

VFusion3D: Learning Scalable 3D Generative Models from Video Diffusion Models

Junlin Han^{1,2*}, Filippos Kokkinos^{1*}, and Philip Torr²

¹ GenAI, Meta

² Torr Vision Group, University of Oxford

* Equal contribution

junlinhan@meta.com, fkokkinos@meta.com, philip.torr@eng.ox.ac.uk

Project page: <https://junlinhan.github.io/projects/vfusion3d.html>




Fig. 1: Comparisons of large 3D reconstruction models. Our VFusion3D reconstructs high-quality and 3D-consistent assets from a single input image.

Abstract. This paper presents a novel method for building scalable 3D generative models utilizing pre-trained video diffusion models. The primary obstacle in developing foundation 3D generative models is the limited availability of 3D data. Unlike images, texts, or videos, 3D data are not readily accessible and are difficult to acquire. This results in a significant disparity in scale compared to the vast quantities of other types of data. To address this issue, we propose using a video diffusion model, trained with extensive volumes of text, images, and videos, as a knowledge source for 3D data. By unlocking its multi-view generative capabilities through fine-tuning, we generate a large-scale synthetic multi-view dataset to train a feed-forward 3D generative model. The proposed model, VFusion3D, trained on nearly 3M synthetic multi-view data, can generate a 3D asset from a single image in seconds and achieves superior performance when compared to current SOTA feed-forward 3D generative models, with users preferring our results over 90% of the time.

1 Introduction

The advent of 3D datasets [7, 9, 10] and the development of advanced neural rendering methods [19, 36–38] have created new possibilities in computer vision and computer graphics. AI for 3D content creation has shown great potential in various fields, such as AR/VR/MR [22], 3D gaming [14, 48], and animation [20]. As a result, there is a growing demand for the development of foundation 3D generative models that can efficiently create various high-quality 3D assets. Despite the clear need and potential, current methods fall short of expectations due to their production of low-quality textures and geometric issues such as floaters.

The primary obstacle in constructing foundation 3D generative models is the scarcity of 3D data available. Acquiring 3D data is a complex and challenging task. Unlike images, which can be easily captured using standard cameras, 3D data requires specialized equipment and capturing techniques. Another option is 3D modeling, a tedious process that can take up to weeks for a single asset. This complexity results in a pool of publicly available 3D assets that is limited, with the largest ones containing up to 10 million assets [7, 9, 63]. In reality, only a small portion of them is usable since a decent percentage of 3D assets are duplicates and many of them lack texture or even a high-quality surface [16, 23, 47].

Conversely, foundation models such as GPT [1, 53, 54] and Diffusion models [3, 8, 45] have consistently showcased remarkable capabilities. These models have proven their prowess in handling intricate tasks, exhibiting advanced problem-solving skills, and delivering exceptional performance across a spectrum of challenges. Such general and strong capabilities are derived from scaling on both data and model size [3, 39, 54]. The primary prerequisite for constructing foundation models is the accessibility of an extensive amount of high-quality training data, often surpassing billions in size. This presents a clear discrepancy between the current state of 3D datasets and the requirements for training foundation 3D generative models.

In this paper, we propose to use of video diffusion models, trained on a large amount of texts, images, and videos, as a 3D data generator. Specifically, we use

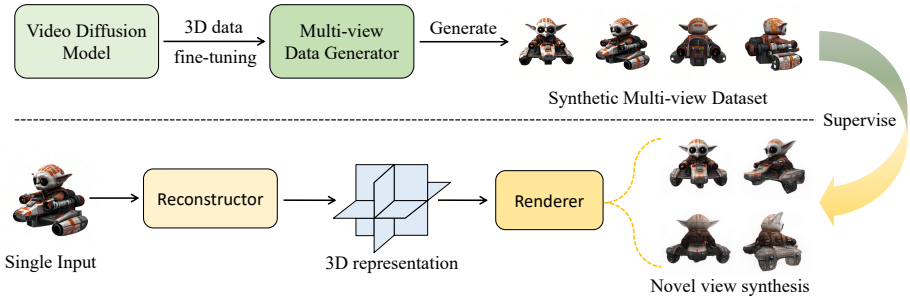


Fig. 2: The pipeline of VFusion3D. We first use a small amount of 3D data to fine-tune a video diffusion model, transforming it into a multi-view video generator that functions as a data engine. By generating a large amount of synthetic data, we train VFusion3D to generate a 3D representation and render novel views.

EMU Video [13] which has been trained with a variety of videos, including those with intricate camera movements and drone footage, and inherently contain cues about the 3D world. This suggests that video diffusion models have a latent understanding of how to generate videos with 3D consistency [34]. By fine-tuning this model using rendered multi-view videos from 100K 3D data, we unlock the model’s inherent capability to generate diverse, 3D-consistent, multi-view videos from text and image prompts. The resulting video diffusion model allows us to scale the 3D data necessary for learning foundation 3D generative models. Using millions of text prompts from web-scale data and a filtering system, we generate a multi-view dataset comprising of 3 million multi-view videos.

We utilize our synthetic multi-view dataset to train a 3D generative model capable of reconstructing 3D assets from single images in a feed-forward manner. Using the recently proposed Large Reconstruction Model (LRM) [16] as our starting point, we introduce a series of training strategies to stabilize the training process. These strategies assist the model in better learning from a substantial volume of synthetic multi-view training samples. Post-training, we further fine-tune with renderings from the 3D dataset, originally employed to fine-tune the video diffusion model, to further improve our 3D generative model. As a result, the proposed model, VFusion3D, can generate high-quality 3D assets from a single image with any viewing angles. We compare VFusion3D against several distillation-based and feed-forward 3D generative models using a user study and automated metrics. Finally, we undertake a comprehensive study to explore, analyze, and discuss topics on 3D data versus synthetic multi-view data, as well as the scaling trends of large 3D generative models.

2 Related Work

Text/Image-to-3D through Distillation or Reconstruction. Learning to generate 3D assets from textual descriptions or single images presents a formidable challenge due to the limited availability of 3D-text paired data. DreamField [17]

suggests using CLIP [43] as a starting point, but textual supervision alone is not enough for 3D generation. DreamFusion [40] optimizes an implicit 3D representation by distilling knowledge from 2D diffusion models [11, 45, 46] using a score-based distillation approach [40]. Such distillation-based approaches has shown promising results and inspired subsequent research [26, 31, 52, 57, 61]. Furthermore, the combination of 2D diffusion models with 3D data to achieve both 3D consistency and visual pleasing is increasingly gaining popularity [25, 28–30, 41, 47]. Concurrent research [34] suggests the reconstruction of 3D Gaussians [19] using multi-view outputs derived from fine-tuned video diffusion models [2, 21]. Such reconstruction-based methods [6, 27, 28, 58] can also leverage multi-view data produced by other methods [29, 47]. With respect to distillation, our work also involves the distillation of 3D knowledge from video diffusion models. However, we distill knowledge from video diffusion models in an explicit way, eliminating the need for score distillation sampling [40].

Feed-forward 3D Generative Models. In the fields of 3D reconstruction and 3D generation, a prominent trend involves learning directly from 3D datasets [12, 16, 31, 44, 49, 50, 59, 62]. Pioneering works [12, 18, 50] attempted to learn 3D representations directly from 3D [10] and multi-view [63] data, but these approaches faced limitations in performance due to low-quality training data or small-scale training data. A series of recent studies have suggested using a tri-plane [5] as the 3D representation, demonstrating that it can scale to larger datasets. These methods train large-scale models to generate a 3D representation from various inputs, including single images [16], textual information [23], posed multi-view images [60], and unposed multi-view images [56]. Typically, these large models [35] learn to predict a tri-plane representation, which is subsequently converted into a NeRF or combined with optimized point generator for Gaussian Splatting [65]. It is also possible to generate 3D Gaussians from multi-view inputs. Then, during inference time, it can employ off-the-shelf text-to-3D [55] or image-to-3D models [47] to generate multi-view inputs, as demonstrated in concurrent work [51]. In this study, we adopt the most general framework, *i.e.*, LRM [15, 16], without architectural modifications. Our objective is not to develop a new architecture. Instead, we strive to enhance the suitability and scalability of general models for training on large-scale synthetic multi-view images.

3 Method

Section 3.1 provides the preliminaries on the EMU Video [13] and LRM [16]. Our approach for transferring a video diffusion model into a 3D multi-view data engine is presented in Section 3.2. The process of gradually improving the LRM into our VFusion3D is described in Section 3.3. Figure 2 illustrates our pipeline.

3.1 Preliminaries

EMU Video. EMU Video [13] is a video diffusion model that builds upon a pre-trained text-to-image diffusion model, EMU [8]. It initializes its weights

from the EMU model and adds new learnable temporal parameters. EMU Video is conditioned on both a text prompt and an image prompt, where the image prompt serves as the first frame and can either be provided or generated. The EMU model is trained on 1.1 billion image-text pairs, and EMU Video is further trained on an additional 34 million video-text pairs. EMU Video is capable of generating high-quality and temporally consistent videos with up to 16 frames.

Large Reconstruction Model. LRM is a large feed-forward model for single image 3D reconstruction. It initially utilizes a pre-trained vision transformer, DINO [4], to encode image features. These features are subsequently used by an image-to-triplane decoder module via cross-attention mechanisms. Camera parameters are also sent to the image-to-triplane decoder after undergoing processing through a small camera embedding network. The tri-plane representation, predicted by the image-to-triplane decoder, is upsampled and reshaped to query 3D point features. These features are then input into a multi-layer perception module to predict RGB and density for volumetric rendering. The training is conducted on 3D multi-view datasets [10, 63] and is supervised using multi-view images with LPIPS loss [64] and MSE loss.

3.2 Video Diffusion Models as 3D Multi-view Data Generator

Prompt Collection and Filtering. We gather object-centric prompts from web-scale datasets and 3D captions. Specifically, we collect textual captions from a large text-image dataset and employ a Llama2-13B [54] with a specially designed prompting template. This template incorporates in-context examples to retain those captions that describe a single object. Additionally, we run Cap3D [33] with a Llama2-70B on multi-view images rendered from some 3D data to collect captions of 3D objects. In total, we collect 4M prompts.

EMU Video Fine-tuning. Our aim is to enable the pre-trained EMU Video model to generate multi-view videos, which will showcase a camera rotating around a 3D object. To accomplish this, we utilize an internal dataset of 100K 3D objects crafted by 3D modeling artists, akin to the 3D data collected in Objaverse [10]. For each asset, we render 16 views with a random elevation within the range of $[0, \pi/4]$. The camera is positioned at uniform intervals ($2\pi/16$ radians) around the object with a constant distance.

These rendered videos are used to fine-tune the EMU Video model. As all videos follow a similar camera trajectory, we do not send the camera parameters into the video diffusion model. Instead, we maintain a fixed camera distance and orientation, randomizing only the elevation. Despite the absence of explicit camera parameters, the first frame (the rendered image with 0 azimuth) can still exhibit varying viewing directions, enabling the EMU Video model to handle condition images with different viewing angles. The EMU Video model is fine-tuned to generate a sequence of views that follows this distribution. The fine-tuning of the EMU Video model is conducted using a standard diffusion model training loss. We freeze all parameters, except for the temporal convolutional and attention layers to ensure that the fine-tuning does not degrade the visual quality of generation.

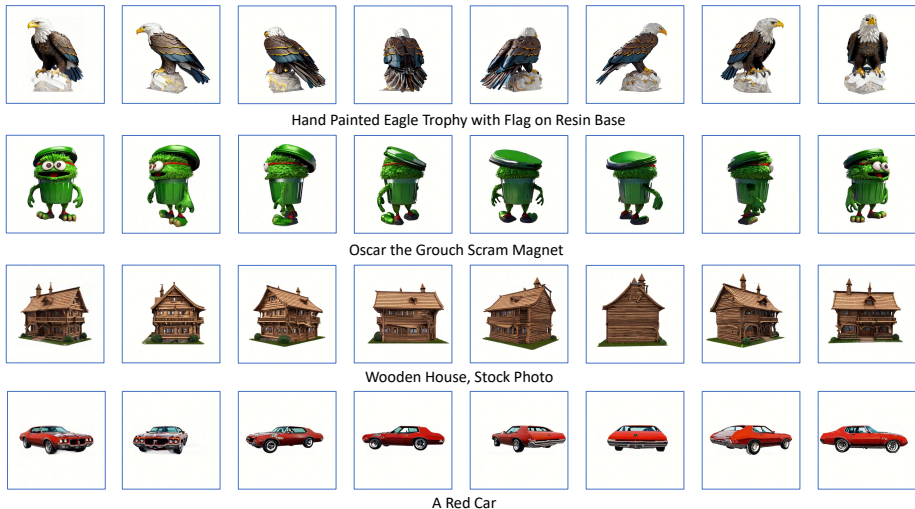


Fig. 3: Samples of multi-view sequences generated by fine-tuned video diffusion model. After fine-tuning with 100K 3D data, the video diffusion model can produce high-quality multi-view sequences, thus functioning as a multi-view data generation engine. The inputs are the initial image and the textual prompt. The last row shows a failure case, for which we have designed a filter to remove such data.

Post Processing and Metadata Preparation. We generated a total of 4 million videos, the majority of which demonstrate 3D-consistency and visual appeal. However, a subset exhibits lower quality and less 3D-consistency, prompting us to design a filter to selectively retain only those of superior quality. To achieve this, we manually labeled 2000 videos, with approximately 1200 classified as high quality. We utilized DINO [4] to extract features from 8 frames, uniformly distributed throughout the video. We subsequently trained a Support Vector Machine to classify video quality based on the averaged DINO features. As a result, we retained 2.7M videos for our synthetic dataset. Examples of generated multi-view sequences and data filtering are shown in Figure 3.

The generated multi-view videos follow a similar camera trajectory to cover a 360° azimuth range, but the elevation depends on the generated condition image, which is less controllable. Therefore, it is necessary to label the elevation for multi-view videos. We used 100K 3D data with ground truth elevation to train our elevation estimator. Similar to the filtering process, we used DINO to extract features from 4 uniform views and averaged their features. A 2-layer MLP was trained on top of this as an elevation predictor. We then use it to label the elevation of all generated multi-view videos.

3.3 VFusion3D

VFusion3D adheres to the architecture of LRM, which is considered the most general method yet for feed-forward 3D creation. We pinpoint the ideal training

strategies and formulate a recipe that enhances suitability and scalability when learning from synthetic multi-view data. Our ultimate objective is to establish learning from video diffusion generated multi-view data as a potential paradigm for training foundation 3D models.

Improved Training Strategies for Synthetic Data. The original training setting of LRM was designed to work with Objaverse [10] and MVImageNet [63] data. However, we found that this setup was not entirely suitable for direct application to our synthetic data, as synthetic data tends to be noisier and does not always maintain full 3D consistency. To stabilize and improve the learning process on synthetic multi-view data, we implemented a series of strategies.

- ***A Multi-stage Training Recipe.*** LRM employs a patch rendering strategy, which renders a small 128×128 resolution patch from a random rendering resolution that ranges from 128 to 384. However, this setting can potentially cause instability in the learning process, particularly in the early stages, and may lead to incorrect optimization directions, resulting in predictions with checkerboard effects. To address this, we suggest to use a multi-stage learning process. This involves gradually increasing the rendering resolution, thereby preparing the model for higher rendering resolutions at each stage. In practice, we use rendering resolutions of 128, 192, 256, 320, and 384. Each stage is trained for 5 epochs, with the exception of the final stage, which is trained for 10 epochs. After training on the 192 resolution stage, we reset the learning rate to the initial learning rate. This adjustment allows the model to better learn from later stages.

- ***Image-level Supervision Instead of Pixel-level Supervision.*** Pixel-level losses, such as L1 and L2, are commonly used. However, they may not be suitable for synthetic multi-view data due to they are strict in pixel-level correspondence. Minor inconsistencies in synthetic multi-view images are inevitable, and when trained with pixel-level loss, these can lead to improper optimization and blurry results. On the other hand, image-level loss methods, such as LPIPS [64], are less stringent as they operate in feature spaces. This inherent flexibility has the potential to strengthen our model, as it accommodates minor inconsistencies in synthetic multi-view images without compromising the optimization process.

- ***Opacity Loss.*** In synthetic multi-view data, small colored patches occasionally appear in images that should have a white background. To address this, we run a saliency detection model [42] to obtain the masks of central objects in our synthetic data. This allows us to apply an opacity loss, which helps maintain focus on the foreground object. Additionally, extra opacity supervision can facilitate the training process, resulting in stronger models. We employ a general method to extract these masks and introduce an opacity loss between rendering density and masks. Combined with our strategy on no pixel-level losses, our approaches efficiently mitigate noises from small color patches in the background.

- ***Camera Noise Injection.*** For object observation, the camera trajectory of our synthetic multi-view data follows a fixed sequence of azimuth changes. This could potentially restrict the model’s generalization capability due to the limited viewing positions. To counteract this, we introduce camera noises during both the 3D data rendering and VFusion3D training stages. We apply a random




Fig. 4: Qualitative results on single image 3D reconstruction. VFusion3D successfully reconstructs 3D objects with strong 3D consistency in both shape and color. In contrast, OpenLRM sometimes fails to infer a reasonable shape, and LGM alters the color of unseen parts.

Method	CLIP Text Similarity \uparrow	CLIP Image Similarity \uparrow
OpenLRM	0.234	0.793
LGM	0.241	0.796
VFusion3D	0.253	0.851

Table 1: Comparative performance in image-to-3D tasks. Overall, VFusion3D exhibits better performance in both text alignment and visual alignment.

minor offset (ranging from -0.05 to 0.05) during the rendering of multi-view images, which are subsequently used to fine-tune the EMU Video. Additionally, we infuse camera noise (varying from -0.02 to 0.02) into both the intrinsic and extrinsic matrices during the training process. This potentially enhances the model’s robustness against incorrect camera parameters and improves its ability to generalize to different viewing angles.

Fine-tuning with 3D Data. Inspired by EMU [8], which employs a small number of visually striking images to fine-tune a pre-trained text-to-image diffusion model with a limited number of batches and iterations. We consider 3D data as visually striking and hypothesize that our model, which is pre-trained with synthetic multi-view sequences, could also benefit from exposure to some 3D samples. Following this procedure, we re-use the 100K data that was used in fine-tuning the EMU Video to fine-tune the VFusion3D model, which was pre-trained on synthetic data. We demonstrate that a combination of synthetic multi-view data and 3D data leads to the best model.




Fig. 5: User study results comparing VFusion3D and previous work. VFusion3D consistently outperforms previous works by considerable margins in both generation quality and image faithfulness.

4 Experiments

4.1 Implementation Details

Our VFusion3D model is trained on 128 NVIDIA A100 (80G) GPUs, and the training process takes approximately 6 days to complete. We use a total batch size of 1024, with 4 multi-view images at a resolution of 128×128 used for supervision per batch, resulting in a total of 4096 multi-view images. The input image has a resolution of 512×512 . The model is trained for 30 epochs with an initial learning rate of 4×10^{-4} , following a cosine annealing schedule with a restart after first 10 epochs. The training begins with a warm-up of 3000 iterations, and the optimizer used is AdamW [32]. More details are presented in the supplementary material.

4.2 Results and Comparisons

Single Image 3D Reconstruction. We benchmark our model against recent large feed-forward methods, including OpenLRM-large [15, 16] and LGM [51]. For evaluation, we collect 25 diverse images that vary in style and shape. We employ BLIP-2 [24] to obtain text captions. All results are rendered at a resolution of 384×384 . We report both the CLIP text similarity scores and CLIP image similarity scores [43]. Quantitative results are displayed in Table 1 and qualitative results are presented in Figure 4. Overall, VFusion3D exhibits superior performance in both text and image alignment, and it typically generates more visually appealing results with high image faithfulness.

Text-to-3D Generation. For text-to-3D generation, we employ EMU to generate images from text prompts, supporting text-image-3D reconstruction for our method. We use a common test set for evaluation, which contains 40 prompts from MVDream [47]. We compare VFusion3D with DreamFusion [40], Magic3D [26], ProlificDreamer [57], MVDream [47], OpenLRM [15, 16], and



Fig. 6: Visual results of text-to-3D generation. Among various methods, VFusion3D demonstrates strong performance. When compared with feed-forward reconstruction models, VFusion3D shows better image fidelity and 3D consistency.

Method	CLIP Text Similarity \uparrow	CLIP Image Similarity \uparrow
DreamFusion	0.261	0.640
Magic3D	0.293	0.687
ProlificDreamer	0.293	0.699
MVDream	0.284	0.688
OpenLRM	0.255	0.826
LGM	0.270	0.832
VFusion3D	0.272	0.899

Table 2: Comparisons of text-to-3D generation methods. Text-to-3D methods typically exhibit stronger text alignment, whereas image-to-3D models often demonstrate better image alignment. VFusion3D yields the highest image similarity scores.

LGM [51]. For text-to-3D models, we use rendered videos from the MVDream website for evaluation. Results are presented in Table 2, and visual samples are shown in Figure 6. When considering feed-forward models, VFusion3D exhibits stronger image and text alignment. Qualitatively, VFusion3D shows superior results compared to other feed-forward models, especially in terms of both color and 3D shape consistency.

Num of Data	Generated Videos Good Rate \uparrow	Trained Models			
		SSIM \uparrow	LPIPS \downarrow	Text Sim \uparrow	Image Sim \uparrow
10K	44.5%	0.822	0.185	0.237	0.721
50K	52.1%	0.823	0.184	0.247	0.753
100K	61.3%	0.824	0.182	0.250	0.759

Table 3: Ablation study on the 3D data required for fine-tuning EMU video. By utilizing more 3D data for fine-tuning EMU video, the ability to generate multi-view sequences in fine-tuned EMU video is improved.

User Study. To assess the overall quality of the generated content and its faithfulness to the input images more accurately, we conducted a user study via Amazon Mechanical Turk. We presented users with two 360° rendered videos - one produced by our VFusion3D model and another by a baseline model, and asked them to indicate a winner. A total of 65 videos (25 for 3D reconstruction, 40 for text-to-3D) were evaluated. Feedback was collected from 5 different users and we present the majority results in the form of a winner rate. Figure 5 displays our results. Our method surpasses SOTA baselines, demonstrating that VFusion3D aligns closely with the original image content and exhibits high visual quality.

4.3 Ablation Study

We conduct an ablation study on several design choices, which include: (1) the specifications for fine-tuning the EMU Video model, (2) the training strategies proposed in VFusion3D, and (3) the number of 3D data for fine-tuning pre-trained VFusion3D. We use two evaluation sets for our study. First, we report SSIM and LPIPS for 3D data evaluation, along with CLIP text similarity score and CLIP image similarity score for text-to-3D evaluation. The 3D data evaluations were conducted on a test dataset comprising 500 unseen 3D models. From different viewing angles, we use one image as input, render 32 novel views, and compare them against their respective ground truths. The text-to-3D evaluation follows the setting of text-to-3D experiments as in section 4.2. The VFusion3D model used in the ablation study 4.3 and analysis 4.4 is a variant that processes input images with a resolution of 224×224 .

How Much 3D Data are Needed for Fine-tuning Video Diffusion Models? We explore the amount of 3D data required for fine-tuning EMU Video. Specifically, we experiment with 10K, 50K, and 100K 3D data. Table 3 presents the results of both the good classification rate of multi-view sequences generated by the fine-tuned EMU Video and the performance of VFusion3D models trained on the synthetic datasets from the respective fine-tuned EMU Video. We generate 500K multi-view sequences before filtering for training VFusion3D variants. As expected, using more 3D data enhances the multi-view sequence generation ability of EMU Video without compromising visual quality. This suggests that our approach could also scale with the collection of more 3D data.

Effect of Improved Training Strategies. We evaluate the impact of the proposed training strategies incorporated in VFusion3D. We use 500K synthetic

Components	SSIM \uparrow	LPIPS \downarrow	CLIP Text Sim \uparrow	CLIP Image Sim \uparrow
Baseline	0.826	0.206	0.223	0.712
+ multi-stage training	0.829	0.168	0.249	0.801
+ no pixel-level loss	0.831	0.167	0.257	0.798
+ opacity supervision	0.831	0.167	0.256	0.802
+ camera noise	0.830	0.169	0.252	0.800

Table 4: Ablation study on improved training strategies. We sequentially incorporate proposed strategies for synthetic data, and the effect of each is validated.

Num of data	SSIM \uparrow	LPIPS \downarrow	CLIP Text Sim \uparrow	CLIP Image Sim \uparrow
No fine-tune	0.832	0.160	0.261	0.839
10K	0.835	0.153	0.261	0.838
30K	0.837	0.149	0.262	0.846
50K	0.839	0.147	0.262	0.834
50K (random)	0.840	0.146	0.261	0.836
50K (16-views)	0.838	0.147	0.262	0.835
100K	0.842	0.143	0.266	0.836

Table 5: Ablation study on settings of 3D data fine-tuning. In general, the use of more 3D data leads to improved performance, irrespective of other settings.

data only for training ablation variants. The results of these variants are presented in Table 4.

Our findings include: (1) The multi-stage training recipe significantly improves the results by stabilizing the training process. (2) Without image-level supervision, such as MSE, we notice a small improvement across all metrics, as well as in the visual results. (3) The inclusion of opacity loss can further enhance performance. (4) While the injection of camera noises does not improve performance according to metrics, further testing reveals that this variant does enhance the model’s robustness. It enables the model to handle a wider range of potentially inaccurate camera matrices. Therefore, the decision to apply it becomes a matter of trade-off.

Settings of 3D Data Fine-tuning. We investigate the settings of 3D data fine-tuning, including sizes, data selections, and number of views. We evaluate datasets of varying sizes - 10K, 30K, 50K, and 100K - where the data is selected based on aesthetic score rankings, from high to low. To determine the aesthetic scores of 3D data, we train an aesthetic score predictor on top of a pre-trained 2D aesthetic score predictor. This predictor uses CLIP to extract features and an MLP for prediction. We average the CLIP features extracted from multi-view images to make it applicable to 3D data. To ascertain whether aesthetic score ranking is beneficial, we further draw 50K 3D data randomly from the 100k 3D data as a variant. By default, all 3D data are rendered to 32 views, but we also explore a variant that only uses 16 views.

Table 5 presents the quantitative results. The conclusions drawn are as follows: (1) Fine-tuning with more 3D data results in stronger models. (2) Unlike 2D images, aesthetic scores are unnecessary in selecting 3D training data. (3)

Data	SSIM \uparrow	LPIPS \downarrow	CLIP Text Sim \uparrow	CLIP Image Sim \uparrow
3D data	0.839	0.161	0.205	0.631
Synthetic data	0.832	0.160	0.261	0.839
Both 3D data	0.842	0.143	0.266	0.836

Table 6: Analysis on 3D data vs. Synthetic Multi-view Data. 3D data is more efficient than synthetic data when it comes to learning to reconstruct common objects. However, large-scale synthetic multi-view data offers the ability to generalize to unusual objects and scenes. Combining the two can yield the best performance.

Even with only 16 views, VFusion3D, when trained with a large number of synthetic data, already learns how to interpolate effectively between different viewing angles. Fine-tuning with 32 views only leads to marginal improvements.

4.4 Analysis and Discussion

In this analysis, our objective is to delve into the intricacies of synthetic multi-view data, exploring their benefits, limitations, and potential in comparison to 3D data. Given our capacity to generate a significant amount of synthetic multi-view data, which supports the data requirements for training foundation 3D generative models, we also present scaling trends of 3D generative models. The evaluation protocol in this section adheres to the ablation study 4.3.

3D Data vs. synthetic multi-view Data. We aim to investigate the performance of models trained on 3D data, synthetic data, and a combination of both (pre-training with synthetic, then fine-tuning with 3D data). For 3D data, we train a VFusion3D model using 150K 3D data from our internal dataset. This includes the same 100K 3D data that were used in the fine-tuning process of the EMU Video. Table 6 presents the quantitative results. It shows that 3D data is more efficient than synthetic multi-view data in teaching the model to reconstruct common objects. Training with 100K 3D data points already matches the performance of 2.7M synthetic data points. However, learning with a limited number of data cannot generalize to uncommon objects or scenes, where large-scale synthetic data provides strong performance in generalization. An additional advantage of synthetic data is that it can be further combined with 3D data fine-tuning to achieve optimal performance.

Scaling Trends for 3D Generative Models. By maintaining a fixed model architecture, we examine the impact of varying training dataset sizes, ranging from 100K, 300K, 500K, 1M, to 2.7M. Our objective is to illustrate scaling trends that shed light on scalable 3D generative modeling. Figure 7 presents these trends. We observe that the generation quality, as measured by LPIPS and CLIP image similarity scores, consistently improves with the size of the synthetic dataset. Given our ability to generate an unlimited amount of synthetic data, this makes our approach highly scalable.

Furthermore, our approach can also scale and improve with several other factors. These include the development of stronger video diffusion models, the availability of more 3D data for fine-tuning the video diffusion model and the




Fig. 7: Scaling trends of VFusion3D on synthetic data. The left and right figures display the LPIPS and CLIP image similarity scores in relation to the dataset size, respectively. The generation quality consistently improves as the dataset size increases.

pre-trained 3D generative model, and the advancement of large 3D feed-forward generative models. All these factors contribute to the scalability of our model, positioning it as a promising avenue for foundation 3D generative models.

Limitations. The fine-tuned video diffusion model is less effective at generating multi-view sequences of specific objects, such as vehicles like cars, bicycles, and motorcycles, and text-related content. Our filtering system excludes most of these less-than-ideal results; however, this could potentially affect the performance of the trained VFusion3D model due to an insufficient amount of data related to vehicles and texts. This is a shortcoming inherited from the pre-trained video diffusion model. With the development of progressively stronger video diffusion models, this limitation should be mitigated.

5 Conclusion

This paper leverages a video diffusion model as a multi-view data generator, thereby facilitating the learning of scalable 3D generative models. We have showcased the potential of video diffusion models to function as a multi-view data engine, capable of generating an even infinite scale of synthetic data to support scalable training. Our proposed model, VFusion3D, which learns from synthetic data, has shown superior performance in the generation of 3D assets. Beyond its current state, VFusion3D is highly scalable and can scale with both the number of synthetic data and 3D data, paving new paths for 3D generative models.

Acknowledgements

Junlin Han is supported by Meta. We would like to thank Jianyuan Wang, Luke Melas-Kyriazi, Yawar Siddiqui, Quankai Gao, Yanir Kleiman, Roman Shapovalov, Natalia Neverova, and Andrea Vedaldi for the insightful discussions and invaluable support.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023) **2**
2. Blattmann, A., Dockhorn, T., Kulal, S., Mendelevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023) **4**
3. Brooks, T., Peebles, B., Homes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C.W.Y., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/research/video-generation-models-as-world-simulators> **2**
4. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021) **5, 6**
5. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: CVPR (2022) **4**
6. Chan, E.R., Nagano, K., Chan, M.A., Bergman, A.W., Park, J.J., Levy, A., Aittala, M., De Mello, S., Karras, T., Wetzstein, G.: Genvs: Generative novel view synthesis with 3d-aware diffusion models (2023) **4**
7. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015) **2**
8. Dai, X., Hou, J., Ma, C.Y., Tsai, S., Wang, J., Wang, R., Zhang, P., Vandenhende, S., Wang, X., Dubey, A., et al.: Emu: Enhancing image generation models using photogenic needles in a haystack. arXiv preprint arXiv:2309.15807 (2023) **2, 4, 8**
9. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., et al.: Objaverse-xl: A universe of 10m+ 3d objects. Advances in Neural Information Processing Systems **36** (2024) **2**
10. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13142–13153 (2023) **2, 4, 5, 7**
11. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. NeurIPS (2021) **4**
12. Erkoç, Z., Ma, F., Shan, Q., Nießner, M., Dai, A.: Hyperdiffusion: Generating implicit neural fields with weight-space diffusion. arXiv preprint arXiv:2303.17015 (2023) **4**
13. Girdhar, R., Singh, M., Brown, A., Duval, Q., Azadi, S., Rambhatla, S.S., Shah, A., Yin, X., Parikh, D., Misra, I.: Emu video: Factorizing text-to-video generation by explicit image conditioning. arXiv preprint arXiv:2311.10709 (2023) **3, 4**
14. Hao, Z., Mallya, A., Belongie, S., Liu, M.Y.: Gancraft: Unsupervised 3d neural rendering of minecraft worlds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14072–14082 (2021) **2**
15. He, Z., Wang, T.: Openlrn: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM> (2023) **4, 9**
16. Hong, Y., Zhang, K., Gu, J., Bi, S., Zhou, Y., Liu, D., Liu, F., Sunkavalli, K., Bui, T., Tan, H.: Lrm: Large reconstruction model for single image to 3d. ICLR (2024) **2, 3, 4, 9**

17. Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: CVPR (2022) [3](#)
18. Jun, H., Nichol, A.: Shap-e: Generating conditional 3d implicit functions. arXiv preprint arXiv:2305.02463 (2023) [4](#)
19. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (2023) [2](#), [4](#)
20. Kolotouros, N., Alldieck, T., Zanfir, A., Bazavan, E., Fieraru, M., Sminchisescu, C.: Dreamhuman: Animatable 3d avatars from text. Advances in Neural Information Processing Systems **36** (2024) [2](#)
21. Kwak, J.g., Dong, E., Jin, Y., Ko, H., Mahajan, S., Yi, K.M.: Vivid-1-to-3: Novel view synthesis with video diffusion models. arXiv preprint arXiv:2312.01305 (2023) [4](#)
22. Li, C., Li, S., Zhao, Y., Zhu, W., Lin, Y.: Rt-nerf: Real-time on-device neural radiance fields towards immersive ar/vr rendering. In: Proceedings of the 41st IEEE/ACM International Conference on Computer-Aided Design. pp. 1–9 (2022) [2](#)
23. Li, J., Tan, H., Zhang, K., Xu, Z., Luan, F., Xu, Y., Hong, Y., Sunkavalli, K., Shakhnarovich, G., Bi, S.: Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. ICLR (2024) [2](#), [4](#)
24. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. arXiv preprint arXiv:2301.12597 (2023) [9](#)
25. Li, W., Chen, R., Chen, X., Tan, P.: Sweetdreamer: Aligning geometric priors in 2d diffusion for consistent text-to-3d. arXiv preprint arXiv:2310.02596 (2023) [4](#)
26. Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3D: High-resolution text-to-3d content creation. arXiv preprint arXiv:2211.10440 (2022) [4](#), [9](#)
27. Liu, M., Shi, R., Chen, L., Zhang, Z., Xu, C., Wei, X., Chen, H., Zeng, C., Gu, J., Su, H.: One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. arXiv preprint arXiv:2311.07885 (2023) [4](#)
28. Liu, M., Xu, C., Jin, H., Chen, L., Xu, Z., Su, H., et al.: One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. arXiv preprint arXiv:2306.16928 (2023) [4](#)
29. Liu, R., Wu, R., Van Hoorick, B., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1-to-3: Zero-shot one image to 3d object. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9298–9309 (2023) [4](#)
30. Liu, Y., Lin, C., Zeng, Z., Long, X., Liu, L., Komura, T., Wang, W.: Syncdreamer: Generating multiview-consistent images from a single-view image. arXiv preprint arXiv:2309.03453 (2023) [4](#)
31. Lorraine, J., Xie, K., Zeng, X., Lin, C.H., Takikawa, T., Sharp, N., Lin, T.Y., Liu, M.Y., Fidler, S., Lucas, J.: Att3d: Amortized text-to-3d object synthesis. arXiv preprint arXiv:2306.07349 (2023) [4](#)
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017) [9](#)
33. Luo, T., Rockwell, C., Lee, H., Johnson, J.: Scalable 3d captioning with pretrained models. arXiv preprint arXiv:2306.07279 (2023) [5](#)
34. Melas-Kyriazi, L., Laina, I., Rupprecht, C., Neverova, N., Vedaldi, A., Gafni, O., Kokkinos, F.: Im-3d: Iterative multiview diffusion and reconstruction for high-quality 3d generation. arXiv preprint arXiv:2402.08682 (2024) [3](#), [4](#)

35. Mercier, A., Nakhli, R., Reddy, M., Yasarla, R., Cai, H., Porikli, F., Berger, G.: Hexagen3d: Stablediffusion is just one step away from fast and diverse text-to-3d generation. arXiv preprint arXiv:2401.07727 (2024) [4](#)
36. Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4460–4470 (2019) [2](#)
37. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99–106 (2021) [2](#)
38. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019) [2](#)
39. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4195–4205 (2023) [2](#)
40. Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: DreamFusion: Text-to-3d using 2d diffusion. arXiv (2022) [4](#), [9](#)
41. Qian, G., Mai, J., Hamdi, A., Ren, J., Siarohin, A., Li, B., Lee, H.Y., Skorokhodov, I., Wonka, P., Tulyakov, S., et al.: Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. arXiv preprint arXiv:2306.17843 (2023) [4](#)
42. Qin, X., Zhang, Z., Huang, C., Dehghan, M., Zaiane, O., Jagersand, M.: U2-net: Going deeper with nested u-structure for salient object detection. vol. 106, p. 107404 (2020) [7](#)
43. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021) [4](#), [9](#)
44. Ren, X., Huang, J., Zeng, X., Museth, K., Fidler, S., Williams, F.: Xcube (x3): Large-scale 3d generative modeling using sparse voxel hierarchies. arXiv preprint arXiv:2312.03806 (2023) [4](#)
45. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022) [2](#), [4](#)
46. Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E., Ghasemipour, S.K.S., Ayan, B.K., Mahdavi, S.S., Lopes, R.G., Salimans, T., Ho, J., Fleet, D.J., Norouzi, M.: Photorealistic text-to-image diffusion models with deep language understanding. arXiv preprint arXiv:2205.11487 (2022) [4](#)
47. Shi, Y., Wang, P., Ye, J., Long, M., Li, K., Yang, X.: Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512 (2023) [2](#), [4](#), [9](#)
48. Sun, C., Han, J., Deng, W., Wang, X., Qin, Z., Gould, S.: 3d-gpt: Procedural 3d modeling with large language models. arXiv preprint arXiv:2310.12945 (2023) [2](#)
49. Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Splatter image: Ultra-fast single-view 3d reconstruction. arXiv preprint arXiv:2312.13150 (2023) [4](#)
50. Szymanowicz, S., Rupprecht, C., Vedaldi, A.: Viewset diffusion:(0-) image-conditioned 3d generative models from 2d data. arXiv preprint arXiv:2306.07881 (2023) [4](#)

51. Tang, J., Chen, Z., Chen, X., Wang, T., Zeng, G., Liu, Z.: Lgm: Large multi-view gaussian model for high-resolution 3d content creation. arXiv preprint arXiv:2402.05054 (2024) [4](#), [9](#), [10](#)
52. Tang, J., Ren, J., Zhou, H., Liu, Z., Zeng, G.: Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. arXiv preprint arXiv:2309.16653 (2023) [4](#)
53. Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., et al.: Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805 (2023) [2](#)
54. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023) [2](#), [5](#)
55. Wang, P., Shi, Y.: Imagedream: Image-prompt multi-view diffusion for 3d generation. arXiv preprint arXiv:2312.02201 (2023) [4](#)
56. Wang, P., Tan, H., Bi, S., Xu, Y., Luan, F., Sunkavalli, K., Wang, W., Xu, Z., Zhang, K.: Pf-lrm: Pose-free large reconstruction model for joint pose and shape prediction. ICLR (2024) [4](#)
57. Wang, Z., Lu, C., Wang, Y., Bao, F., Li, C., Su, H., Zhu, J.: Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. arXiv preprint arXiv:2305.16213 (2023) [4](#), [9](#)
58. Wu, R., Mildenhall, B., Henzler, P., Park, K., Gao, R., Watson, D., Srinivasan, P.P., Verbin, D., Barron, J.T., Poole, B., et al.: Reconfusion: 3d reconstruction with diffusion priors. arXiv preprint arXiv:2312.02981 (2023) [4](#)
59. Xu, D., Yuan, Y., Mardani, M., Liu, S., Song, J., Wang, Z., Vahdat, A.: Agg: Amortized generative 3d gaussians for single image to 3d. arXiv preprint arXiv:2401.04099 (2024) [4](#)
60. Xu, Y., Tan, H., Luan, F., Bi, S., Wang, P., Li, J., Shi, Z., Sunkavalli, K., Wetstein, G., Xu, Z., et al.: Dmv3d: Denoising multi-view diffusion using 3d large reconstruction model. ICLR (2024) [4](#)
61. Yi, T., Fang, J., Wang, J., Wu, G., Xie, L., Zhang, X., Liu, W., Tian, Q., Wang, X.: Gaussiandreamer: Fast generation from text to 3d gaussians by bridging 2d and 3d diffusion models. arXiv preprint arXiv:2310.08529 (2023) [4](#)
62. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021) [4](#)
63. Yu, X., Xu, M., Zhang, Y., Liu, H., Ye, C., Wu, Y., Yan, Z., Zhu, C., Xiong, Z., Liang, T., et al.: Mvimnet: A large-scale dataset of multi-view images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9150–9161 (2023) [2](#), [4](#), [5](#), [7](#)
64. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) [5](#), [7](#)
65. Zou, Z.X., Yu, Z., Guo, Y.C., Li, Y., Liang, D., Cao, Y.P., Zhang, S.H.: Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. arXiv preprint arXiv:2312.09147 (2023) [4](#)