) U-Net conditioning*, where a source-view image  $x_i$  will go through a frozen Auto-Encoder. The output latent map  $f_i^{AE}$  is then concatenated with a noisy latent feature map  $z_t$  from the previous diffusion timestep  $t + 1$  as input to the U-Net. *ii) Attention conditioning*, where the source-view image  $x_i$  will go through a frozen CLIP image encoder. The output CLIP embedding  $f_i^{CLIP}$  is concatenated with the relative pose  $\pi_i$  and then fed into the attention modules in the U-Net. Zero123 [18] can thus be formulated as:

$$z_{t-1} = \epsilon_{\theta} (f_i^{AE} \oplus z_t, f_i^{CLIP} \oplus \pi_i, t), \quad (1)$$

where  $\epsilon_{\theta}$  denotes the U-Net,  $t$  the timestep in the diffusion process and symbol  $\oplus$  denotes the concatenation operator. We found that the Attention conditioning mainly ignores the CLIP features and only propagates pose information as shown in Figure 4. More examples are available in the supplemental materials. We hypothesize that

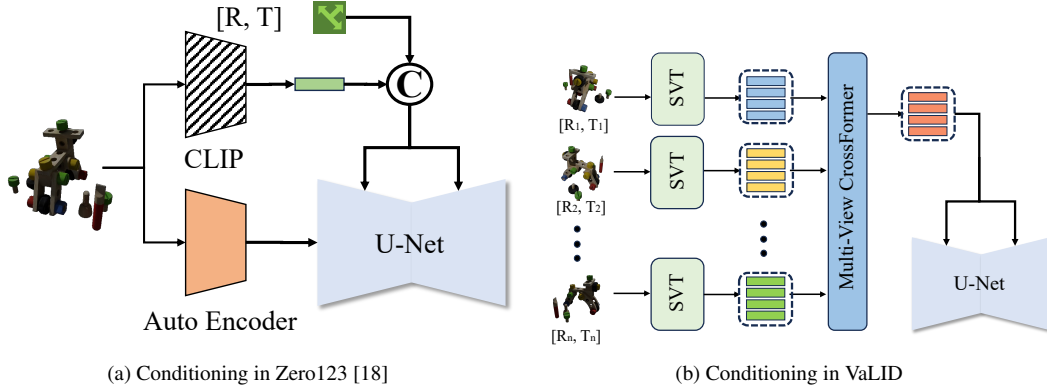


Figure 3. View conditioning in prior work (a) vs. in VaLID (b). By learning a joint appearance pose representation, VaLID seamlessly handles variable length input, which is not possible in a disentangled representation of appearance and pose as in Zero123 [18].

the CLIP encoding, a high-level single token image embedding, is too coarse to retain image details. Thus Eq. (1) reverts to  $z_{t-1} \approx \epsilon_{\theta}(f_i^{AE} \oplus z_t, \pi_i, t)$ , where we see that the extracted features of input source-view image  $f_i^{AE}$  and relative pose  $\pi_i$  are injected into the U-Net at different stages. This works well when only a single source-view image is available as in Zero123 [18]. However, when multiple source-view images are available, this disentangled appearance-pose conditioning strategy would require the model to align the poses with the source-view features at a later stage in the network.

**Learning joint appearance and pose conditioning.** We propose to entangle appearance and pose features together. Since only U-Net conditioning encodes appearance information, one straightforward solution is that we feed both source-view images and corresponding poses to the U-Net input. The problem here is that the number of input views would be inherently tied to the first convolutional layer in the U-Net, making it inflexible. Furthermore, when multiple source-view images are available, feeding them directly into the U-Net will produce inconsistent results, which are shown in Figure 5. Thus we utilize the flexibility of the attention mechanism and opt for feeding both source-view images and corresponding poses to the attention modules in U-Net. In this way, we need to replace the CLIP image encoder by another architecture to better encode source-view images. We achieve this with the SVT Module which can extract fine spatial information from source-view images. Our final architecture can be formulated as:

$$z_{t-1} = \epsilon_{\theta}(z_t, \mathbf{SVT}(x_1, \pi_1, \dots, x_{N_v}, \pi_{N_v}), t). \quad (2)$$

Then the network is trained by optimizing a simplified variational lower-bound

$$\mathcal{L} = \mathbb{E}[||\epsilon_t - z_{t-1}||^2]. \quad (3)$$

where  $\epsilon_t \sim \mathcal{N}(0, 1)$ .




Figure 4. Conditioning limitations in Zero123 [18]: Conditioning is dominated by the U-Net input (UC) whereas CLIP embedding is ignored (CLIP).

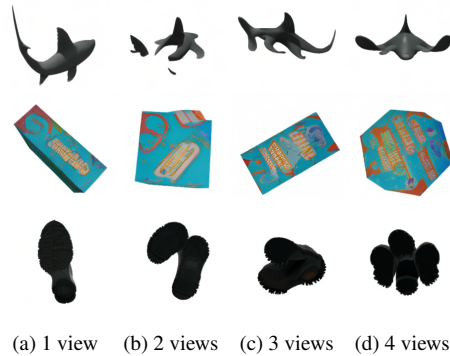


Figure 5. We can observe without the stage 2 training, there will exist heavy inconsistency in the outcomes when we feed multiple source-view images to the proposed architecture.

**Source View Tokenization Module** In the Source View Tokenization (SVT) Module, each input source-view pose image  $(x_i, \pi_i)$  is processed separately and converted into pose-image tokens encoding both appearance and geometry information. We extract features for each input source-view image  $x_i$  with a ViT encoder ( $\Phi$ ), pre-trained with the Masked Auto-Encoder (MAE) task [9]. The input image

will be converted into small image patches  $\mathcal{X} = \{x_j\}_{j=1}^{N_p}$  where  $N_p$  is the number of image patches. Positional embeddings are added to the images as in [3] for spatial awareness. These image patches are then fed into the ViT encoder and converted into image tokens:

$$\mathbf{M} = \Phi(\mathcal{X}), \quad \mathbf{M} = \{m_j\}_{j=1}^{N_p} \quad (4)$$

At this stage, each image token  $m_j$  only contains appearance and 2D spatial information whereas camera pose information is missing. To inject such information, the 3D poses are appended to the end of each token:

$$\tilde{\mathbf{M}} = \{\tilde{m}_j\}_{j=1}^{N_p}, \text{ where } \tilde{m}_j = m_j \oplus \pi^s. \quad (5)$$

where  $\pi^s$  is the corresponding pose to each image token  $m_j$ . For now, each pose-image token  $\tilde{m}_j$  is aware of both appearance information and spatial information (2D spatial and 3D pose). Compared to the CLIP image encoder, the Source View Tokenization can extract rich spatial information and thus better describe image details. The proposed architecture is shown in Figure 2. Later, the extracted tokens from multiple source views can be fused together in our Multi-view Token Cross Former module described in the next Section.

## 3.2. Multi View Fusion

### 3.2.1 Multi-view Cross Former

In the Source View Tokenization Module, we extract pose-image tokens  $\tilde{\mathbf{M}}_i$  for each input source-view image. As mentioned earlier, inconsistent results are produced when providing multi-views directly to the U-Net (Figure 5). Alternatively feeding all source view image tokens will increase the computation proportionally to the number of images, and causes out-of-memory issues. Inspired by text-to-image diffusion models [27], where a fixed number of tokens have been shown to perform favorably, we introduce a fixed number (equal to 64) of learnable seed tokens  $S_0$ , initialized as  $S_0 \sim \mathcal{N}(0, 1)$  at the beginning of training. These are used as queries in the Multi-view Cross Former block, visualized in Figure 6. Since these initial tokens  $S_0$  are shared across training examples, they learn to be query token biases, which aim to extract relevant information from pose-image tokens  $\tilde{\mathbf{M}}_i$ . The Multi-view Cross Former block, in particular, is first computed with an attention operation followed by a residual:

$$S'_l = \text{Attn}_l(Q, K, V) + S_{l-1}. \quad (6)$$

Then the output tokens will pass through a feed-forward layer  $\text{FFW}_l$  which consists of layer normalization, a linear layer, GeLU activation, another linear layer sequentially, followed by a final residual connection:

$$S_l = \text{FFW}_l(S'_l) + S'_l. \quad (7)$$

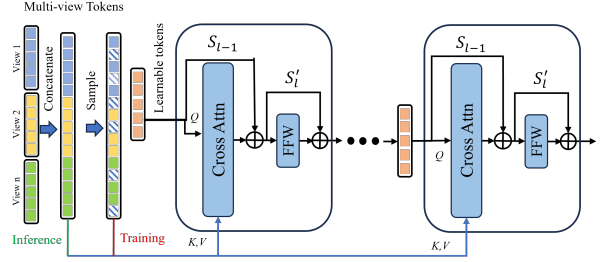


Figure 6. Multi-view Cross Former. During **training** time (stage 2), we randomly sample image tokens from all available source-view images. With this strategy, we can make the Multi-view Cross Former receives multi-view information while restrict training cost. In the **inference** time, we feed all the tokens to include all information for a better performance.

The query, key, and value in the attention block are calculated with  $Q = S_{l-1}$  and  $K = V = \tilde{\mathbf{M}}_1 \oplus \dots \oplus \tilde{\mathbf{M}}_{N_v}$ . The output tokens  $S_l$  will be used as new target-view seeds for the following layer and the whole block will be repeated  $L$  times. Finally, the target-view seeds  $S_L$  will be fed into attention modules at each timestep of the diffusion process to produce realistic images at the target view. With this design, the final target-view seeds  $S_L$  learn to warp the source view tokens into the target-view tokens that are being used for conditioning the diffusion process. This also enables our model to always feed a fixed number of query tokens into the U-Net attention module which improves efficiency a lot.

### 3.2.2 Efficient Training and Inference

As feeding multiple source-view images during training is still expensive, we split the whole training procedure into two steps. In the first stage, the proposed architecture is trained to produce target-view images given a single source-view image. After this stage, our model already has the ability to produce realistic target-view images. As multi-view information aggregation only takes effect in the Multi-view Cross Former, we freeze all modules apart from the Multi-view Cross Former in the second stage of the training. With this strategy in the second stage of training, all the training resources will be used to make multi-view information work harmoniously to produce consistent outcomes. Additionally, we propose a sampling strategy during the second stage training where we feed a variable number of views to the model and randomly sample pose-image tokens from them. By reducing the number of pose-image tokens fed into the Multi-view Cross Former, the training cost is further reduced.

**Stage 1: Single-view Optimization** We build the proposed model on top of a pretrained stable-diffusion model. Thus we can assume that the U-Net has already the ability to produce realistic images. With this assumption, we only

optimize the attention modules in the U-Net. Thus in this stage, the Source View Tokenization Module, Multi-view Cross Former, and attention modules in the U-Net are optimized with a single source-view image as input.

**Stage 2: Variable-view Optimization** After stage 1 training, we only finetune the Multi-view Cross Former where multi-view information fusion takes place, highlighted in Figure 2. In each training iteration, variable input views are fed into the model while only partial pose-image tokens are fed into the Multi-view Cross Former for optimization. With this design, we not only improve the training efficiency due to reducing the number of used pose-image tokens but also improve the robustness of the proposed architecture when only partial information is available. At inference time, we can choose to feed all pose-image tokens to the Multi-view Cross Former to provide more information to the model. We can alternatively sample fewer tokens for efficiency purposes though we find this also hurts performance (shown in Sec. 4.3).

## 4. Experiments

### 4.1. Experimental Setting

**Dataset** We follow Zero123’s [18] experimental setting and fine-tune our diffusion model on the Objaverse dataset [2] which contains more than 800K high-quality 3D object CAD models. For a fair comparison, we use the processed rendering data provided by Zero123 [18] which are 12 random views for each object. For evaluation, we mainly focus on the performance of out-of-distribution data. Thus we conducted the evaluation on the Google Scanned Objects dataset [4] (GSO), which contains high-quality scanned household items, and RTMV [34] dataset, which provides more complex scenes each with around 20 random objects. In both training and evaluation, all the images are resized to  $256 \times 256$  resolution with a white background. To be compliant with the visual transformer architecture in the Source View Tokenization Module, we conduct the center crop thus the final input size will be  $224 \times 224$ .

**Evaluation Setting** On the Google Scanned Objects dataset, 24 views are randomly sampled in 3D as the target view. To demonstrate the performance of our method with variable-length input source-view images, multiple source-view images are rendered following some predefined rules. A reference view (view 0) is randomly picked up first with polar angle  $60^\circ$  for each object. Then 3 other views (view 1,2,3) will be clockwise selected by only varying azimuth (azimuth angle interval as  $90^\circ$ ). In this setting, the available input information is controlled increasing gradually proportional to the number of input source-view images. Because the RTMV dataset already provides rendered images, we conducted the evaluation following the Zero123 [18] experimental setting for fair comparison. To quanti-

tatively evaluate the quality of generated images, we use Peak signal-to-noise ratio (PSNR), Structural Similarity Index (SSIM) [38], and Learned Perceptual Image Patch Similarity (LPIPS) [48] to measure similarity between rendered images and ground truth images.

**Baselines** We mainly compare our method to the previous state-of-the-art method, Zero123 [18] which can produce realistic images at the target view iteratively with a diffusion model. As the authors only released the pretrained weights, we reproduced the experiment results with our own implemented inference code. Although Neural Radiance Field [22] (NeRF) is widely adopted for novel view synthesis, it only works well when a large number of images are available. Hence, to further evaluate the effectiveness of our proposed method, we compare it to DietNeRF [12], a technique that regularizes NeRF using a CLIP image-to-image consistency loss. DietNeRF has demonstrated its capability to produce high-quality results even when only a limited number of images are available. As a comparison, InfoNeRF also targets at scenarios where only limited images are available but does not use any pretrained model. It can produce realistic images at the target view by making reconstructed scenes compact along individual rays and consistent across rays in the neighborhood. This is achieved by the proposed entropy constraint and spatial smoothness constraint. Additionally, single-view 3D reconstruction with the help of a novel view synthesis diffusion model (image-to-3D) attracts more and more attention recently. It optimizes a NeRF with a diffusion model trained for novel view synthesis. We use threestudio [7] which combines Zero123 and NeRF as another baseline method to represent these methods.

### 4.2. Quantitative comparison of methods with multiple source view image as input.

We first compare our method to the baseline methods quantitatively with two different experiment settings, a single source view image as input and multiple source view images ( $N = 4$ ) as input. The former experiments are conducted on both Google Scanned Objects dataset and RTMV dataset whereas the latter experiments are only conducted on Google Scanned Objects dataset. Thus we maintain the same experiment setting for the first experiment to Zero123 for a fair comparison while enable more precise control of source views in the second experiments.

For the experiment results with a single source view image as input, which are shown in Tab. 1, the proposed method can outperforms all previous methods on both GSO and RTMV datasets. On GSO dataset, our method can outperform NeRF-based methods with a large margin (DietNeRF and InfoNeRF). This is because NeRF-based methods usually cannot handle such a severe situation where only a single image is available. As for image-to-3D

	DietNeRF	InfoNeRF	3studio	Z123	VaLID		DietNeRF	InfoNeRF	3studio	Z123	VaLID
PSNR $\uparrow$	10.38	10.29	12.29	19.00	20.03	PSNR $\uparrow$	6.52	6.55	-	8.89	9.03
SSIM $\uparrow$	0.647	0.837	0.800	0.865	0.881	SSIM $\uparrow$	0.439	0.612	-	0.515	0.519
LPIPS $\downarrow$	0.460	0.29	0.289	0.115	0.091	LPIPS $\downarrow$	0.609	0.497	-	0.432	0.432

(a) GSO dataset

(b) RTMV dataset

Table 1. Quantitative comparison of methods with a single source view image as input.

	DietNeRF	InfoNeRF	Z123-C	T-Avg	VaLID	View num	1	2	3	4
PSNR $\uparrow$	11.23	13.04	19.81	15.83	21.31	PSNR $\uparrow$	20.04	20.41	21.05	21.31
SSIM $\uparrow$	0.697	0.802	0.877	0.839	0.895	SSIM $\uparrow$	0.881	0.884	0.891	0.895
LPIPS $\downarrow$	0.409	0.182	0.102	0.182	0.069	LPIPS $\downarrow$	0.091	0.085	0.073	0.069

(a)  $N = 4$  source view images are input

(b) Impact of variable length input views.

Table 2. Quantitative comparison of methods with multiple source view images as input on GSO dataset. We introduce two new baselines to process multiple source-view images. ‘Z123-C’ denotes Zero123 select the closest source view to the target view as input. ‘T-Avg’ denotes apply average pooling to merge the tokens from different source views for the proposed method.

method (3studio), it can outperform NeRF-based methods with the aid of Zero123 which can provide additional supervision at unobserved areas in NeRF optimization. But for novel view synthesis task, its performance is far from the adopted diffusion model (Zero123). We assume this is because Zero123 cannot produce multi-view consistent outcomes, which leads to a blur reconstructed 3D meshes. The proposed method can achieve a 5% improvement on PSNR compared to previous state-of-the-art method Zero123 [18]. The improvement on SSIM and LPIPS are also obvious. As for the RTMV dataset, we follow the Zero123 [18] experiment setting because this dataset already provides rendered images. For each scene, the first view is used as the source view and the following views are used as the target view. Compared to the GSO dataset, this dataset contains more objects for each scene thus more challenging. Even in this situation, our method can still outperform other baseline methods. One thing needs to be noticed is that 3studio cannot produce meaningful results on this dataset with a single source view image as input. This may because Zero123 cannot produce reasonable outcomes where heavily multi-view inconsistency may exist. What’s worse, no advanced regularization techniques are used. Thus it cannot reconstruct the 3D meshes correctly.

For the experiment results with multiple source view images as input ( $N = 4$ ), which are shown in Tab. 2, the proposed method can outperforms all previous methods on GSO dataset. Compared to the experiment results where only a single source view image is available, NeRF-based methods achieve better performance with more available information. Even in this case, their performances are still much lower than the performance of the proposed method with only a single source view image as input. Besides NeRF-based methods, we introduce a new baseline method based on Zero123 [18]. As Zero123 can only receive a single source view image as input, we always select the closest source view to the target view as input. Similarly, although a

better performance is achieved, it is still lower than the performance of the proposed method with only a single source view image as input. For the proposed method, we also try to implemented a token averaging module for multi-view information fusion which resulted in lower PSNR (15.83 vs 21.31 dB). We hypothesize this is due to tokens with the same 2D position being unaligned across views. With all the designs, by increasing the input source-view images, the performances on all metrics increase gradually and achieve the best performance when  $N = 4$  views are available which are shown in Tab. 2b. This demonstrates the capability of the proposed method to utilize information from multiple source-view images when available. Furthermore, the proposed architecture has the flexibility to receive variable-length inputs and thus is more suitable for realistic applications or when the model receives input images progressively.

Finally, we show some qualitative results in Figure 7. Apart from the quality of each individually generated image, the consistency among different target views is also important which is usually missing in previous work. To demonstrate the consistency of the generated images among different target views, we select a predefined camera trajectory to demonstrate the qualitative results. From these qualitative results, we can observe the proposed method not only can produce more realistic images compared to previous methods but also can maintain a better consistency among different target views even no explicit geometry constraints are applied. Especially when multiple source-view images are available.

### 4.3. Ablation Study

In the ablation study, we show the influence of random token sampling strategy (RTS) in both inference and training times. In the proposed method, RTS is only used in training to reduce GPU memory for high training efficiency. Here, we apply RTS for inference stage is to evaluate the ro-




Figure 7. Qualitative results on GSO dataset. We can observe our method can produce high-quality outcomes at the target view compared to the previous state-of-the-art methods. Furthermore, the generated images from our method also maintain better consistency across different views. This is demonstrated by generated images along a predefined camera trajectory. Especially when multiple source-view images are available, the quality of generated images improves a lot. More examples can be found in the supplemental materials.




Figure 8. Token Sampling setting in the stage 2 training (a)-(c). The efficiency analysis of the proposed architecture is shown in (d).

bustness of the proposed model. First, we show the effect of different sampling settings, involved in stage 2 of training in Figure 8 (a)-(c). It can be observed that in stage 2 of training, tokens sampling from two source views (Sample (2 views)) actually achieve quite close performance, compared to feeding all the tokens to Multi-view Cross Former (Token Concat (2 views)). Moreover, the experiments show that providing more source-view images to token sampling improves the quality of the generated images. However, increasing up to six views (Sample (6 views)) doesn't show improvement in comparison to four views (Sample (4 views)). This may be because a fixed number of learnable tokens is applied in Multi-view Cross Former. When the number of input tokens becomes too large, it will face a bottleneck to collect all useful information.

The computational burden of preserving (without sampling) all pose-image tokens for different numbers of source-view images is shown in Figure 8 (d). It can be observed that the introduced Multi-view Cross Former shows a significant reduction in the required Multiply-Accumulate Operations (MACs), especially when the number of source views increased. This also contributes to ignoring a considerable number of pose-image tokens via sampling does not show a performance loss.

More ablation study can be found in the supplement ma-

terial.

## 5. Conclusion

This paper presents a novel diffusion-based framework for novel view image synthesis, named as Variable-Length Input Diffusion model (VaLID). Compared to previous diffusion-based methods, which can only receive a single image as input, VaLID can accept variable number of views as input in both training and inference time. This flexibility makes the proposed model more suitable for realistic applications where multiple views but with a variable number are usually available or progressively presents to model. The multi-view fusion is achieved by an appearance-pose entanglement conditioning strategy. To handle the information inconsistency among different views, a Multi-view Cross Former module has been introduced which is highly efficient, both in terms of reducing the number of tokens, which conditioned the diffusion process, to have a fix-length set and in terms of training strategy. The proposed method outperforms previous methods both quantitatively and qualitative which demonstrates its effectiveness.

## References

- [1] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo,



- Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022.
- [2] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023.
- [3] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [4] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *International Conference on Robotics and Automation*, pages 2553–2560. IEEE, 2022.
- [5] Emilien Dupont, Bautista Miguel Angel, Alex Colburn, Aditya Sankar, Carlos Guestrin, Josh Susskind, and Qi Shan. Equivariant neural rendering. In *International Conference on Machine Learning*, 2020.
- [6] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [7] Yuan-Chen Guo, Ying-Tian Liu, Ruizhi Shao, Christian Laforte, Vikram Voleti, Guan Luo, Chia-Hao Chen, Zi-Xin Zou, Chen Wang, Yan-Pei Cao, and Song-Hai Zhang. threestudio: A unified framework for 3d content generation. <https://github.com/threestudio-project/threestudio>, 2023.
- [8] Yuxuan Han, Ruicheng Wang, and Jiaolong Yang. Single-view view synthesis in the wild with learned adaptive multi-plane images. In *ACM SIGGRAPH*, 2022.
- [9] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.
- [10] Philipp Henzler, Jeremy Reizenstein, Patrick Labatut, Roman Shapovalov, Tobias Ritschel, Andrea Vedaldi, and David Novotný. Unsupervised learning of 3d object categories from videos in the wild. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4698–4707, 2021.
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [12] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *IEEE International Conference on Computer Vision*, pages 5885–5894, 2021.
- [13] Yash Kant, Aliaksandr Siarohin, Michael Vasilkovsky, Riza Alp Guler, Jian Ren, Sergey Tulyakov, and Igor Gilitschenski. Invs : Repurposing diffusion inpainters for novel view synthesis. In *SIGGRAPH Asia 2023 Conference Papers*, 2023.
- [14] Johannes Kopf, Kevin Matzen, Suhil Alsian, Ocean Quigley, Francis Ge, Yangming Chong, Josh Patterson, Jan-Michael Frahm, Shu Wu, Matthew Yu, et al. One shot 3d photography. *ACM Transactions on Graphics (TOG)*, 39(4):76–1, 2020.
- [15] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [16] Kai-En Lin, Lin Yen-Chen, Wei-Sheng Lai, Tsung-Yi Lin, Yi-Chang Shih, and Ravi Ramamoorthi. Vision transformer for nerf-based view synthesis from a single input image. In *WACV*, 2023.
- [17] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. *arXiv preprint arXiv:2309.17261*, 2023.
- [18] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. *arXiv preprint arXiv:2303.11328*, 2023.
- [19] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Learning to generate multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023.
- [20] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, and Wenping Wang. Wonder3d: Single image to 3d using cross-domain diffusion. *arXiv:2310.15008*, 2023.
- [21] Luke Melas-Kyriazi, Christian Ruppert, Iro Laina, and Andrea Vedaldi. Realfusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023.
- [22] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [23] Fangzhou Mu, Jian Wang, Yicheng Wu, and Yin Li. 3d photo stylization: Learning to generate stylized novel views from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16273–16282, 2022.
- [24] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022.
- [25] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *International Conference on Learning Representations*, 2023.

- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [27] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [28] Mehdi S. M. Sajjadi, Henning Meyer, Etienne Pot, Urs Bergmann, Klaus Greff, Noha Radwan, Suhani Vora, Mario Lucic, Daniel Duckworth, Alexey Dosovitskiy, Jakob Uszkoreit, Thomas Funkhouser, and Andrea Tagliasacchi. Scene Representation Transformer: Geometry-Free Novel View Synthesis Through Set-Latent Scene Representations. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2022.
- [29] Yichun Shi, Peng Wang, Jialong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023.
- [30] Vincent Sitzmann, Semon Rezchikov, William T. Freeman, Joshua B. Tenenbaum, and Fredo Durand. Light field networks: Neural scene representations with single-evaluation rendering. In *Neural Information Processing Systems*, 2021.
- [31] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, JMLR Proceedings, pages 2256–2265, 2015.
- [32] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Viewset diffusion: (0-)image-conditioned 3d generative models from 2d data. *International Conference on Computer Vision*, 2023.
- [33] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa. Mvdifffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. In *Neural Information Processing Systems*, 2023.
- [34] Jonathan Tremblay, Moustafa Meshry, Alex Evans, Jan Kautz, Alexander Keller, Sameh Khamis, Thomas Müller, Charles Loop, Nathan Morrical, Koki Nagano, et al. Rtmv: A ray-traced multi-view synthetic dataset for novel view synthesis. *arXiv preprint arXiv:2205.07058*, 2022.
- [35] Hung-Yu Tseng, Qinbo Li, Changil Kim, Suhub Alsian, Jia-Bin Huang, and Johannes Kopf. Consistent view synthesis with pose-guided diffusion models. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2023.
- [36] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 548–557, 2020.
- [37] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021.
- [38] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [39] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. In *International Conference on Learning Representations*, 2023.
- [40] Haohan Weng, Tianyu Yang, Jianan Wang, Yu Li, Tong Zhang, C. L. Philip Chen, and Lei Zhang. Consistent123: Improve consistency for one image to 3d object synthesis. *arXiv preprint arXiv:2310.08092*, 2023.
- [41] DeJia Xu, Yifan Jiang 0001, Peihao Wang, Zhiwen Fan, Yi Wang, and Zhangyang Wang. Neurlift-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*, pages 4479–4489. IEEE, 2023.
- [42] DeJia Xu, Yifan Jiang, Peihao Wang, Zhiwen Fan, Humphrey Shi, and Zhangyang Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXII*, pages 736–753. Springer, 2022.
- [43] Jiayu Yang, Ziang Cheng, Yunfei Duan, Pan Ji, and Hongdong Li. Consistnet: Enforcing 3d consistency for multi-view images diffusion. *arXiv*, 2023.
- [44] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [45] Jialong Ye, Peng Wang, Kejie Li, Yichun Shi, and Heng Wang. Consistent-1-to-3: Consistent image to 3d view synthesis via geometry-aware diffusion models. *arXiv preprint arXiv:2310.03020*, 2023.
- [46] Lin Yen-Chen, Pete Florence, Jonathan T Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. Nerf-supervision: Learning dense object descriptors from neural radiance fields. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 6496–6503. IEEE, 2022.
- [47] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, 2021.
- [48] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.