

Bolt3D: Generating 3D Scenes in Seconds

Stanislaw Szymanowicz^{1,2} Jason Y. Zhang¹ Pratul Srinivasan³
 Ruiqi Gao³ Arthur Brussee³ Aleksander Holynski³
 Ricardo Martin-Brualla¹ Jonathan T. Barron³ Philipp Henzler¹

¹Google ²Visual Geometry Group, University of Oxford ³Google DeepMind

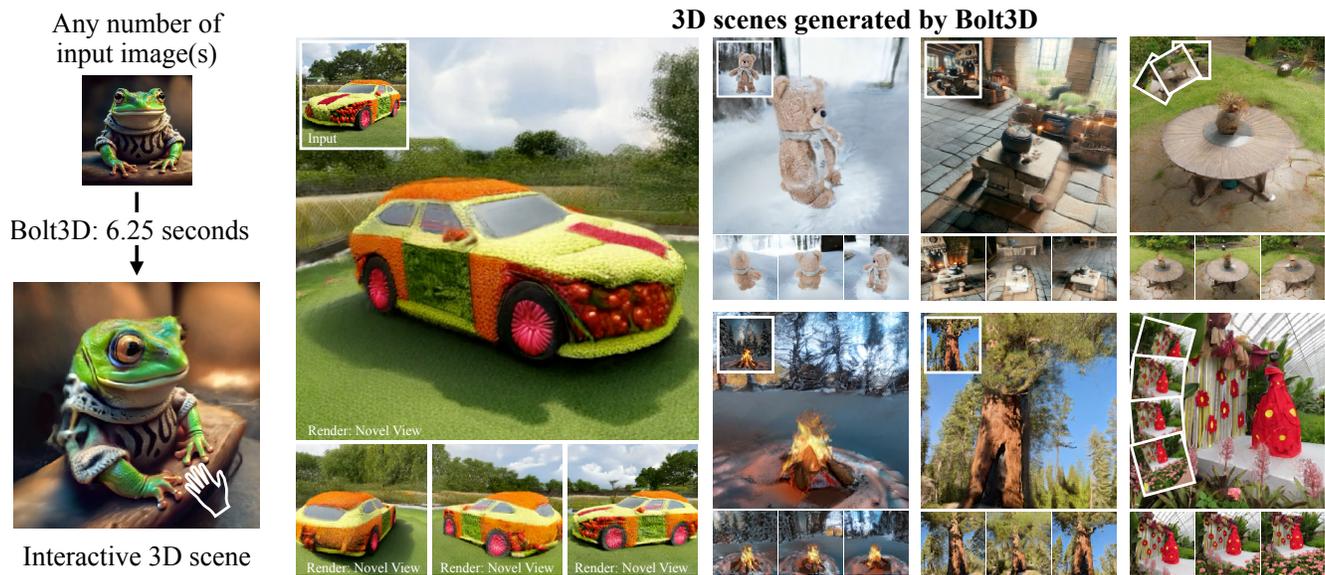


Figure 1. Given an arbitrary number of input images, Bolt3D directly outputs a 3D representation which can be rendered at interactive frame-rates. Operating in a feed-forward manner, generation takes mere seconds. Bolt3D features a latent diffusion model with a scalable 2D architecture, trained on large-scale appearance and geometry data, enabling generation of full 360° scenes from one or multiple input images. We invite the reader to explore these scenes in the interactive viewer available on the project website.

Abstract

We present a latent diffusion model for fast feed-forward 3D scene generation. Given one or more images, our model Bolt3D directly samples a 3D scene representation in less than seven seconds on a single GPU. We achieve this by leveraging powerful and scalable existing 2D diffusion network architectures to produce consistent high-fidelity 3D scene representations. To train this model, we create a large-scale multiview-consistent dataset of 3D geometry and appearance by applying state-of-the-art dense 3D reconstruction techniques to existing multiview image datasets. Compared to prior multiview generative models that require per-scene optimization for 3D reconstruction, Bolt3D reduces the inference cost by a factor of up to 300×. Project website: szymanowiczs.github.io/bolt3d.

1. Introduction

Modern image and video generative models generate compelling high-quality visual content, but these models sample 2D images, rather than an underlying 3D scene. The ability to directly generate 3D content instead would enable numerous applications, such as interactive visualization and editing. However, scaling modern diffusion-based generative models to generate detailed 3D scenes remains a significant challenge for the research community, primarily due to two reasons. First, representing and structuring (possibly unbounded) 3D data to enable training a diffusion model that generates full scenes at high resolution is an unsolved problem. Second, “ground truth” 3D scenes are extremely scarce compared to the abundant 2D image and video data used to train state-of-the-art generative models. As a result, many recent 3D generative models are limited to syn-

thetic objects [32, 77, 79, 92] or partial “forward-facing” scenes [31, 74, 83, 84]. Models that scale to real, full 360° scenes use camera-conditioned multiview or video diffusion models to turn input image(s) into a large “dataset” of synthetic observations [14, 38], from which an explicit 3D representation (such as a neural [41] or 3D Gaussian [28] radiance field) is then recovered via test-time optimization. While this approach is capable of producing high-quality 3D content, it is impractical; both sampling hundreds of augmented images with the multiview diffusion model and optimizing a 3D representation to match these images are slow and compute-intensive.

In this paper, we present a latent diffusion model for fast feed-forward 3D scene generation from one or more images. Our model, called Bolt3D, leverages the tremendous progress made in scalable 2D diffusion model architectures to generate an explicit 3D scene representation. We represent 3D scenes as sets of 3D Gaussians, stored in multiple 2D grids in which each cell stores the parameters of one pixel-aligned Gaussian [7, 61] (“Splatter Images” [61]). Crucially, unlike prior work, we use *more Splatter Images* than input views, and we *generate* Splatter Images using a diffusion model, which enables generating content for unobserved regions of the scene.

Our generation process consists of two parts: denoising color and position of each Gaussian and subsequently regressing each Gaussian’s opacity and shape. Given a set of posed input images and target camera poses, our model jointly predicts the scene appearance (pixel colors) viewed by the *target* cameras as well as per-pixel 3D coordinates of scene points in *all* cameras (both input and target). We enable Bolt3D to predict accurate high-resolution per-pixel 3D geometry by designing and training a Geometry Variational Auto-Encoder (VAE [12]) that resembles architectures used in image generation, but which we train from scratch using geometry data.

Unlike 2D image datasets, real-world 3D scene datasets are small and limited. To address this, we create a large-scale geometry dataset by running a robust Structure-from-Motion framework [29] on large-scale multi-view image datasets. We use this new dataset to train our geometric VAE and diffusion model.

Finally, to obtain a renderable 3D representation (which 3D point coordinates alone do not directly provide), we train a Gaussian head network that takes the full set of high-resolution images, geometry maps, and camera poses predicted by the denoising network as input and predicts the remaining properties of 3D Gaussians (opacities, shapes), and refined colors. The Gaussian head is supervised with rendering losses, which enables our full model’s output to render high-quality novel views.

We demonstrate that Bolt3D outperforms prior single- and few-view feed-forward 3D regression methods and syn-

thesizes detailed 3D scene content, even in ambiguous regions that are not observed in any input image. Furthermore, we show that Bolt3D reduces inference cost up to 300× compared to multi-view image generation methods, which typically require per-scene optimization to reconstruct 3D models.

2. Related work

Feed-forward 3D regression. Reconstructing a 3D scene from an input image (or from multiple images) is sometimes formulated as a feed-forward regression problem [8, 13, 15, 16, 23, 30, 40, 50, 64, 66, 67, 71, 75, 82]. Most relevant to us are recent approaches that output 3D Gaussians [28], a representation favored for its real-time rendering speed. Splatter Image [61] and pixelSplat [7] were the first to associate Gaussians with the pixels of one or two input images, respectively. Follow-up works iterated on this approach with improved architectures [63, 81, 91], depth conditioning [62, 78], explicit feature matching [9], and removing camera pose requirements [80, 94]. Although these regression-based methods are capable of accurately reconstructing observed regions, they tend to produce blurry results in unseen regions. In contrast, Bolt3D is a generative approach and is thereby capable of generating unobserved regions of the scene.

Reconstruction via 2D image generation. A common paradigm for addressing ambiguity in few-view 3D reconstruction is to generate multiple views of the scene, from which 3D can be recovered. Zero-1-to-3 [35] and 3DiM [73] did this for individual objects, and follow-up work improved the conditioning and sampling mechanism in a variety of ways [6, 17, 22, 52, 76, 85, 97]. Joint sampling improves sample consistency, and can be achieved by multi-view generators [14, 22, 24, 36, 56, 69, 95], or video models [10, 18, 34, 38, 59, 68]. Alternatively, other methods have explored using estimated depth to improve geometric consistency [24, 37, 43, 83, 84, 87]. The outputs of the generation process are 2D images, and reconstructing a renderable 3D asset requires an optimization-based reconstruction pipeline [28, 41]. This two-stage approach produces high-quality results, but requires minutes or hours of optimization time for real-world scenes. Our approach leverages a diffusion model to generate multiple views of the scene, but in our approach the 3D scene is a direct output from our model, so no optimization stage is required and the cost of inference is therefore decreased by multiple orders of magnitude.

3D Generation. Directly generating 3D representations circumvents the limitations of the previous works, by combining the ability of generative models to handle ambiguity while avoiding a costly distillation/optimization stage. Early diffusion-based models denoise a voxel grid that pa-

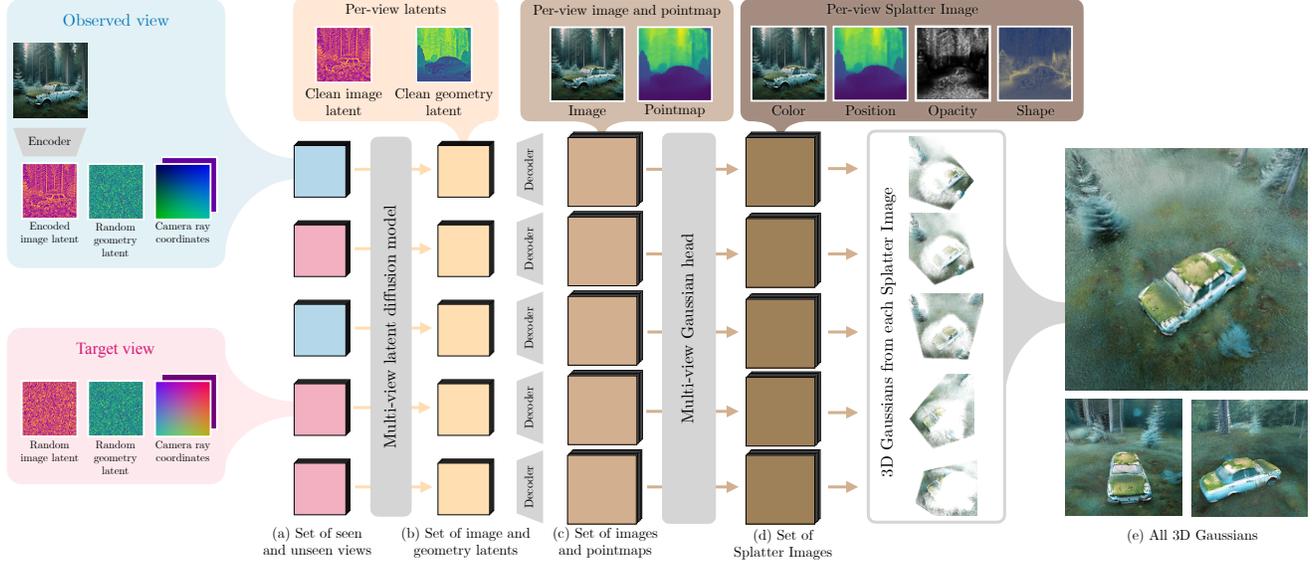


Figure 2. **Method.** Bolt3D takes as input one or more posed, observed images, and a set of target poses (a), and outputs a renderable 3D scene (e). First, we use a multi-view latent diffusion model to sample per-view latent appearance and geometry (b). The appearance and geometry latents are independently decoded to full-resolution images and pointmaps (c) using a pre-trained image VAE decoder and our trained geometry decoder, respectively. Next, a multi-view Gaussian head predicts the opacities and scales of pixel-aligned 3D Gaussians, and refines the predicted colors. Together with the pointmap from (c), these values form Splatter Images [61] (d), which can be combined to create a complete 3D Gaussian representation of the scene (e).

parameterizes a radiance field [42] and require 3D supervision. The difficulty in obtaining such 3D data motivated the formulation of a denoising objective on 2D images with a rendering bottleneck [1, 2, 60, 65], similar to early GAN [25]-based works that directly output voxel grids [19], radiance fields [44, 53, 86] or tri-planes [4, 5] using differentiable rendering to directly train on 2D images. Follow-up works improve the architecture [79] and the 3D representation [32, 70], but such approaches have shown limited success beyond individual objects and small baseline camera motion. Some more recent works [31, 54] exploit temporal locality to generate videos given camera trajectories. A second body of work aims to directly learn a ‘natively 3D’ latent space, showing remarkable success on bounded objects [32, 39, 48, 77, 90, 92], but lagging behind in quality when applied to 3D scenes with backgrounds [26, 47, 55, 89] due to limitations of the autoencoder. We opt for a view-centric representation and analyze the architecture and training recipe of the autoencoder to achieve very high visual fidelity. Most similar to our work is latentSplat [74] which learns a latent geometry representation similar to ours. However, latentSplat is limited to single object categories, is only applicable to 2-view reconstruction and features a VAE-GAN framework. Our model works on any object or scene, can take any number of images as input and builds on a powerful latent diffusion model. We also demonstrate that our method performs better.

3. Preliminaries

Latent Diffusion Models. The key component of our method is a latent diffusion model [21, 58]. A diffusion model defines a forward process by gradually adding Gaussian noise to a data point \mathbf{x}_0 : $p(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\alpha_t\mathbf{x}_0, \sigma_t)$, until \mathbf{x}_t is close to a Gaussian. A denoiser model $\hat{\mathbf{x}}_\theta(\mathbf{x}_t; t)$ is learned to predict the clean sample \mathbf{x}_0 given the noisy sample \mathbf{x}_t at time step t by minimizing a weighted ℓ_2 loss:

$$\theta = \arg \min_{\theta} \mathbb{E}_{p_{\mathbf{x}_0}, p(\mathbf{x}_t|\mathbf{x}_0)} w(t) \|\hat{\mathbf{x}}_\theta(\mathbf{x}_t; t) - \mathbf{x}_0\|_2^2 \quad (1)$$

Other parametrizations of the denoiser model have been proposed, such as predicting the noise added to the clean data or a combination of the clean data and the noise like the v-prediction [51]. A latent diffusion model first compresses the data into a lower-dimensional latent space and then builds a diffusion model in the latent space. A two-stage training procedure is typically applied: the compression is first learned by a VAE and fixed, followed by learning the diffusion model.

3D Gaussian Representation. We leverage a set of 3D Gaussians as the scene representation: $\mathcal{G} = \{\boldsymbol{\mu}_i, \sigma_i, \boldsymbol{\Sigma}_i, \mathbf{c}_i\}_{i=1}^G$ where $\boldsymbol{\mu}_i \in \mathbb{R}^3$ is the mean of the Gaussian, $\sigma_i \in [0, 1]$ its opacity, $\boldsymbol{\Sigma}_i \in \mathbb{R}^{3 \times 3}$ its covariance matrix and $\mathbf{c}_i \in \mathbb{R}^3$ is the isotropic color. These Gaussians can be rendered to a camera efficiently using the splatting operation [28].

Few-view 3D reconstruction. The goal of few view reconstruction is to recover the representation \mathcal{G} given a small collection of N input views, where “views” are paired images \mathbf{I} and their camera poses $\mathbf{\Pi}$. In our case, we assume $N \in \{1 - 4\}$. With such limited input views and therefore insufficient coverage of the scene, pure optimization-based reconstruction becomes an ill-posed problem, as many representations \mathcal{G} can explain the same set of posed images $(\mathbf{I}, \mathbf{\Pi})$. Therefore, we instead formulate the task as a generation problem: learn and sample from a conditional distribution of the 3D scene given the input views: $p(\mathcal{G}|\mathbf{I}, \mathbf{\Pi})$.

4. Method

Bolt3D takes a single or multiple images and their camera poses as input, and outputs a 3D Gaussian representation. The model consists of two components: a latent diffusion model that generates more views and the 3D location per pixel (i.e., a 3D pointmap) for each view, and a feed-forward Gaussian head model that takes the outputs from the diffusion model and predicts full parameters of one colored 3D Gaussian per pixel (i.e. a Splatter Image) for each view. Below we describe the 3D Gaussian representation (Sec. 4.1), the latent diffusion model (Sec. 4.2), the feed-forward Gaussian head model (Sec. 4.3) and the training process (Sec. 4.4). We present an overview in Fig. 2.

4.1. 3D Representation

Assuming K views that can be observed or generated, we leverage a 3D Gaussian representation consisting of K Splatter Images [61], i.e., pixel-aligned colored 3D Gaussians for each view. Each 3D Gaussian contains four properties: color, 3D position (mean of the Gaussian), opacity, and covariance matrix. In contrast to prior reconstruction-based methods, we leverage a generative approach that hallucinates new content in a wider range of the scene than what is covered by the input views.

Factorized sampling. We factorize the generation of Gaussian parameters into two parts: first, a latent diffusion model is leveraged to generate the color and 3D position. Then, a feed-forward Gaussian head model takes the color and 3D position as input and predicts the opacity, covariance matrix and refined color. The motivation is that we can easily get data of the colors and 3D positions by collecting captured images and running dense Structure-from-Motion to serve as the target of the latent diffusion model. On the other hand, finding direct supervision for the covariance matrices and opacities is non-trivial. However, given the colors and 3D positions generated from the diffusion model, the covariance matrix and opacity are much less ambiguous, and therefore can be modeled by a deterministic mapping function that is supervised with a rendering loss.

4.2. Geometric multi-view latent diffusion model

We train a multi-view latent diffusion model that jointly models images and 3D pointmaps. Specifically, the model takes one or more images \mathbf{I}^{cond} and their camera poses $\mathbf{\Pi}^{\text{cond}}$ as input. Given multiple target camera poses $\mathbf{\Pi}^{\text{tgt}}$, the model learns to capture the joint distribution of the target images \mathbf{I}^{tgt} , the target pointmaps \mathbf{P}^{tgt} and the source view pointmaps \mathbf{P}^{cond} :

$$p(\mathbf{I}^{\text{tgt}}, \mathbf{P}^{\text{tgt}}, \mathbf{P}^{\text{cond}} | \mathbf{I}^{\text{cond}}, \mathbf{\Pi}^{\text{cond}}, \mathbf{\Pi}^{\text{tgt}}) \quad (2)$$

Model architecture. We finetuned our model from a pre-trained multi-view image diffusion model, which itself was finetuned from text-to-image latent diffusion model, to maximally maintain the generalization ability of the model while being trained on limited amount of multi-view data with 3D positions. The images and geometry (i.e., pointmaps and cameras) are encoded and decoded by two separate VAEs that $8\times$ downsample the input signals spatially (i.e. from 512×512 in the pixel space to 64×64 in the latent space). The image VAE is pre-trained and frozen while the geometry VAE is trained from scratch. The camera pose is parametrized as a 6-dim raymap that encodes the ray origin and direction at each spatial location. At the input of the latent diffusion, the image latent, geometry latent and a raymap of the same size of the latents are channel-wise concatenated. We use a \mathbf{v} -parametrization and a \mathbf{v} -prediction loss for the diffusion model [51]. The model is pre-trained on 8 and fine-tuned on 16 views in total, and 1-3 views are randomly sampled as input views. During sampling we generate 16 views in total.

Geometry VAE. We train a geometry VAE to jointly encode the pointmap \mathbf{P} and camera raymap \mathbf{r} of a view into a geometry latent:

$$\boldsymbol{\mu}, \boldsymbol{\sigma} = \mathcal{E}(\mathbf{P}, \mathbf{r}), \mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}), \hat{\mathbf{P}}, \hat{\mathbf{r}} = \mathcal{D}(\mathbf{z}), \quad (3)$$

where \mathcal{E} is a convolutional encoder and \mathcal{D} is a transformer decoder. The model is optimized by minimizing the following training objective, which is a combination of the standard VAE objective and a geometry-specific loss:

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_1 \mathcal{L}_{\text{KL}} + \lambda_2 \mathcal{L}_{\text{grad}}. \quad (4)$$

The reconstruction loss \mathcal{L}_{rec} is given by

$$\mathcal{L}_{\text{rec}} = \|\mathbf{w}(\hat{\mathbf{P}} - \mathbf{P})\|_2^2 + \|\hat{\mathbf{r}} - \mathbf{r}\|_2^2, \quad (5)$$

where \mathbf{w} is a per-pixel weighting depending on the distance from the point to the center of the scene in the local camera frame (see supplement for more details). Intuitively, this encourages accurate geometry while accounting for lower confidence further from the camera. \mathcal{L}_{KL} is given by:

$$\mathcal{L}_{\text{KL}} = -\frac{1}{2} \sum_k (1 + \log \sigma_k - \mu_k^2 - \sigma_k). \quad (6)$$

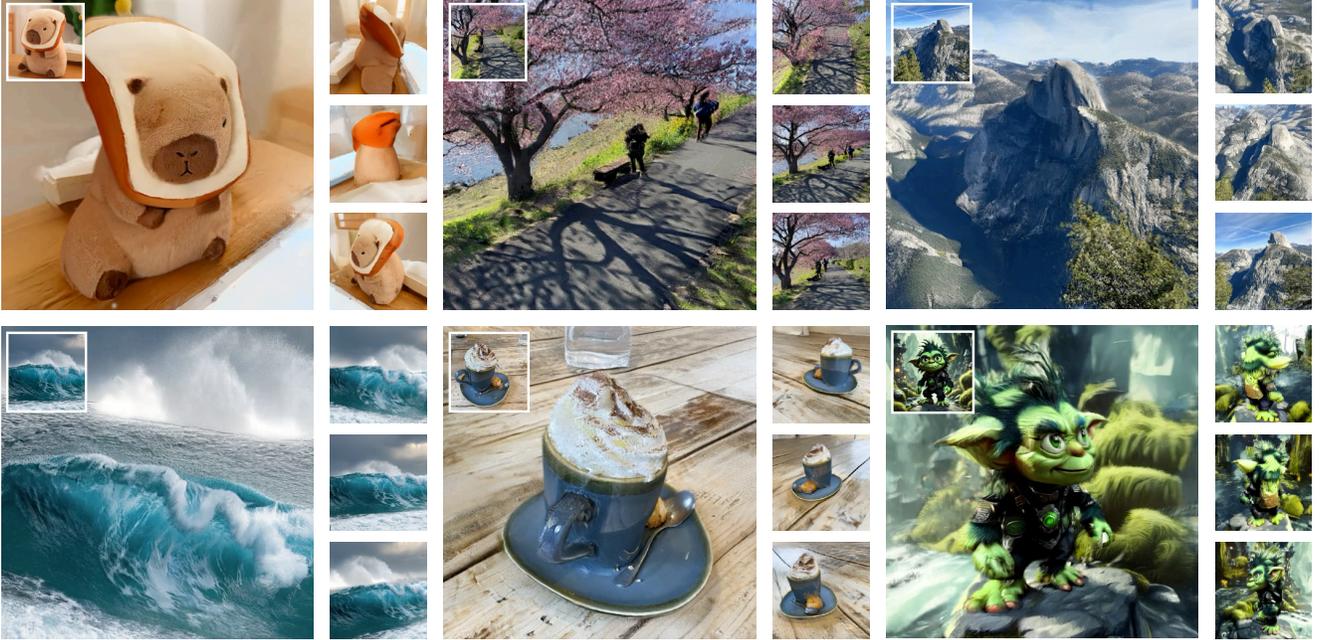


Figure 3. **Qualitative results.** We show renders of our 3D scenes reconstructed from just one input image (top left corner in each image) in a feed-forward manner. Inference takes only 7 seconds on a single GPU.

Finally, we add an ℓ_2 reconstruction error of the vertical and horizontal gradients of the pointmap, to improve boundary sharpness in the decoded pointmap:

$$\mathcal{L}_{\text{grad}} = \|\partial_u \hat{\mathbf{P}} - \partial_u \mathbf{P}\|_2^2 + \|\partial_v \hat{\mathbf{P}} - \partial_v \mathbf{P}\|_2^2. \quad (7)$$

4.3. Gaussian Head

Given cameras and generated images and pointmaps, we train a multi-view feedforward Gaussian head model to output the refined color, opacities and covariance matrices of 3D Gaussians stored in Splatter Images. We first calibrate the generated pointmaps to be pixel aligned under the corresponding camera coordinate system. That is, we transform the 3D points into the camera coordinate, keep z and set x, y based on the camera ray and z . While the VAE decodes color and pointmaps independently for each view, we found that implementing a multi-view Gaussian head model is crucial. A U-ViT [3] architecture is applied with $4\times$ patchification before the transformer blocks, see Supplementary for more details. The Gaussian head takes 8 views as input, and is trained using photometric losses (i.e. an L2 loss and a perceptual loss [93]) on rendered images from 4 subsampled input views and 8 novel views.

4.4. Training

Data. Our method requires supervision of dense, multi-view consistent pointmaps associated with the multi-view images. We leverage a recent state-of-the-art dense reconstruction and matching method, MAST3R [29], running the depth and feature estimation followed by bundle adjustment

of all pixels for 20-25 images per scene. This setup allows for complete scene coverage and outputs multi-view consistent 3D pointmaps. While the data is not perfect, we find that the residual noise and geometric imperfections are minor enough to not affect our method significantly (partially due to using rendering losses, in addition to geometry losses). We run MAST3R on all scenes from CO3D [46], MVImg [88], RealEstate10K (RE10K) [96] and DL3DV-7K [33] forming a dataset of around 300k multi-view consistent 3D scenes. We use standard train-test splits in CO3D, MVImg and RE10K, and we use the first 6K scenes in DL3DV for training and last 1K for testing. In addition, we leverage synthetic object datasets (Objaverse [11] and a high-quality internal object dataset) with their corresponding pointmap renderings. We train on a mix of these datasets, sampling the real scenes with equal probability from all 4 real datasets and sampling the synthetic datasets 1:2 compared to sampling real data.

Training Protocol. We train our model in 3 stages:

1. **Geometry VAE.** We first train the model at 256×256 resolution for 3 million iterations, then finetune at 512×512 for another 250k iterations.
2. **Gaussian head.** Given ground truth color and autoencoded geometry, the Gaussian head is trained for 100k iterations with rendering losses to output Splatter Images.
3. **Latent diffusion model.** We initialized our latent diffusion model from CAT3D [14] and trained it for 700k iterations on the 8-view setup before finetuning it for 70k iterations on 16-views. See supplement for more details.

5. Experiments

We begin our experiments with presenting the results of our method in Fig. 3, and we encourage the reader to view the videos and interactive visualizations on the project website which include high-quality reconstructions on a wide range of inputs.

Our experiments are divided into 4 sections. First, we illustrate that modeling ambiguity using a generative model is crucial for few-view reconstruction by evaluating against state-of-the-art regression-based approaches. Second, we show that our approach for modeling ambiguity outperforms recent approaches for feed-forward 3D generation from few input images. Third, we evaluate the speed-quality trade-off between state-of-the-art optimization-based methods and our method. Finally, we analyze the Geometry VAE and show its crucial role in the performance of our method. For Geometry VAE and gaussian decoder ablations we refer to the supplement.

When comparing against prior works we evaluate our method on the number of views other methods were designed for. We evaluate performance at center crops of 512×512 resolution unless stated otherwise.

Metrics. We quantify 3D reconstruction quality with standard metrics for novel-view synthesis: PSNR, SSIM, LPIPS [93] and FID [20] to measure pixel-wise, patch-wise, perceptual and distribution similarity, respectively.

5.1. Comparison to 3D Regression.

We compare our method with state-of-the-art feed-forward Gaussian Splat regression methods: Flash3D [62] (single-view) and DepthSplat [78] (few-view).

Protocol. We use 1-view RE10K for comparison against Flash3D by using the split from [14, 76] and using the first frame in each video as the conditioning frame, with the target frames as specified in the split. This split contains larger camera motion (90 frames) than that used in the original Flash3D paper, thus evaluating the reconstruction ability beyond small (± 30 frames) camera motion. Additionally, we include an evaluation on CO3D, also adapted from [14, 76] to evaluate both methods under stronger self-occlusion, as is typical for single-view reconstruction. For comparison against DepthSplat [78], we use DL3DV, which features large camera motion, and we evaluate both methods in the 2-view and 4-view settings, using source and target views from [78]. We run evaluations on the scenes using the overlap of their and our testing sets.

Results. We find that Bolt3D outperforms both methods, as shown in Tab. 1. The strong performance of our method compared to feed-forward reconstruction highlights the importance of modeling ambiguity, as we show in Fig. 4. This observation is consistent with the fact that the biggest gain

	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow
1-view RE10K	Flash3D	17.40	0.699	0.419	96.9
	Ours	21.03	0.805	0.257	55.5
1-view CO3D	Flash3D	14.43	0.552	0.608	174.8
	Ours	16.78	0.562	0.505	97.5
2-view DL3DV	DepthSplat	16.25	0.515	0.465	95.9
	Ours	17.75	0.551	0.392	64.5
4-view DL3DV	DepthSplat	19.48	0.638	0.327	58.8
	Ours	20.64	0.653	0.310	48.2

Table 1. **Comparison to regression-based methods.** Our generative approach improves performance across multiple datasets and numbers of input views. The biggest gain is seen in the 1-view setting where ambiguity is the greatest.

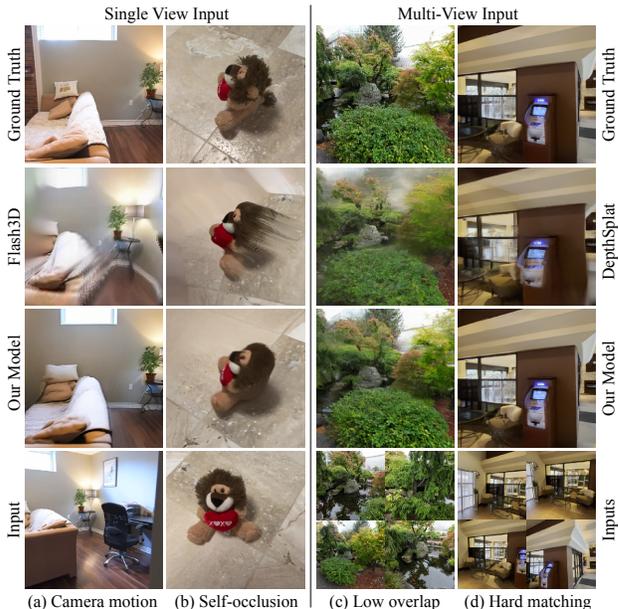


Figure 4. **Modeling ambiguity** is advantageous in both single- and few-view reconstruction. Given just a single input image, our generative framework synthesizes realistic unobserved outside the observed field-of-view (a) and in occluded regions (b). Bolt3D also performs well in challenging scenarios given multiple input images that have a small overlap (c) or cases where feature matching is difficult (d).

in performance is observed in the 1-view setting, where ambiguity is the largest.

5.2. Comparison to Feed-Forward 3D Generation.

We compare our method to two recent feed-forward 3D generative methods capable of reconstructing real scenes: LatentSplat [74] and Wonderland [31].

Protocol. When comparing to LatentSplat, we evaluate at 256×256 resolution on the datasets they proposed: RealEstate10k and the hydrants and teddybear categories of CO3D, with their extrapolation splits. When comparing to Wonderland, we follow their quantitative evaluation proto-

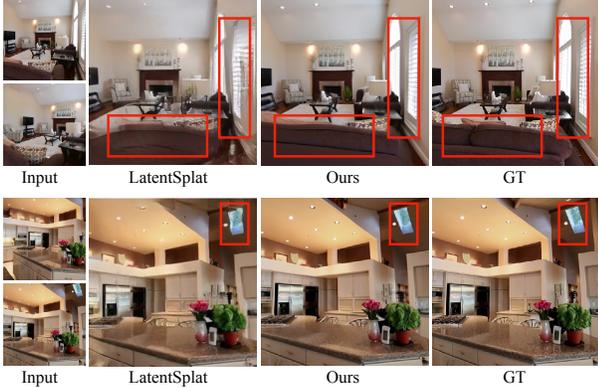


Figure 5. Our latent diffusion framework generates more realistic details the VAE-GAN framework proposed by LatentSplat.

	Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow
1-view Re10k	Wonderland*	17.15	0.550	0.292	-
	Ours	21.54	0.747	0.234	11.30
2-view RE10K	LatentSplat	22.62	0.777	0.196	2.79
	Ours	23.13	0.806	0.166	3.66
2-view CO3D ted.	LatentSplat	17.71	0.533	0.434	71.12
	Ours	18.94	0.605	0.393	76.24
2-view CO3D hyd.	LatentSplat	15.78	0.306	0.426	48.03
	Ours	17.38	0.437	0.390	76.12

Table 2. **Comparison to 3D generative methods.** Our diffusion-based framework with explicit 3D supervision achieves the best image-level metrics. * denotes results reported in [31]

col on RealEstate10K, i.e. randomly sampling 1000 testing scenes, using the first frame of the video as the source frame and measuring the rendering quality of 14 views following the source view given that one input view.

Results. We observe that our method outperforms LatentSplat qualitatively (Fig. 5) and quantitatively (Tab. 2) on image-level metrics. We suspect that latentsplat’s adversarial loss gives it an advantage in the distribution-level FID metric. Fig. 5 illustrates that our diffusion model generates higher-quality details than LatentSplat’s VAE-GAN. We also demonstrate that our method outperforms Wonderland in Tab. 2. We hypothesize this is due to explicit modeling of geometry in the autoencoder and in the diffusion model. It is also worth noting that Wonderland uses a video model and takes 5 minutes to generate a scene, while ours generates 16 splatter images in around 6 seconds.

5.3. Optimization-based 3D Reconstruction.

We explore the speed-quality trade-off versus a state-of-the-art optimization-based method, CAT3D [14]. We also include qualitative comparisons to RealmDreamer [57].

Protocol. We report the performance of CAT3D in a 3-cond setting evaluated on 512×512 center-crops. We use the same datasets and splits as proposed by the original paper. We measure inference cost in GPU-minutes spent on

	Method	3-view				Infer. Cost
		PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	gpu-min \downarrow
RE10	CAT3D	29.56	0.937	0.134	13.75	77.28
	Ours	27.00	0.905	0.154	27.40	0.25
LLFF	CAT3D	22.06	0.745	0.194	37.54	80.00
	Ours	18.75	0.562	0.341	96.61	0.25
DTU	CAT3D	19.97	0.809	0.202	41.76	72.00
	Ours	18.59	0.738	0.312	67.49	0.25
CO3D	CAT3D	20.85	0.673	0.329	44.33	73.60
	Ours	19.41	0.628	0.416	79.60	0.25
Mip.	CAT3D	16.62	0.377	0.515	91.82	73.60
	Ours	15.67	0.309	0.540	124.17	0.25

Table 3. **Comparison to optimization-based methods.** Our method performs competitively with state-of-the-art optimization-based methods, while reducing the inference cost $300\times$.

end-to-end reconstruction process. In addition, we collect qualitative results shared on the official project websites of RealmDreamer and CAT3D.

Inference Cost. Our method takes 6.25 seconds to reconstruct one scene on a single H100 NVIDIA GPU or 15 seconds on an A100. (We report A100 times in Tab. 3, to be comparable to CAT3D). CAT3D takes 5 seconds to generate 80 images on 16 GPUs, and needs 640-800 generated images per scene, followed by 4 minutes of reconstruction, amounting to around 5 minutes, depending on the dataset.

Results. In Tab. 3 we observe that our method shows strong performance, and it does so while requiring $300\times$ less compute for inference. Qualitatively, our method produces very high quality reconstructions for a range of scenes, and in Fig. 6 we illustrate that our method generates better results than RealmDreamer, and sometimes sharper results than CAT3D, especially on the 3DGS variant and in backgrounds or fine details. This is due to the optimization process in CAT3D regressing to the mean in the case of inconsistent generations, especially in the less robust 3DGS optimization process. While our method does not outperform CAT3D, it is still capable of generating high-quality 3D scenes from a diverse range of inputs (as seen in the supplementary material), and we argue that a $300\times$ reduction in inference cost well justifies a small drop in quality.

5.4. Image VAEs generalize poorly to geometry

Typically, VAEs are pre-trained on RGB images for latent diffusion models that generate images [49]. We show that these VAEs do not work well for unbounded geometry.

Metrics. We use metrics commonly used for depth estimation [45, 72]—relative error (AbsRel) between the target z-component of the pointmap z and prediction \hat{z} , $\text{Rel.} = \|\hat{z} - z\| / z$, and the prediction threshold accuracy, $\delta_\epsilon = \max(\hat{z}/z, z/\hat{z}) < \epsilon$ with $\epsilon = 1.01$. Additionally, we measure the re-projection error: the mean Euclidean distance on the

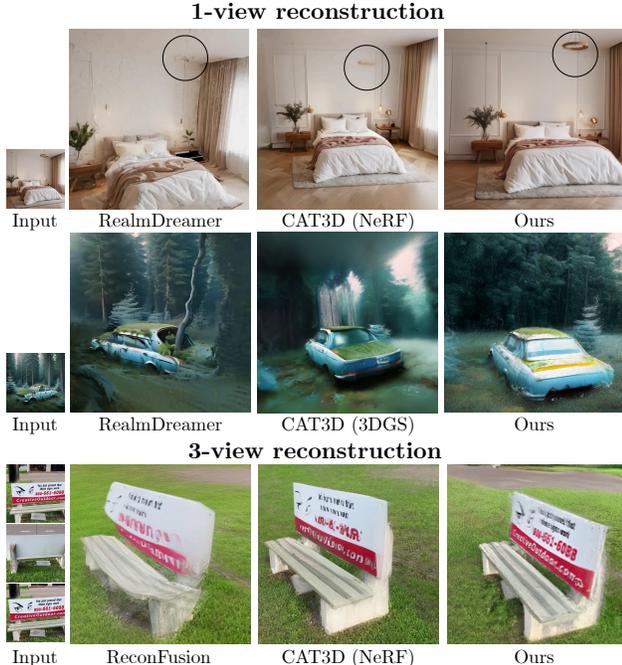


Figure 6. **Qualitative comparison to optimization-based methods.** The reconstruction quality of our method compares favorably with that of optimization-based methods, and sometimes even exhibits sharper details, while using 300× less compute.

	Synth. (Bounded)/Real (Unbounded)		
	Rel. × 100↓	$\delta_{1.01}$ ↑	Δuv (px)↓
Im AE + mean-z scaling	1.03/17.9	67.1/24.2	9.72/160
Im AE + max-xyz scaling	4.56/9.74	19.8/15.0	95.6/907
Im AE + nonlinear scaling	1.34/15.8	58.8/40.3	16.9/245
Geo AE, Conv decoder	0.530/0.684	87.2/81.4	3.94/4.66
Geo AE, ViT decoder (Ours)	0.636/0.670	82.4/81.5	3.85/2.69

Table 4. **Comparison of VAEs Trained for Images (Im AE) vs. Geometry (Geo AE) on Pointmaps.** Our autoencoder trained specifically for geometry tasks performs much better than pre-trained image auto-encoders, regardless of how points are scaled. Our transformer-based decoder outperforms the convolution-based one on real data.

image plane when re-projecting the point back to the source camera $\Delta uv = \frac{1}{N} \sum_i^N \|(u_i, v_i) - (\hat{u}_i, \hat{v}_i)\|_2$, where u and v are the x and y coordinates, respectively, in pixels on the image. We measure performance separately for synthetic and real data, both at 512×512 resolution.

Scaling. Our geometry data is not metric, so the scale of the scene is arbitrary. Properly normalizing the data before inputting it to the VAE can significantly impact performance. We consider 3 methods for scaling data. We experiment with (1) scaling the scene such that mean scene depth is 1, (2) scaling the scene such that the maximum coordinate value in any direction is 1 and (3) applying a nonlinear contraction function (the sigmoid function) before inputting it to the encoder, and applying its inverse after decoding. Our

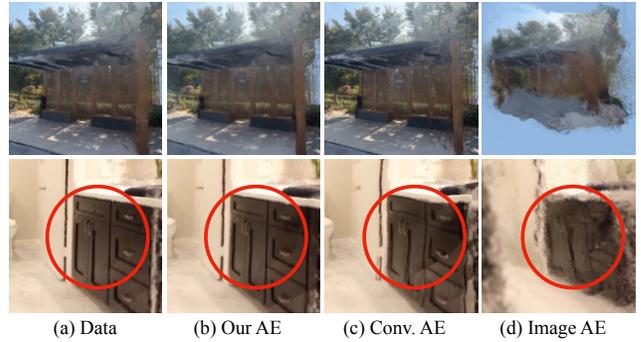


Figure 7. Given input geometry (a), our Auto-Encoder (b) accurately reconstructs input geometry. Using a convolutional decoder (c) results in inaccurate geometry and using a pre-trained Image Auto-Encoder (d) fails catastrophically for both indoor and large-scale outdoor scenes.

autoencoder uses (1).

Results. In Tab. 4 we find that pre-trained image autoencoders work reasonably well for synthetic, bounded data. This conclusion aligns with prior work showing that VAEs pre-trained on RGB images can be applied effectively to bounded relative depth $d \in [0, 1]$ [27]. However, we find that regardless of the scaling method, autoencoders trained on images struggle to autoencode pointmaps. Meanwhile, our autoencoder trained on geometry maintains high accuracy (80% of points are within 1% of ground truth) for both synthetic and real data. We visualize 2D pointmaps and point cloud renders in Fig. 7 and observe that the image autoencoder fails catastrophically in outdoor scenes, and exhibits obvious inaccuracies when autoencoding indoor scenes. While the convolutional decoder performs comparably to the transformer quantitatively, we empirically observe in Fig. 7 that it introduces jarring artifacts, e.g., bent lines. This analysis illustrates that our Geometry AE and its architecture are instrumental to high-quality results on unbounded scenes from our method.

6. Conclusion

We have presented Bolt3D, a fast feed-forward method that generates detailed 3D scenes in less than 7 seconds. To enable this capability, we propose a 3D scene representation that allows for denoising high-resolution 3D scenes using powerful 2D latent diffusion architectures. We create a large-scale 3D scene dataset to train Bolt3D, and demonstrate that it accurately models ambiguity, enabling high-quality 1-view reconstructions where regression-based methods fail. Our proposed feed-forward approach is capable of reconstructing a wide variety of scenes and reduces the cost of 3D generation by 300× compared to existing optimization-based methods, opening up opportunities for 3D content creation at scale.

Acknowledgments We would like to express our deepest gratitude to Ben Poole for helpful suggestions, guidance, and contributions. We also thank George Kopanas, Sander Dieleman, Matthew Burruss, Matthew Levine, Daniel Duckworth, Peter Hedman, Songyou Peng, Rundi Wu, Alex Trevithick, Hadi Alzayer, David Charatan, Jiapeng Tang and Akshay Krishnan for valuable discussions and insights. Finally, we extend our gratitude to Shlomi Fruchter, Kevin Murphy, Mohammad Babaeizadeh, Han Zhang and Amir Hertz for training the base text-to-image latent diffusion model. This work was done while S. Szymanowicz was on an internship at Google. He is supported by the EPSRC Doctoral Training Partnerships Scholarship (DTP) EP/R513295/1 and the Oxford-Ashton Scholarship.

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. *CVPR*, 2023. 3
- [2] Titas Anciukevičius, Fabian Manhardt, Federico Tombari, and Paul Henderson. Denoising diffusion via image-based rendering. In *ICLR*, 2024. 3
- [3] Fan Bao, Shen Nie, Kaiwen Xue, Yue Cao, Chongxuan Li, Hang Su, and Jun Zhu. All are worth words: A vit backbone for diffusion models. *CVPR*, 2023. 5
- [4] Miguel Angel Bautista, Pengsheng Guo, Samira Abnar, Walter Talbott, Alexander Toshev, Zhuoyuan Chen, Laurent Dinh, Shuangfei Zhai, Hanlin Goh, Daniel Ulbricht, Afshin Dehghan, and Josh Susskind. Gaudi: A neural architect for immersive 3d scene generation. *NeurIPS*, 2022. 3
- [5] 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. *CVPR*, 2022. 3
- [6] Eric R. Chan, Koki Nagano, Matthew A. Chan, Alexander W. Bergman, Jeong Joon Park, Axel Levy, Miika Aitala, Shalini De Mello, Tero Karras, and Gordon Wetzstein. GeNVS: Generative Novel View Synthesis with 3D-Aware Diffusion Models. *ICCV*, 2023. 2
- [7] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelSplat: 3D Gaussian Splats from Image Pairs for Scalable Generalizable 3D Reconstruction. *CVPR*, 2024. 2
- [8] 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 *ICCV*, 2021. 2
- [9] Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-Jen Cham, and Jianfei Cai. MVSplat: Efficient 3D Gaussian Splatting from Sparse Multi-View Images. *ECCV*, 2024. 2
- [10] Yuedong Chen, Chuanxia Zheng, Haofei Xu, Bohan Zhuang, Andrea Vedaldi, Tat-Jen Cham, and Jianfei Cai. MVSplat360: Feed-forward 360 Scene Synthesis from Sparse Views. *NeurIPS*, 2024. 2
- [11] 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 *CVPR*, 2023. 5
- [12] Max Welling Diederik P. Kingma. Auto-encoding variational bayes. *ICLR*, 2014. 2
- [13] Yilun Du, Cameron Smith, Ayush Tewari, and Vincent Sitzmann. Learning to render novel views from wide-baseline stereo pairs. *CVPR*, 2023. 2
- [14] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul P. Srinivasan, Jonathan T. Barron, and Ben Poole. CAT3D: Create Anything in 3D with Multi-View Diffusion Models. *NeurIPS*, 2024. 2, 5, 6, 7
- [15] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV*, 2019. 2
- [16] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. AtlasNet: A Papier-Mâché Approach to Learning 3D Surface Generation. In *CVPR*, 2018. 2
- [17] Jiatao Gu, Alex Trevithick, Kai-En Lin, Joshua M Susskind, Christian Theobalt, Lingjie Liu, and Ravi Ramamoorthi. NerfDiff: Single-image View Synthesis with NeRF-guided Distillation from 3D-Aware Diffusion. *ICML*, 2023. 2
- [18] Hao He, Yinghao Xu, Yuwei Guo, Gordon Wetzstein, Bo Dai, Hongsheng Li, and Ceyuan Yang. Cameractrl: Enabling camera control for text-to-video generation. *ICLR*, 2025. 2
- [19] Philipp Henzler, Niloy J. Mitra, and Tobias Ritschel. Escaping plato’s cave: 3d shape from adversarial rendering. *ICCV*, 2019. 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. *NeurIPS*, 2018. 6
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising Diffusion Probabilistic Models. *NeurIPS*, 2020. 3
- [22] Lukas Höllein, Aljaž Božič, Norman Müller, David Novotny, Hung-Yu Tseng, Christian Richardt, Michael Zollhöfer, and Matthias Nießner. Viewdiff: 3d-consistent image generation with text-to-image models. *CVPR*, 2024. 2
- [23] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. LRM: Large Reconstruction Model for Single Image to 3D. *ICLR*, 2024. 2
- [24] Hanzhe Hu, Zhizhuo Zhou, Varun Jampani, and Shubham Tulsiani. Mvd-fusion: Single-view 3d via depth-consistent multi-view generation. In *CVPR*, 2024. 2
- [25] Mehdi Mirza Bing Xu David Warde-Farley Sherjil Ozair Aaron Courville Yoshua Bengio Ian J. Goodfellow, Jean Pouget-Abadie. Generative adversarial networks. *NeurIPS*, 2014. 3
- [26] Jun Gao Sanja Fidler Andreas Geiger Karsten Kreis Katja Schwarz, Seung Wook Kim. Learning 3d-aware latent diffusion models in view space. *ICLR*, 2024. 3

- [27] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing Diffusion-Based Image Generators for Monocular Depth Estimation. *CVPR*, 2024. 8
- [28] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. on Graphics (TOG)*, 2023. 2, 3
- [29] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding Image Matching in 3D with MAST3R. *ECCV*, 2024. 2, 5
- [30] Jiaxin Li, Zijian Feng, Qi She, Henghui Ding, Changhu Wang, and Gim Hee Lee. MINE: Towards Continuous Depth MPI with NeRF for Novel View Synthesis. *ICCV*, 2021. 2
- [31] Hanwen Liang, Junli Cao, Vidit Goel, Guocheng Qian, Sergei Korolev, Demetri Terzopoulos, Konstantinos N Platanotis, Sergey Tulyakov, and Jian Ren. Wonderland: Navigating 3D Scenes from a Single Image. *arXiv preprint arXiv:2412.12091*, 2024. 2, 3, 6, 7
- [32] Chenguo Lin, Panwang Pan, Bangbang Yang, Zeming Li, and Yadong Mu. DiffSplat: Repurposing Image Diffusion Models for Scalable 3D Gaussian Splat Generation. *ICLR*, 2025. 2, 3
- [33] Lu Ling, Yichen Sheng, Zhi Tu, Wentian Zhao, Cheng Xin, Kun Wan, Lantao Yu, Qianyu Guo, Zixun Yu, Yawen Lu, et al. DL3DV-10K: A Large-scale Scene Dataset for Deep Learning-based 3D Vision. In *CVPR*, 2024. 5
- [34] Fangfu Liu, Wenqiang Sun, Hanyang Wang, Yikai Wang, Haowen Sun, Junliang Ye, Jun Zhang, and Yueqi Duan. Reconx: Reconstruct any scene from sparse views with video diffusion model. *arXiv:2408.16767*, 2024. 2
- [35] 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. *ICCV*, 2023. 2
- [36] 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. *ICLR*, 2024. 2
- [37] Baorui Ma, Huachen Gao, Haoge Deng, Zhengxiong Luo, Tiejun Huang, Lulu Tang, and Xinlong Wang. You See it, You Got it: Learning 3D Creation on Pose-Free Videos at Scale. *CVPR*, 2025. 2
- [38] Luke Melas-Kyriazi, Iro Laina, Christian Rupprecht, Natalia Neverova, Andrea Vedaldi, Oran Gafni, and Filippos Kokkinos. IM-3D: Iterative Multiview Diffusion and Reconstruction for High-Quality 3D Generation. *ICML*, 2024. 2
- [39] Xuyi Meng, Chen Wang, Jiahui Lei, Kostas Daniilidis, Jiatao Gu, and Lingjie Liu. Zero-1-to-G: Taming Pretrained 2D Diffusion Model for Direct 3D Generation. *arXiv preprint arXiv:*, 2024. 3
- [40] Michael Oechsle, Andreas Geiger, Michael Niemeyer, Lars Mescheder. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. *CVPR*, 2020. 2
- [41] 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. *ECCV*, 2020. 2
- [42] Norman Müller, Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulo, Peter Kotschieder, and Matthias Nießner. DiffRF: Rendering-guided 3D Radiance Field Diffusion. *CVPR*, 2023. 3
- [43] Norman Müller, Katja Schwarz, Barbara Rössle, Lorenzo Porzi, Samuel Rota Bulò, Matthias Nießner, and Peter Kotschieder. Multidiff: Consistent novel view synthesis from a single image. *CVPR*, 2024. 2
- [44] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. *CVPR*, 2021. 3
- [45] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ICCV*, 2021. 7
- [46] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. *ICCV*, 2021. 5
- [47] Xuanchi Ren, Yifan Lu, Hanxue Liang, Jay Zhangjie Wu, Huan Ling, Mike Chen, Francis Fidler, Sanja and Williams, and Jiahui Huang. Scube: Instant large-scale scene reconstruction using voxplats. *NeurIPS*, 2024. 3
- [48] Barbara Rössle, Norman Müller, Lorenzo Porzi, Samuel Rota Bulò, Peter Kotschieder, Angela Dai, and Matthias Nießner. L3DG: Latent 3D Gaussian Diffusion. *SIGGRAPH Asia*, 2024. 3
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. *CVPR*, 2022. 7
- [50] 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. *CVPR*, 2022. 2
- [51] Tim Salimans and Jonathan Ho. Progressive Distillation for Fast Sampling of Diffusion Models. *ICLR*, 2022. 3, 4
- [52] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, et al. ZeroNVS: Zero-Shot 360-Degree View Synthesis from a Single Image. *CVPR*, 2024. 2
- [53] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *NeurIPS*, 2020. 3
- [54] Katja Schwarz, Norman Müller, and Peter Kotschieder. Generative gaussian splatting: Generating 3d scenes with video diffusion priors. *arXiv:2503.13272*, 2025. 3
- [55] Bradley Brown Seung Wook Kim, Kangxue Yin, Karsten Kreis, Katja Schwarz, Daiqing Li, Robin Rombach, Antonio Torralba, and Sanja Fidler. Neuralfield-ldm: Scene generation with hierarchical latent diffusion models. *CVPR*, 2023. 3
- [56] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. *arXiv:2308.16512*, 2023. 2

- [57] Jaidev Shriram, Alex Trevithick, Lingjie Liu, and Ravi Ramamoorthi. Realdreamer: Text-driven 3d scene generation with inpainting and depth diffusion. *3DV*, 2025. 7
- [58] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *ICLR*, 2021. 3
- [59] Wenqiang Sun, Shuo Chen, Fangfu Liu, Zilong Chen, Yueqi Duan, Jun Zhang, and Yikai Wang. Dimensionx: Create any 3d and 4d scenes from a single image with controllable video diffusion. *ICCV*, 2025. 2
- [60] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Viewset Diffusion: (0-)Image-Conditioned 3D Generative Models from 2D Data. *ICCV*, 2023. 3
- [61] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter Image: Ultra-Fast Single-View 3D Reconstruction. *CVPR*, 2024. 2, 3, 4
- [62] Stanislaw Szymanowicz, Eldar Insafutdinov, Chuanxia Zheng, Dylan Campbell, Joao Henriques, Christian Rupprecht, and Andrea Vedaldi. Flash3D: Feed-Forward Generalisable 3D Scene Reconstruction from a Single Image. *3DV*, 2025. 2, 6
- [63] Jiayang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. *ECCV*, 2024. 2
- [64] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree Generating Networks: Efficient Convolutional Architectures for High-resolution 3D Outputs. In *ICCV*, 2017. 2
- [65] Ayush Tewari, Tianwei Yin, George Cazenavette, Semon Rezhikov, Josh Tenenbaum, Frédo Durand, Bill Freeman, and Vincent Sitzmann. Diffusion with Forward Models: Solving Stochastic Inverse Problems Without Direct Supervision. *NeurIPS*, 2023. 3
- [66] Shubham Tulsiani, Tinghui Zhou, Alexei A Efros, and Jitendra Malik. Multi-view Supervision for Single-view Reconstruction via Differentiable Ray Consistency. In *CVPR*, 2017. 2
- [67] Shubham Tulsiani, Richard Tucker, and Noah Snavely. Layer-structured 3d scene inference via view synthesis. In *ECCV*, 2018. 2
- [68] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitrii Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: Novel Multi-view Synthesis and 3D Generation from a Single Image using Latent Video Diffusion. *ECCV*, 2024. 2
- [69] Matthew Wallingford, Anand Bhattad, Aditya Kusupati, Vivek Ramanujan, Matt Deitke, Aniruddha Kembhavi, Roozbeh Mottaghi, Wei-Chiu Ma, and Ali Farhadi. From an Image to a Scene: Learning to Imagine the World from a Million 360° Videos. *NeurIPS*, 2025. 2
- [70] Bo Wang, Yifan Zhang, Jian Li, Yang Yu, Zhenping Sun, Li Liu, and Dewen Hu. SplatFlow: Learning Multi-frame Optical Flow via Splatting. *IJCV*, 132(8):3023–3045, 2024. 3
- [71] 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. *CVPR*, 2021. 2
- [72] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. DUST3R: Geometric 3D Vision Made Easy. *CVPR*, 2024. 7
- [73] Daniel Watson, William Chan, Ricardo Martin-Brualla, Jonathan Ho, Andrea Tagliasacchi, and Mohammad Norouzi. Novel View Synthesis with Diffusion Models. *ICLR*, 2023. 2
- [74] Christopher Wewer, Kevin Raj, Eddy Ilg, Bernt Schiele, and Jan Eric Lenssen. latentsplat: Autoencoding variational gaussians for fast generalizable 3d reconstruction. In *ECCV*, 2024. 2, 3, 6
- [75] Chao-Yuan Wu, Justin Johnson, Jitendra Malik, Christoph Feichtenhofer, and Georgia Gkioxari. Multiview compressive coding for 3D reconstruction. *CVPR*, 2023. 2
- [76] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P. Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. ReconFusion: 3D Reconstruction with Diffusion Priors. *CVPR*, 2024. 2, 6
- [77] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3D Latents for Scalable and Versatile 3D Generation. *CVPR*, 2025. 2, 3
- [78] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. DepthSplat: Connecting Gaussian Splatting and Depth. *CVPR*, 2025. 2, 6
- [79] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, et al. DMV3D: Denoising Multi-view Diffusion Using 3D Large Reconstruction Model. *ICLR*, 2024. 2, 3
- [80] Botao Ye, Sifei Liu, Haofei Xu, Li Xueting, Marc Pollefeys, Ming-Hsuan Yang, and Peng Songyou. No Pose, No Problem: Surprisingly Simple 3D Gaussian Splats from Sparse Unposed Images. *ICLR*, 2025. 2
- [81] Xu Yinghao, Shi Zifan, Yifan Wang, Chen Hansheng, Yang Ceyuan, Peng Sida, Shen Yujun, and Wetzstein Gordon. Grm: Large gaussian reconstruction model for efficient 3d reconstruction and generation. *ECCV*, 2024. 2
- [82] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. PixelNeRF: Neural radiance fields from one or few images. *CVPR*, 2021. 2
- [83] Hong-Xing Yu, Haoyi Duan, Junhwa Hur, Kyle Sargent, Michael Rubinstein, William T Freeman, Forrester Cole, Deqing Sun, Noah Snavely, Jiajun Wu, et al. Wonderjourney: Going from Anywhere to Everywhere. In *CVPR*, 2024. 2
- [84] Hong-Xing Yu, Haoyi Duan, Charles Herrmann, William T Freeman, and Jiajun Wu. Wonderworld: Interactive 3D Scene Generation from a Single Image. *CVPR*, 2025. 2
- [85] Jason J. Yu, Fereshteh Forghani, Konstantinos G. Derpanis, and Marcus A. Brubaker. Long-term photometric consistent novel view synthesis with diffusion models. In *ICCV*, 2023. 2
- [86] Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. *CVPR*, 2021. 3

- [87] Wangbo Yu, Jinbo Xing, Li Yuan, Wenbo Hu, Xiaoyu Li, Zhipeng Huang, Xiangjun Gao, Tien-Tsin Wong, Ying Shan, and Yonghong Tian. ViewCrafter: Taming Video Diffusion Models for High-fidelity Novel View Synthesis. *arXiv:2409.02048*, 2024. 2
- [88] Xianggang Yu, Mutian Xu, Yidan Zhang, Haolin Liu, Chongjie Ye, Yushuang Wu, Zizheng Yan, Tianyou Liang, Guanying Chen, Shuguang Cui, and Xiaoguang Han. MVImgNet: A Large-scale Dataset of Multi-view Images. In *CVPR*, 2023. 5
- [89] Yang Yuanbo, Shao Jiahao, Li Xinyang, Shen Yujun, Geiger Andreas, and Liao Yiyi. Prometheus: 3d-aware latent diffusion models for feed-forward text-to-3d scene generation. *CVPR*, 2025. 3
- [90] Biao Zhang, Jiapeng Tang, Matthias Niessner, and Peter Wonka. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Trans. on Graphics (TOG)*, 42(4), 2023. 3
- [91] Kai Zhang, Sai Bi, Hao Tan, Yuanbo Xiangli, Nanxuan Zhao, Kalyan Sunkavalli, and Zexiang Xu. GS-LRM: Large Reconstruction Model for 3D Gaussian Splatting. *ECCV*, 2024. 2
- [92] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. CLAY: A Controllable Large-scale Generative Model for Creating High-quality 3D Assets. *ACM Trans. on Graphics (TOG)*, 2024. 2, 3
- [93] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *CVPR*, 2018. 5, 6
- [94] Shangzhan Zhang, Jianyuan Wang, Yinghao Xu, Nan Xue, Christian Rupprecht, Xiaowei Zhou, Yujun Shen, and Gordon Wetzstein. Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views. *CVPR*, 2025. 2
- [95] Jensen (Jinghao) Zhou, Hang Gao, Vikram Voleti, Aaryaman Vasishtha, Chun-Han Yao, Mark Boss, Philip Torr, Christian Rupprecht, and Varun Jampani. Stable virtual camera: Generative view synthesis with diffusion models. *arXiv:2503.14489*, 2025. 2
- [96] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Trans. Graph*, 37, 2018. 5
- [97] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. *CVPR*, 2023. 2