

MultiDiff: Consistent Novel View Synthesis from a Single Image

Norman Müller¹ Katja Schwarz¹ Barbara Rössle² Lorenzo Porzi¹ Samuel Rota Bulò¹
 Matthias Nießner² Peter Kotschieder¹

Meta Reality Labs Zurich¹ Technical University of Munich²

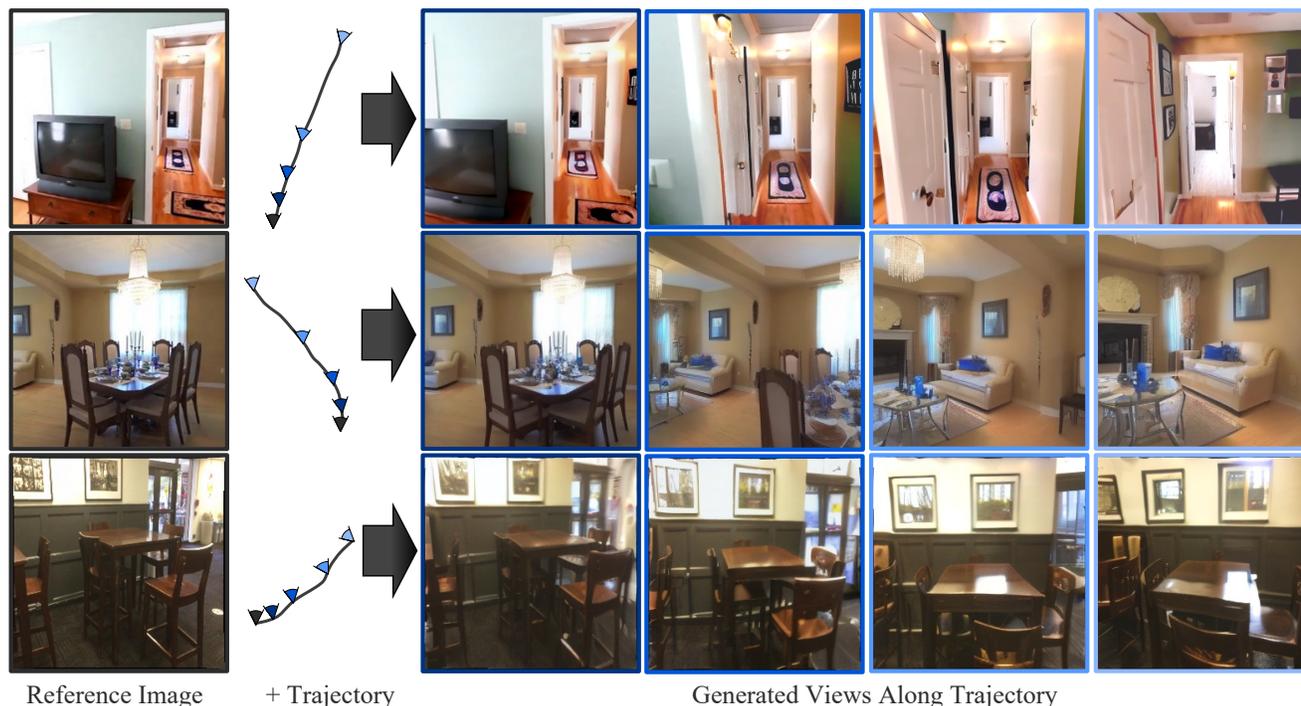


Figure 1. Given a single input image, MultiDiff synthesizes consistent novel views following a desired camera trajectory. These synthesized views harmonize well even in areas unseen from the reference view. Examples from RealEstate10K [81] (top two rows) and ScanNet [9] (bottom row) test sets demonstrate that our model can handle large camera changes and challenging perspectives.

Abstract

We introduce *MultiDiff*, a novel approach for consistent novel view synthesis of scenes from a single RGB image. The task of synthesizing novel views from a single reference image is highly ill-posed by nature, as there exist multiple, plausible explanations for unobserved areas. To address this issue, we incorporate strong priors in form of monocular depth predictors and video-diffusion models. Monocular depth enables us to condition our model on warped reference images for the target views, increasing geometric stability. The video-diffusion prior provides a strong proxy for 3D scenes, allowing the model to learn continuous and pixel-accurate correspondences across generated images. In contrast to approaches relying on autoregressive image generation that are prone to drifts and error accumulation, *MultiDiff* jointly synthesizes a sequence of frames yielding high-quality and multi-view consistent re-

sults – even for long-term scene generation with large camera movements, while reducing inference time by an order of magnitude. For additional consistency and image quality improvements, we introduce a novel, structured noise distribution. Our experimental results demonstrate that *MultiDiff* outperforms state-of-the-art methods on the challenging, real-world datasets *RealEstate10K* and *ScanNet*. Finally, our model naturally supports multi-view consistent editing without the need for further tuning.

1. Introduction

In this work, we address the challenging and highly ill-posed task of view extrapolation from a single image. The goal is to synthesize a set of multiple novel views that are diverse and in themselves consistent. As input, our method only requires a single input image and a user-defined free-

form camera trajectory that may deviate substantially from the reference view. Providing a solution to this problem unlocks applications in virtual & augmented reality and 3D content creation, where generating immersive and multi-view coherent scenes is paramount.

Many existing, state-of-the-art approaches for novel view synthesis are reconstruction-based (*e.g.*, by optimizing a Neural Radiance Field [36] from a fixed number of input views), and are thus inherently limited in generating high-quality novel views for areas without sufficient training coverage. In contrast, we leverage diffusion-based, generative approaches [21, 59, 61–63], that are capable of producing high-quality, single images or individual, simple 3D objects, due to their ability of learning powerful (conditional) image priors. Despite significant progress, these models are still unable to synthesize several, multi-view consistent views of large scenes. This is largely due to the lack of inherent 3D modeling capabilities, the absence of large-scale 3D ground truth datasets, but also the ill-posed nature of the problem, requiring more sophisticated methodological advances. Ultimately, we are aiming for a solution that i) generates seamlessly aligned and multi-view consistent output images w.r.t. a given input image, ii) maintains both high variability and fidelity in occluded regions and previously unseen areas, and iii) extends to camera trajectories well beyond the provided input reference image viewpoint or a simplistic 360° panoramic view.

Some recent works have approached consistent view extrapolation by leveraging an autoregressive approach: *Look Outside the Room* [47] is a transformer-based approach combined with locality constraints w.r.t. the input cameras for enforcing consistency among generated frames. Similarly, *Pose-Guided Diffusion Models* [69] apply attention along epipolar lines to condition a diffusion model. PhotoNVS [78] also proposes an autoregressive attempt where the diffusion model is conditioned on a reference view and a specialized representation for relative camera geometry. A significant drawback of autoregressive models is their tendency to error accumulation [31, 52]. Repeatedly conditioning the model on its previously generated frames can turn minor output deficiencies quickly into undesirable and semantically meaningless results – particularly on longer-term trajectories. In contrast, *Diffusion with Forward Models* [66] (DFM) trains a diffusion model to directly sample from the distribution of 3D scenes, inherently improving 3D consistency. However, DFM is computationally expensive, limited to low image resolutions, slow at inference, and cannot directly integrate 2D diffusion priors. The goal of our work is to overcome both the main limitations of autoregressive works and enabling fast and significantly more stable, long-term generation of novel views.

To this end, we propose MultiDiff, a novel and improved, latent diffusion model-based approach for novel view syn-

thesis, given a single reference image and a pre-defined target camera trajectory as input. We address the challenge of generating pixel-aligned, multi-view consistent image sequences by incorporating strong and complementary priors, significantly constraining the ill-posed nature of the task. Geometric stability is improved by integrating a monocular depth prior, where we condition our model on warped reference images for desired novel views, using off-the-shelf but potentially noisy monocular depth estimators. We also introduce a structured noise distribution for improving multi-view consistency, applying the aforementioned warping procedure to the reference image noise and hence generating correlated 3D noise in all overlapping target views. By integrating a video diffusion model prior, we are able to compensate for missing and geometrically inconsistent reference image warpings due to potential issues with the monocular depth estimator. Video priors provide a strong proxy for 3D scene understanding, enhancing temporal consistency by largely reducing flickering artifacts – particularly for long-trajectory view synthesis. However, their lack of explicit camera control makes their integration nontrivial for view extrapolation.

In order to avoid error propagation issues as observed with autoregressive models, we synthesize entire sequences of novel views in a concurrent and efficient way. Finally, due to our conditioning, we can additionally edit our generated scenes, allowing for direct and intuitive interaction with our model. We summarize our main contributions as follows:

- We address the ill-posed view extrapolation problem by integrating priors from monocular depth estimators and video diffusion models for learning pixel-wise correspondences using novel techniques for spatial-aware conditioning across predicted sequences.
- We simultaneously predict multiple frames for a target sequence, overcoming error accumulation of autoregressive methods, while retaining higher resolution at reduced computational costs compared to methods directly sampling from the distribution of 3D scenes.
- By introducing a novel structured noise distribution, we obtain more multi-view consistent sampling results.

2. Related Works

Image and Video Diffusion. Diffusion Models (DMs) [21, 60, 64] are powerful generative models that have achieved state-of-the-art results in unconditional as well as class- and text-guided image synthesis [2, 11, 12, 16, 18, 23, 39, 40, 44, 46, 50, 53, 71]. Recently, DMs have been extended to the task of video synthesis [8, 14, 22, 24, 34, 56]. While recent video DMs can be conditioned on different modalities such as text or images [8, 17, 73], they do not enable explicit control the camera viewpoint in the generated videos. Nonetheless, the temporal consistency learnt by these models is a powerful

prior that we can leverage to tackle the task of novel view synthesis in an underconstrained setting. Specifically, we use the publicly available VideoCrafter1 [8] to initialize the correspondence attention layers in our pipeline.

Regression-Based Models for Novel View Synthesis.

The goal of novel view synthesis (NVS) is to produce realistic images of a given instance or scene from previously unseen camera viewpoints. Earlier approaches require hundreds of posed training images per instance and optimize each instance individually [10, 33, 35–37, 57, 67]. By learning priors across multiple training scenes, more recent works enable NVS from only one or a few images at inference [7, 13, 19, 29, 41, 42, 49, 54, 58, 68, 72, 77]. These methods optimize a regression objective, i.e. an L1 or L2 loss to reconstruct the training images. While this allows for impressive results on interpolation near input views, regression-based NVS approaches struggle with reconstruction ambiguity and longer-range extrapolations [6]. As our goal is to synthesize novel views far beyond observed views, we instead train a generative model.

Generative Models for Novel View Synthesis. To better model reconstruction ambiguity and long-range view extrapolation, multiple recent works deploy generative models for NVS. Earlier works use GANs [27, 28, 30, 43, 75], VAEs [31], or autoregressive models [47, 48, 51]. Interestingly, GeoGPT [51] directly models long-range 3D correspondences between source and target views with an autoregressive transformer, demonstrating that an intermediate 3D representation may not be needed for NVS from a single image. More recently, diffusion models have achieved impressive results on object-centric data [1, 32, 38, 55, 74, 76, 82]. While these works focus on relatively constrained camera motions around a single object, another line of work addresses scenes with arguably more complex camera trajectories [3, 6, 15, 25, 26, 65, 66, 69, 78]. MVDiffusion [65] performs image synthesis conditioned on depth maps of a given mesh, jointly generating all images of the trajectory. To increase consistency, cross-view interactions are modelled by correspondence-aware attention layers that require given pixel-to-pixel correspondences using GT geometry during training and inference. Our approach instead aims at learning those multi-view correspondences, which allows us to synthesize novel views along a trajectory given just a single RGB image, without the need for any geometric information about the target views. This renders our method applicable to a wider range of scenarios, where no prior 3D reconstruction is available. DFM [66] trains a diffusion model to directly sample from the distribution of 3D scenes. Modeling the scene with a 3D representation is inherently 3D-consistent, but is computationally expensive and in practice limits DFM to lower image resolutions and slow inference. Pose-Guided Diffusion [69] and Pho-

toNVS [78] train a pose-conditioned 2D diffusion model to autoregressively generate frames along a given camera trajectory. However, especially for long trajectories, autoregressive sampling is prone to error accumulation, leading to common struggles with loop closure when taking a Markov assumption and slow inference as it cannot be parallelized. Hence, we do not use autoregressive sampling but generate all images jointly, enabling the model to learn short- and long-term correspondences between views. In stark contrast to MVDiffusion that also performs joint frame synthesis, we only use depth from an off-the-shelf monocular depth estimator with no geometric cues about the target views. Multi-Diff can therefore generate novel views from a single input image only. The learnt correspondence attention enables our model to achieve better consistency than state-of-the-art autoregressive approaches while achieving higher image quality than related works.

3. Method

Given a single *reference* image \mathbf{I}_{ref} , our goal is to generate semantically plausible, consistent novel views along a camera trajectory $\mathcal{C} := \{\mathbf{c}^n\}_{n=1}^N$, where each camera pose \mathbf{c}^n is relative to the camera of the reference image. To this end, we propose a pose-conditional 2D diffusion model with correspondence attention, i.e., attention layers that *jointly* operate on all generated views of the trajectory. A key challenge in novel-view synthesis for the highly under-constrained single-image setting is to achieve consistency in the lack of explicit correspondence supervision. We therefore leverage strong priors that excel at related tasks. Most importantly, we note that the task of video generation is closely related to our problem setting, where temporal consistency is an intrinsic objective.

In the following, we explain how we can integrate and adjust a video prior in conjunction with depth and image priors to enable free viewpoint control. Next, we provide a detailed explanation of our conditioning mechanism and the correspondence attention which adds viewpoint control to our pipeline. Lastly, we introduce structured noise, which ports approximate correspondences between frames to obtain more consistent synthesis results. Our pipeline is illustrated in Figure 2.

Video Prior We build our generative model on top of VideoCrafter [8]. VideoCrafter trains a denoising 3D U-Net in a fixed latent space, using a pretrained image encoder \mathcal{E} and a pretrained image decoder \mathcal{D} to map to and from latent space, respectively. At the core of VideoCrafter is a 3D U-Net with alternating spatial layers and temporal attention. The spatial layers process each frame in a batch individually while the temporal attention operates on all frames jointly. This pretrained 3D U-Net architecture is a well-suited initialization for the task of NVS as the temporal layers already

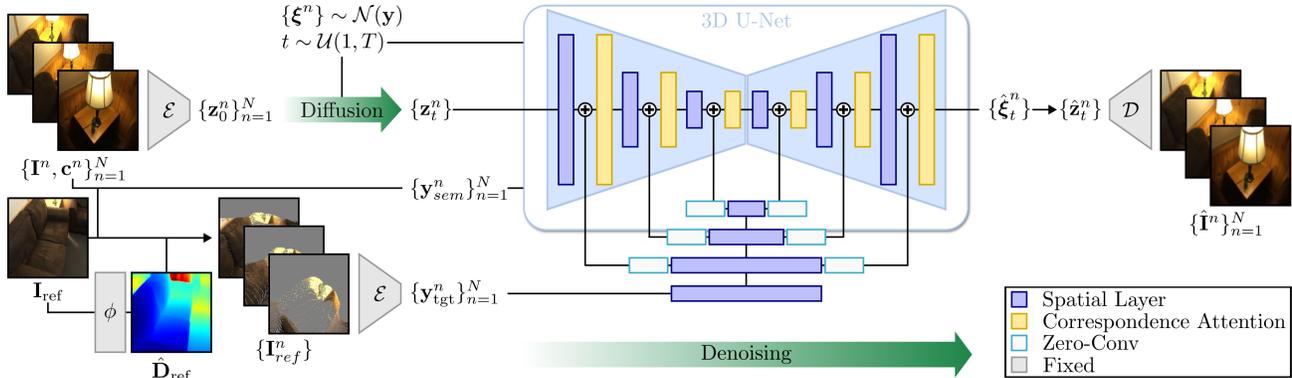


Figure 2. MultiDiff is pose-conditional diffusion model for novel view synthesis from a single image. The diffusion model is trained in the latent space of a fixed auto-encoder with encoder \mathcal{E} and decoder \mathcal{D} and is conditioned on a reference image \mathbf{I}_{ref} and a camera trajectory $\{\mathbf{c}^n\}$. Specifically, we embed N posed target images $\{\mathbf{I}^n\}_{n=1}^N$ into latent space, apply forward diffusion according to a timestep t and structured noise $\{\xi^n\}$, and train a 3D U-Net to predict $\{\xi^n\}$ from the noisy inputs $\{\mathbf{z}_t^n\}$. For each sample n , the U-Net’s prediction $\hat{\xi}_t^n$ is used to reconstruct the denoised sample $\hat{\mathbf{z}}_t^n$ which can then be decoded into the predicted target image $\hat{\mathbf{I}}^n$. We condition the U-Net on the reference image by warping \mathbf{I}_{ref} to the novel views using depth $\hat{\mathbf{D}}_{\text{ref}}$ from a pretrained estimator ϕ . We condition the U-Net on the reference image by warping \mathbf{I}_{ref} to the novel views using depth $\hat{\mathbf{D}}_{\text{ref}}$ from a pretrained estimator ϕ . The warps $\{\mathbf{I}_{\text{ref}}^n\}$ are encoded into latent representations $\{\mathbf{y}_{\text{tgt}}^n\}$ and injected into the U-Net in a ControlNet inspired manner. We further condition the model directly on the camera pose and an embedding of the reference image as part of the semantic condition $\{\mathbf{y}_{\text{sem}}^n\}$.

provide an inductive bias towards (temporal) consistency. During training, we nevertheless finetune all layers of the U-Net for the novel view synthesis task, where instead of ensuring temporal consistency, the attention layers should establish correspondences between multiple views. Hence, we refer to this type of attention as *correspondence attention*.

Novel-view synthesis In order to generate novel views that adhere to the given camera trajectory \mathcal{C} , we need to condition our pipeline on the target camera poses $\mathbf{c}^n \in \mathcal{C}$. A naïve approach to integrating control over the camera viewpoint is to directly condition the 3D U-Net on \mathbf{c}^n , e.g., via cross-attention. In practice, we concatenate them to the semantic condition of VideoCrafter that consists of an embedding of the reference image and the framerate of the input sequence yielding the semantic conditioning \mathbf{y}_{sem} .

However, this form of guidance alone is too weak to deliver satisfactory novel-view synthesis results (see Sec. 4.2). We therefore integrate a monocular depth prior in order to constrain the highly ill-posed nature of the task. In our experiments, we use ZoeDepth [4] pretrained on ScanNet [9] and refer to the supplementary material for ablations about alternative monocular estimators. We use the depth \mathbf{D}_{ref} estimated from the reference image \mathbf{I}_{ref} to implement a warping function Ψ^n that enables warping images from the camera of the reference image to any other camera $\mathbf{c}^n \in \mathcal{C}$. We denote by $\mathbf{I}_{\text{ref}}^n := \Psi_n(\mathbf{I}_{\text{ref}})$ the reference image warped to camera \mathbf{c}^n and by $\mathbf{M}^n := \Psi_n(\mathbf{1})$ the mask indicating the area of valid warped pixels in camera \mathbf{c}^n .

To facilitate learning the 3D correspondences across the spatial features, for each view n , we encode $\mathbf{I}_{\text{ref}}^n$ into latent space via \mathcal{E} and stack the mask \mathbf{M}^n , suitably resized, along

the channel dimension. The resulting tensor is denoted \mathbf{y}_{tgt} . Inspired by ControlNet [79], we create a copy of the down-sampling layers of diffusion U-Net to extract features from \mathbf{y}_{tgt} , but we prepend a convolutional layer to cope with the additional mask channel.

The intermediate feature maps are then processed with zero-initialized convolutions and added to the outputs of all spatial layers of the 3D U-Net. Note that this differs from the procedure proposed in ControlNet, which only inserts the feature maps into the decoder. We further do not freeze the layers of VideoCrafter to enable learning the correspondence attention. In initial experiments we found that finetuning all layers jointly results in better performance than using a fixed video prior.

The warping operation is implemented by leveraging an off-the-shelf monocular depth estimator and thus error-prone and incomplete. By also passing the reference image and camera poses to the network in the semantic conditioning \mathbf{y}_{sem} , we enable our approach to follow the provided trajectory even in absence of overlap with the reference image. We refer to our ablation Sec. 4.2 for a discussion about the importance of the individual design decisions. In the rest of the section we summarize with \mathbf{y} all quantities we condition our model on, namely reference image \mathbf{I}_{ref} , camera trajectories \mathcal{C} , and all derived ones (estimated depth, warped reference images, corresponding masks).

Structured noise distribution $\mathcal{N}(\mathbf{y})$. Images of a 3D scene captured from different point of views exhibit strong correlations. Hence, it is beneficial to inject similar correlations in the noise ϵ that is used by our diffusion model to synthesize the different camera views, which would otherwise be a standard normal multi-variate. This helps en-

forcing more consistent outputs [45]. Since the correlations are mainly driven by geometric constraints, we leverage the warping function Ψ_n introduced in the previous paragraph to warp a standard normal multi-variate ϵ^0 to all other camera views in \mathcal{C} , while filling the gaps with independent Gaussian noise. This yields per-view noise $\xi^n := \mathbf{M}^n \odot \Psi_n(\epsilon^0) + (1 - \mathbf{M}^n) \odot \epsilon^n$, where ϵ^n is a standard normal multi-variate and \mathbf{M}^n is the suitably-resized warp-validity mask. This process yields $\xi := (\xi^1, \dots, \xi^N)$, which is regarded as a sample of the structured noise distribution we denoted by $\mathcal{N}(\mathbf{y})$.

Training Objective. Let $V := \{(\mathbf{I}^0, \mathbf{c}^0), \dots, (\mathbf{I}^N, \mathbf{c}^N)\}$ be a ground-truth, posed video sequence, where \mathbf{I}^n and \mathbf{c}^n are the n th image and camera pose in the sequence, respectively. We assume \mathbf{I}^0 to be the reference image, *i.e.* $\mathbf{I}_{\text{ref}} := \mathbf{I}^0$, and assume all cameras to be relative to \mathbf{c}^0 . We encode all target images of the sequence into a joint latent representation $\mathbf{z} := (\mathbf{z}^1, \dots, \mathbf{z}^N)$, where $\mathbf{z}^n := \mathcal{E}(\mathbf{I}^n)$, and \mathbf{y} is the conditioning information encompassing the encoded reference image, camera poses and warped reference images described earlier in the section. The denoising training objective takes the following form for the training example V :

$$\mathcal{L}(\theta; V) := \mathbb{E}_{\substack{\xi \sim \mathcal{N}(\mathbf{y}) \\ t \sim \mathcal{U}(1, T)}} \left[\|\xi - \varepsilon_\theta(\mathbf{z} \oplus_t \xi; \mathbf{y}, t)\|^2 \right], \quad (1)$$

where t is sampled from a uniform distribution $\mathcal{U}(1, T)$ and ξ is noise sampled from the structured noise distribution $\mathcal{N}(\mathbf{y})$. The term $\mathbf{z} \oplus_t \xi := \sqrt{\alpha_t} \mathbf{z} + \sqrt{1 - \alpha_t} \xi$ perturbs \mathbf{z} with noise ξ according to a variance-preserving formulation with parameters α_t , from which ε_θ , *i.e.* our denoising 3D U-Net with weights θ , is required to recover ξ . Our model is optimized using Adam by minimizing the training loss function averaged over random batches of video sequences sampled from a given dataset.

Inference. At inference time, we assume to be given a reference image \mathbf{I}_{ref} and a sequence of cameras \mathcal{C} relative to it, which we use to compute the conditions \mathbf{y} . We generate a video sequence from our model by using the DDIM schedule [61], *i.e.* starting from $\mathbf{z}_T \sim \mathcal{N}(\mathbf{y})$ we iterate the following equation

$$\begin{aligned} \mathbf{z}_{t-1} &:= \sqrt{\alpha_{t-1}}(\mathbf{z}_t \ominus_t \varepsilon_\theta(\mathbf{z}_t; \mathbf{y}, t)) \\ &\quad + \sqrt{1 - \alpha_{t-1}} \varepsilon_\theta(\mathbf{z}_t; \mathbf{y}, t), \end{aligned} \quad (2)$$

until we obtain \mathbf{z}_0 by setting $\alpha_0 := 1$. The term $\mathbf{z}_t \ominus_t \epsilon := \frac{\mathbf{z}_t - \sqrt{1 - \alpha_t} \epsilon}{\sqrt{\alpha_t}}$ recovers \mathbf{z} from \mathbf{z}_t assuming noise ϵ . The final result \mathbf{z}_0 entails the synthesized views in latent space for all cameras in \mathcal{C} , from which we compute the counterparts in pixel space by applying the decoder \mathcal{D} . Note that MultiDiff can generate all images of the sequence simultaneously. However, sometimes the novel view has little or no overlap

	Method	Short-term				Long-term			
		PSNR \uparrow	LPIPS \downarrow	FID \downarrow	KID \downarrow	FID \downarrow	KID \downarrow	FVD \downarrow	mTSED \uparrow
128px	DFM [66]	18.10	0.299	36.37	0.010	31.20	0.007	120.2	0.972
	Text2Room [27]	15.45	0.370	34.41	0.008	84.10	0.048	163.4	0.932
	PhotoNVS [78]	15.66	0.376	26.39	0.006	42.99	0.016	117.5	0.907
	MultiDiff (Ours) w/o SN	16.21	0.335	27.26	0.005	30.28	0.006	107.5	0.936
	MultiDiff (Ours)	16.41	0.318	25.30	0.003	28.25	0.004	94.37	0.941
256px	Text2Room [27]	14.88	0.458	35.41	0.009	91.92	0.050	178.6	0.837
	PhotoNVS [78]	15.01	0.452	26.75	0.005	45.08	0.017	130.4	0.801
	MultiDiff (Ours) w/o SN	15.55	0.412	29.49	0.008	33.71	0.010	116.4	0.849
	MultiDiff (Ours)	15.65	0.393	25.90	0.004	30.15	0.006	105.9	0.855

Table 1. Quantitative comparison on RealEstate10K [81] test sequences. Our model achieves higher image quality than state-of-the-art baselines and comparable consistency compared to DFM.

with the reference image, making the warped reference image, *i.e.* condition $\mathbf{y}_{\text{tgt}}^n$ less informative. To further refine the results, we can run the sampling again on the generated sequence, but now use the warp of the closest generated image in $\mathbf{y}_{\text{tgt}}^n$ which in practice this slightly improves consistency.

4. Experiments

In this section, we evaluate the performance of our method for the task of consistent novel view synthesis from a single reference image.

Datasets We compare our methods against state-of-the-art approaches on RealEstate10K [81] and ScanNet [9]. Both datasets provide video sequences together with registered camera parameters. RealEstate10K is a large dataset of real estate recordings gathered from YouTube. The clips typically feature smooth camera movement with little to no camera roll or pitch. Most frames further show considerable coverage of the respective rooms. Following previous works [47, 78], we center-crop and downsample the videos to 256px resolution. ScanNet consists of 1513 handheld captures of indoor environments. The camera trajectories follow a scan-pattern which can contain rapid changes and variation of camera orientation. The resulting frames encompass close-up object captures as well as wide room recordings, leading to heavy occlusions and an overall diverse data distribution. The aforementioned features make ScanNet extremely challenging for novel view synthesis from a single image and our evaluations in Sec. 4 indicate that additional priors are very beneficial in this setting. We resize the images to 256×256 and remark that ScanNet contains 3D meshes that we use for MVDiffusion as it requires predefined correspondences between frames.

Evaluations We evaluate our approach in terms of image fidelity and consistency of the generated outputs. Similar to [47], we consider both short-term and long-term view synthesis. Specifically, we randomly select 1k sequences with 200 frames from the test set and evaluate the 50th generated frame for short-term and the 200th generated frame for long-term view synthesis for RealEstate10K. Due to the faster camera motion, on ScanNet instead we choose the

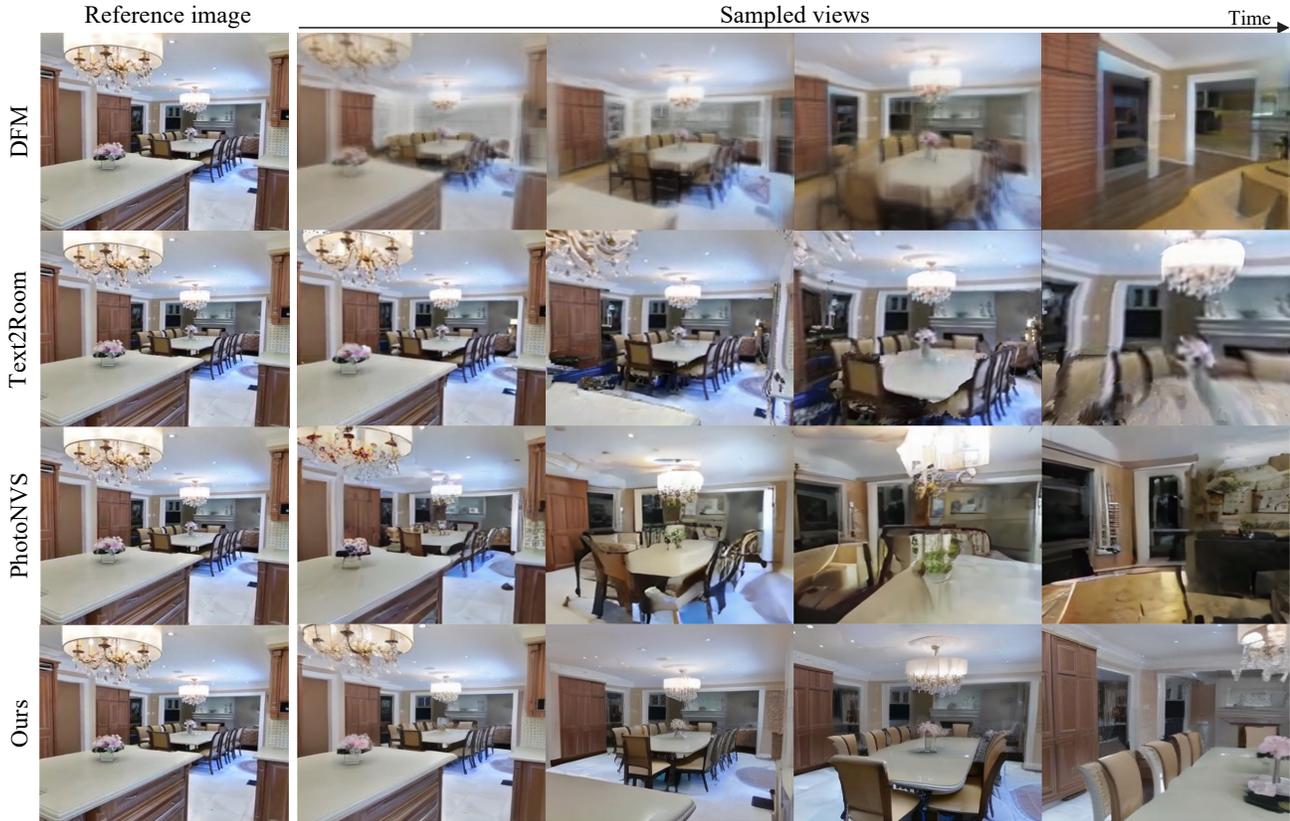


Figure 3. Novel views following ground-truth trajectories (right) given a reference view (left) on RealEstate10K. Through our joint multi-frame prediction combined with effective priors and conditioning, our sequence of novel views is highly realistic and view-consistent compared to the baselines, which show severe degradation over time.

	Method	Short-term				Long-term			
		PSNR \uparrow	LPIPS \downarrow	FID \downarrow	KID \downarrow	FID \downarrow	KID \downarrow	FVD \downarrow	mTSED \uparrow
128px	MVDiffusion [65]	13.14	0.439	43.28	0.013	43.58	0.013	186.6	0.506
	DFM [66]	16.59	0.444	75.19	0.036	111.9	0.069	167.2	0.912
	Text2Room [27]	15.01	0.452	39.87	0.008	82.44	0.0041	173.1	0.812
	PhotoNVS [78]	15.23	0.440	49.19	0.019	75.23	0.038	89.04	0.479
	MultiDiff (Ours) w/o SN	15.29	0.372	40.36	0.008	43.61	0.011	80.71	0.752
	MultiDiff (Ours)	15.50	0.356	38.44	0.007	42.41	0.010	74.10	0.776
256px	MVDiffusion [65]	12.88	0.502	50.18	0.017	51.60	0.018	230.1	0.361
	Text2Room [27]	14.32	0.514	46.69	0.014	93.09	0.058	201.1	0.631
	PhotoNVS [78]	14.61	0.542	63.21	0.033	96.85	0.059	134.2	0.263
	MultiDiff (Ours) w/o SN	14.80	0.445	47.10	0.013	50.84	0.016	119.3	0.529
	MultiDiff (Ours)	15.00	0.431	43.84	0.010	47.11	0.013	114.9	0.576

Table 2. Quantitative comparison on ScanNet [9] test sequences. Our approach outperform all baselines at 256px resolution and shows significantly higher image fidelity compared to DFM.

25th frame for short-term and 100th for long-term evaluation. In the short-term setting, we report Peak Signal-to-Noise Ratio (PSNR) and perceptual similarity (LPIPS) [80] as standard metrics for novel view synthesis. To evaluate the extrapolation capacities in regard of image fidelity, we evaluate Fréchet Inception Distance [20] (FID) and Kernel Inception Distance [5] (KID) for long-term settings. To measure the video-consistency of the generated trajectory images, we compute Fréchet Video Distance (FVD) [70] scores. Further, we follow [78] and report the symmetric epipolar distance (SED) to quantify faithfulness with respect to the provided camera trajectory, i.e., relative pose

accuracy. Here, we compute the mean thresholded symmetric epipolar distance (mTSED) over the pixel thresholds [1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0] and refer to the supplementary for detailed results.

Baselines We compare our approach against the state-of-the-art approaches for scene synthesis from a single reference image, including DFM [66], PhotoNVS [78], Text2Room [25] and MVDiffusion [65]. As MVDiffusion is purely text-conditional, we incorporate the reference image during inference as follows. We use DDIM inversion to obtain the noise corresponding to the reference image and include it in the batch during sampling. Due to the global awareness of MVDiffusion, information from the reference image can propagate to all generated views. Please see supplementary material for more details. DFM trains a diffusion model to directly sample from the distribution of 3D scenes. Unlike our approach, DFM cannot directly integrate 2D diffusion priors and does not generalize well to out-of-domain inputs as our experiments on ScanNet indicate. PhotoNVS [78] trains a pose-conditioned 2D diffusion model to iteratively predict the next frame for a given camera trajectory. Text2Room [25] uses an autoregressive approach of predicting depth and leveraging a

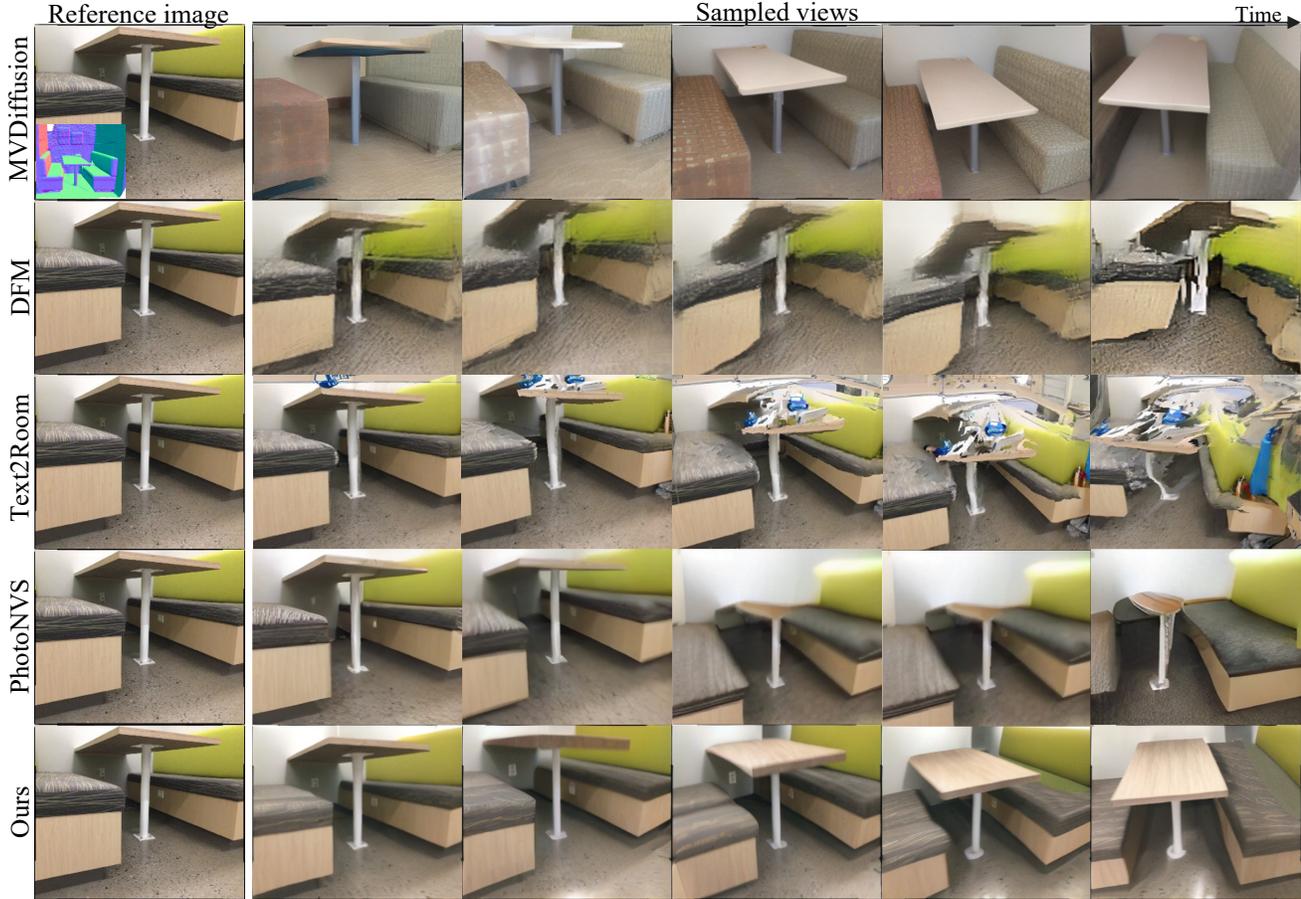


Figure 4. Generated views along ScanNet [9] test sequence (right) given a reference view (left). Our method simultaneously generates sequences of novel views that are both more realistic and more view-consistent than the baselines, DFM and PhotoNVS, which suffer from a considerable performance drop across large view point changes. Although MVDiffusion uses sensor depth input, the generated views are much less consistent with the reference image (e.g., colors of the cushions), compared to our generations, which do not rely on sensor depth.

depth-conditional T2I model to generate new views that are used to update a textured mesh. In contrast, MultiDiff generates multiple frames from the input image in parallel, resulting in better long-term view synthesis and faster inference: To synthesize a 128×128 frame, PhotoNVS requires ≈ 45 s, DFM ≈ 17 s while ours only takes ≈ 1 s.

4.1. Consistent Novel View Synthesis

Comparison against state of the art We quantitatively evaluate our approach on the task of consistent novel-view synthesis from a single reference on RealEstate10K [81] in Tab. 1 and ScanNet in Tab. 2. Since DFM does not support higher resolutions than 128px due to memory limitations, whereas the other methods run at a default resolution of 256px, we perform separate analyses at both resolutions. On RealEstate10K, we observe that our method achieves consistently better FID and KID scores on both short-term, as well as long-term evaluations: The short-term FID compared to DFM improves from 36.37 to 25.30 (at 128px), while the long-term FID improves by 33% compared to

PhotoNVS. Moreover, our model outperforms all baselines in terms of FVD and achieves comparable results on LPIPS and mTSED with respect to DFM. We note that the PixelNeRF [77] representation of DFM leads to highly consistent results, therefore good scores on pixel-level metrics like short-term PSNR, however, this comes at the cost of sharpness (as reflected in FID/KID).

By leveraging strong image- and video-diffusion priors, our method achieves clear improvements over the baselines on ScanNet: As shown in Tab. 2, MultiDiff outperforms MVDiffusion on short- and long-term metrics, indicating our model’s ability to learn long-term correspondences even without relying on ground-truth geometry. In comparison to DFM, Text2Room and PhotoNVS, we observe strong photometric short- and long-term improvements over all baselines. Figs. 3 and 4 show qualitative comparisons on RealEstate10K and ScanNet, respectively. It stands out that our method synthesises realistic and consistent novel views even across large viewpoint changes, where the quality of the baselines drops noticeable.

Method	Short-term			Long-term		
	PSNR \uparrow	LPIPS \downarrow	FID \downarrow	FID \downarrow	FVD \downarrow	mTSED \uparrow
MultiDiff no prior	14.29	0.493	63.56	85.30	236.8	0.587
MultiDiff no vid.	14.68	0.552	37.05	38.43	214.9	0.728
MultiDiff no warp	13.65	0.557	47.42	58.30	181.1	0.484
MultiDiff no pose	15.53	0.417	27.84	32.25	120.26	0.624
MultiDiff (Ours)	15.65	0.393	25.90	30.15	105.9	0.855

Table 3. Ablation of individual components of our pipeline on RealEstate10K [81] test sequences at 256px resolution.

4.2. Ablations

We show the contributions of individual components of our approach in Tab. 3 and refer to the supplementary material for more qualitative comparisons.

Importance of priors As described in Sec. 3, we initialize our model with weights obtained by training on large-scale image and video datasets. To study the importance of those priors for the task of consistent novel-view synthesis from a single image, we ablate them one by one: As shown in Tab. 3, training from scratch (“MultiDiff no prior”) leads to strong degradation of image quality, as well as overall consistency. Removing the video diffusion prior (“MultiDiff no vid.”) has strong influence on the long-term consistency (mTSED decreases by 12.7%), as well as the video quality (FVD increases by more than 120%). We further ablate the monocular depth estimates on the reference image as condition to our model in “MultiDiff no warp” (Tab. 3). The drop in mTSED from 85.5% to 48.4% indicates that the model without reference warps is not able to closely adhere to the input trajectory. Besides the depth-warpings of the reference image, our method uses relative camera poses to synthesize images from the desired target poses. When removing this modality (“MultiDiff no pose”), we notice effect on long-term generation becomes apparent, where there is minimal to no warp-guidance to inform about the desired camera poses, hence mTSED decreases from 0.85 to 0.62.

Importance of structured noise As described in Sec. 3, we introduce structured noise by warping the initial noise consistently between target views according to the depth estimates of the reference image. We measure the effect of noise warping in Tab. 1 and Tab. 2 (“MultiDiff w/o SN”) on RealEstate10K and ScanNet trajectories. On both datasets, we observe that the structured noise leads to significantly more consistent and higher quality synthesis results. We show the effect of noise-warping in Fig. 5.

4.3. Consistent Editing

In contrast to existing works such as DFM or PhotoNVS, our approach directly supports consistent editing without task-specific training. During training, our model is tasked to synthesize consistent novel views even in absence of meaningful reference warps. By masking an area in a reference image that should not be warped, our model naturally



Figure 5. Without structured noise (“MultiDiff w/o SN”), the color of the dining table is not maintained w.r.t. the reference image.

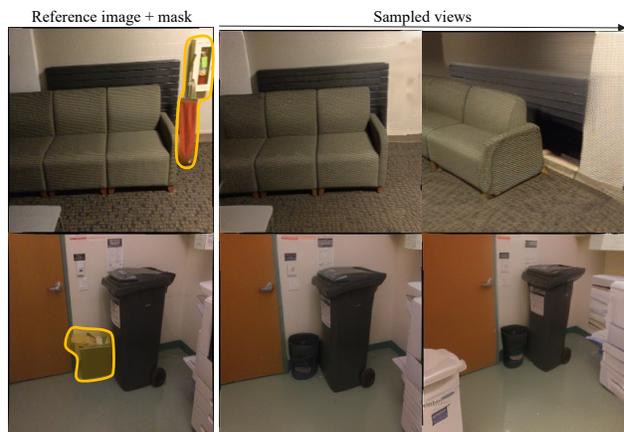


Figure 6. Consistent masking-based editing results on ScanNet test images.

performs consistent completion in those regions. We show examples on ScanNet test images in Fig. 6 and refer to the supplementary material for more qualitative results.

5. Conclusions

In this paper, we introduce MultiDiff, a novel approach for view extrapolation from a single input image. We identify video priors as a powerful proxy for this setting and demonstrate how they can be incorporated and adapted by converting temporal attention to *correspondence attention*. With monocular depth cues, we facilitate learning improved correspondences by conditioning our model on reference views warped w.r.t. the target camera trajectory. Our experiments on RealEstate10k and ScanNet show significant improvements over relevant baselines, with particular gains on long-term sequence generation and overall inference speed.

Acknowledgements

Matthias Nießner was supported by the ERC Starting Grant Scan2CAD (804724).

References

- [1] Titas Anciukevicius, Zexiang Xu, Matthew Fisher, Paul Henderson, Hakan Bilen, Niloy J. Mitra, and Paul Guerrero. RenderDiffusion: Image diffusion for 3D reconstruction, inpainting and generation. *arXiv*, 2022. 3
- [2] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. ediffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 2
- [3] Miguel Ángel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander T Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, Afshin Dehghan, and Joshua M. Susskind. GAUDI: A neural architect for immersive 3D scene generation. *arXiv preprint arXiv:2207.13751*, 2022. 3
- [4] Shariq Farooq Bhat, Reiner Birkl, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth, 2023. 4
- [5] Mikołaj Bińkowski, Dougal J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *International Conference on Learning Representations*, 2018. 6
- [6] Eric R. Chan, Koki Nagano, Matthew A. Chan, Alexander W. Bergman, Jeong Joon Park, Axel Levy, Miika Aittala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. GeNVS: Generative novel view synthesis with 3D-aware diffusion models. In *CVPR*, 2023. 3
- [7] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *CVPR*, 2021. 3
- [8] VideoCrafter contributors. Videocrafter. *Github*. Accessed October 15, 2023 [Online] <https://github.com/AI-Lab-CVC/VideoCrafter>. 2, 3
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 1, 4, 5, 6, 7
- [10] Peng Dai, Yinda Zhang, Zhuwen Li, Shuaicheng Liu, and Bing Zeng. Neural point cloud rendering via multi-plane projection. In *CVPR*, 2020. 3
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 2
- [12] Tim Dockhorn, Arash Vahdat, and Karsten Kreis. Score-based generative modeling with critically-damped langevin diffusion. In *ICLR*, 2022. 2
- [13] Yilun Du, Cameron Smith, Ayush Tewari, and Vincent Sitzmann. Learning to render novel views from wide-baseline stereo pairs. In *CVPR*, 2023. 3
- [14] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. In *ICCV*, 2023. 2
- [15] Rafail Fridman, Amit Abecasis, Yoni Kasten, and Tali Dekel. Scenescape: Text-driven consistent scene generation. In *NeurIPS*, 2023. 3
- [16] Shuyang Gu, Dong Chen, Jianmin Bao, Fang Wen, Bo Zhang, Dongdong Chen, Lu Yuan, and Baining Guo. Vector quantized diffusion model for text-to-image synthesis. In *CVPR*, 2022. 2
- [17] Xianfan Gu, Chuan Wen, Jiaming Song, and Yang Gao. Seer: Language instructed video prediction with latent diffusion models. *arXiv preprint arXiv:2303.14897*, 2023. 2
- [18] Yingqing He, Shaoshu Yang, Haoxin Chen, Xiaodong Cun, Menghan Xia, Yong Zhang, Xintao Wang, Ran He, Qifeng Chen, and Ying Shan. Scalecrafter: Tuning-free higher-resolution visual generation with diffusion models. *arXiv preprint arXiv:2310.07702*, 2023. 2
- [19] Philipp Henzler, Jeremy Reizenstein, Patrick Labatut, Roman Shapovalov, Tobias Ritschel, Andrea Vedaldi, and David Novotny. Unsupervised learning of 3d object categories from videos in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4709, 2021. 3
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. 6
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 2
- [22] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 2
- [23] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. 23:47–1, 2022. 2
- [24] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *ICLR*, 2022. 2
- [25] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *ICCV*, 2023. 3, 6
- [26] Seung Wook Kim, Bradley Brown, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchical latent diffusion models. In *CVPR*, 2023. 3
- [27] Jing Yu Koh, Honglak Lee, Yinfei Yang, Jason Baldrige, and Peter Anderson. Pathdreamer: A world model for indoor navigation. In *ICCV*, 2021. 3
- [28] Jing Yu Koh, Harsh Agrawal, Dhruv Batra, Richard Tucker, Austin Waters, Honglak Lee, Yinfei Yang, Jason Baldrige, and Peter Anderson. Simple and effective synthesis of indoor 3D scenes. *AAAI*, 2023. 3
- [29] Jonás Kulhánek, Erik Derner, Torsten Sattler, and Robert Babuska. Viewformer: Nerf-free neural rendering from few images using transformers. In *ECCV*, 2022. 3

- [30] Zhengqi Li, Qianqian Wang, Noah Snavely, and Angjoo Kanazawa. Infinitenature-zero: Learning perpetual view generation of natural scenes from single images. In *ECCV*, 2022. 3
- [31] Andrew Liu, Richard Tucker, Varun Jampani, Ameesh Makadia, Noah Snavely, and Angjoo Kanazawa. Infinite nature: Perpetual view generation of natural scenes from a single image. In *ICCV*, 2021. 2, 3
- [32] 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. In *ICCV*, 2023. 3
- [33] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 3
- [34] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *CVPR*, 2023. 2
- [35] Moustafa Meshry, Dan B. Goldman, Sameh Khamis, Hugues Hoppe, Rohit Pandey, Noah Snavely, and Ricardo Martin-Brualla. Neural rerendering in the wild. In *CVPR*, 2019. 3
- [36] 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. In *European conference on computer vision*, pages 405–421. Springer, 2020. 2
- [37] Norman Müller, Andrea Simonelli, Lorenzo Porzi, Samuel Rota Buló, Matthias Nießner, and Peter Kotschieder. Autorf: Learning 3d object radiance fields from single view observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [38] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Peter Kotschieder, and Matthias Nießner. Diffri: Rendering-guided 3d radiance field diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4328–4338, 2023. 3
- [39] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. 2
- [40] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. 2022. 2
- [41] Michael Niemeyer, Lars M. Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 3
- [42] Michael Niemeyer, Jonathan T. Barron, Ben Mildenhall, Mehdi S. M. Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. 3
- [43] Dario Pavllo, David Joseph Tan, Marie-Julie Rakotosaona, and Federico Tombari. Shape, pose, and appearance from a single image via bootstrapped radiance field inversion. In *CVPR*, 2023. 3
- [44] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: improving latent diffusion models for high-resolution image synthesis. *CoRR*, arxiv preprint arxiv:2307.01952, 2023. 2
- [45] Haonan Qiu, Menghan Xia, Yong Zhang, Yingqing He, Xintao Wang, Ying Shan, and Ziwei Liu. Freenoise: Tuning-free longer video diffusion via noise rescheduling, 2023. 5
- [46] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [47] Xuanchi Ren and Xiaolong Wang. Look outside the room: Synthesizing a consistent long-term 3D scene video from a single image. In *CVPR*, 2022. 2, 3, 5
- [48] Chris Rockwell, David F Fouhey, and Justin Johnson. Pixel-synth: Generating a 3D-consistent experience from a single image. In *ICCV*, 2021. 3
- [49] Barbara Roessle, Jonathan T. Barron, Ben Mildenhall, Pratul P. Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [50] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2
- [51] Robin Rombach, Patrick Esser, and Björn Ommer. Geometry-free view synthesis: Transformers and no 3D priors. In *ICCV*, 2021. 3
- [52] Stephane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Int. Conf. Art. Intell. Stat.* PMLR, 2011. 2
- [53] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022. 2
- [54] 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 *CVPR*, 2022. 3
- [55] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, and Jiajun Wu. ZeroNVS: Zero-shot 360-degree view synthesis from a single real image. *arXiv preprint arXiv:2310.17994*, 2023. 3
- [56] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. In *ICLR*, 2023. 2

- [57] Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, 2019. 3
- [58] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3d-structure-aware neural scene representations. *arXiv preprint arXiv:1906.01618*, 2019. 3
- [59] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. pages 2256–2265. PMLR, 2015. 2
- [60] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2
- [61] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 2, 5
- [62] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *CoRR*, abs/1907.05600, 2019.
- [63] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. *Advances in neural information processing systems*, 33:12438–12448, 2020. 2
- [64] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2021. 2
- [65] Shitao Tang, Fuayng Zhang, Jiacheng Chen, Peng Wang, and Furukawa Yasutaka. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. *arXiv preprint 2307.01097*, 2023. 3, 6
- [66] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezchikov, Joshua B. Tenenbaum, Frédo Durand, William T. Freeman, and Vincent Sitzmann. Diffusion with forward models: Solving stochastic inverse problems without direct supervision. *NeurIPS*, 2023. 2, 3, 5, 6
- [67] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: image synthesis using neural textures. *ACM TOG*, 2019. 3
- [68] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering. In *ICCV*, 2021. 3
- [69] Hung-Yu Tseng, Qinbo Li, Changil Kim, Suhub Alsisan, Jia-Bin Huang, and Johannes Kopf. Consistent view synthesis with pose-guided diffusion models. In *CVPR*, 2023. 2, 3
- [70] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*, 2018. 6
- [71] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *NeurIPS*, 2021. 2
- [72] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 3
- [73] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *arXiv preprint arXiv:2306.02018*, 2023. 2
- [74] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel view synthesis with diffusion models. In *ICLR*, 2023. 3
- [75] Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. SynSin: End-to-end view synthesis from a single image. In *CVPR*, 2020. 3
- [76] Jianfeng Xiang, Jiaolong Yang, Binbin Huang, and Xin Tong. 3d-aware image generation using 2d diffusion models. In *ICCV*, 2023. 3
- [77] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelNeRF: Neural radiance fields from one or few images. In *CVPR*, pages 4578–4587, 2021. 3, 7
- [78] Jason J. Yu, Fereshteh Forghani, Konstantinos G. Derpanis, and Marcus A. Brubaker. Long-term photometric consistent novel view synthesis with diffusion models. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023. 2, 3, 5, 6
- [79] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models, 2023. 4
- [80] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [81] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. In *ACM TOG*, 2018. 1, 5, 7, 8
- [82] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *CVPR*, 2023. 3