



X-Ray: A Sequential 3D Representation For Generation

Tao Hu¹ Wenhang Ge^{2*} Yuyang Zhao^{1*} Gim Hee Lee¹

¹ Department of Computer Science, National University of Singapore

² Hong Kong University of Science and Technology (Guangzhou)

taohu@nus.edu.sg

gimhee.lee@nus.edu.sg

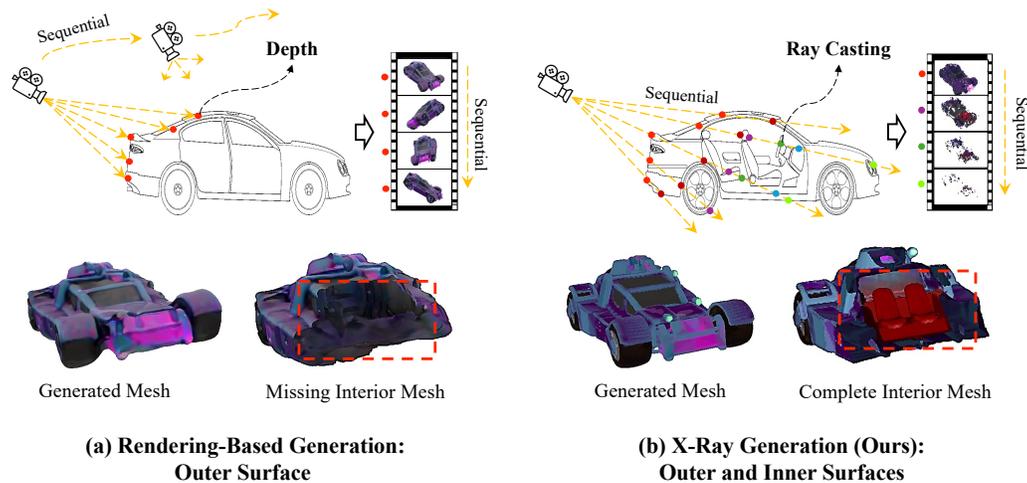


Figure 1: Comparison between the rendering-based 3D generation [49, 14] and our proposed X-Ray generation [49, 14]. The competitors focus on the visible outer surface in multiple camera views. In contrast, our model can sense both the visible and hidden surface in single camera view and generate the outer and inner surfaces of objects. An example of missing mesh interior from rendering-based 3D synthesis vs. complete mesh interior from our X-Ray generator are shown in the bottom row.

Abstract

We introduce X-Ray, a novel 3D sequential representation inspired by the penetrability of x-ray scans. X-Ray transforms a 3D object into a series of surface frames at different layers, making it suitable for generating 3D models from images. Our method utilizes ray casting from the camera center to capture geometric and textured details, including depth, normal, and color, across all intersected surfaces. This process efficiently condenses the whole 3D object into a multi-frame video format, motivating the utilization of a network architecture similar to those in video diffusion models. This design ensures an efficient 3D representation by focusing solely on surface information. Also, we propose a two-stage pipeline to generate 3D objects from X-Ray Diffusion Model and Upsampler. We

*Co-second authors.

demonstrate the practicality and adaptability of our X-Ray representation by synthesizing the complete visible and hidden surfaces of a 3D object from a single input image. Experimental results reveal the state-of-the-art superiority of our representation in enhancing the accuracy of 3D generation, paving the way for new 3D representation research and practical applications. Our project page is in <https://tau-yihouxiang.github.io/projects/X-Ray/X-Ray.html>.

1 Introduction

General, accurate, and efficient 3D representations are three of the most critical requirements for 3D generation [23, 29, 32, 21]. The significance of this goal stems from the ever-expanding array of applications reliant on 3D technology, ranging from virtual reality and augmented reality to computer-aided design and beyond. Previous approaches to 3D representation such as meshes, point clouds, voxels, Neural Radiance Fields (NeRF) [36, 45, 58, 24, 15] and 3D Gaussian Splatting [21] possess unique strengths respectively, but face challenges in concurrently satisfying the three requirements for 3D synthesis. Specifically, meshes are widely used in 3D modeling, while they are constrained by their topology when describing complex objects, which limits their generative capacity. Point clouds offer a more flexible capture of the object geometries but lack continuous and dense feature extraction [12, 42]. Voxels simplify spatial reasoning at the cost of significant rising memory complexity with increasing resolution. Neural representations, such as NeRF [36] and 3D Gaussian Splatting [21], offer an impressive leap in rendering photorealistic scenes. Nevertheless, the 3D object are predicted by multi-view images with a relatively long optimized period.

Recently, rendering-based 3D generative methods [26, 57, 48, 14, 53, 49, 18] have gained significant attention for their ability to achieve general and even efficient 3D generation by incorporating neural representations [36, 21] with Transformers [52] or Diffusion Models [43, 2]. However, these methods have a critical limitation: they cannot completely generate objects that include both visible and hidden surfaces. This limitation arises from the methodological design of rendering-based 3D generation, which relies on the 2D supervision of rendered images. As a result, these methods primarily focus on reconstructing the visible external surfaces of objects, while neglecting the internal hidden surfaces. This oversight leads to incomplete or unrealistic reconstructions, as shown in Fig. 1 (a).

In this paper, we propose the X-Ray representation to overcome this limitation of incomplete generation while maintaining the efficiency and generalization required for 3D generation. As illustrated in Fig. 1 (b), our X-Ray, inspired by x-ray imaging in the medical field, can see through the entire object. It efficiently captures and stores information about both visible and hidden surfaces. Consequently, the hidden interior of the object can be fully reconstructed.

Our X-Ray is designed to capture the shape (depth and normal) and appearance attributes (color) along all the sequentially intersected surfaces through ray casting. We transpose the collected slim grid voxels into a multi-frame surface representation, which significantly reduces the data footprint while preserving essential detail. Moreover, the compatibility of our X-Ray's data structure with sequential 3D representation in video formats, as illustrated in Fig. 2, opens novel pathways for leveraging video diffusion models in 3D generation. Specifically, by treating our X-Ray representation as sequences of frames, we first harness the power of the video diffusion model [2] and then utilize the video upsampler [7] to generate 3D objects from low to high resolution. As a result, our approach yields high-quality results while inheriting the advanced capabilities and efficiency of video processing.

We demonstrate the advantages of our X-Ray representation through comprehensive experiments, showcasing its superiority, especially completeness in 3D object generation. We train and evaluate our method in image-to-3D reconstruction task and pure 3D generation task. The experimental results reveal that the proposed X-Ray achieves a significant leap forward in the quality of 3D object generation, positioning it as a feasible solution to longstanding challenges in the field. The main contribution of the paper can be summarized as follows:

1. We present X-Ray, a novel 3D representation that encode the whole visible and hidden surfaces in to video format through ray casting algorithm for maintaining generalization, accuracy and efficiency.
2. We propose the generative model of our X-Ray via video diffusion model and video upsampling model, enabling low-to-high generation of 3D objects from single images.



Figure 2: Samples of our X-Ray 3D sequential representation. Given a viewpoint, we capture the 3D attributes multi-layer surface frames, including hit H , depth D , normal N , and color C , in a video format. Noted that the number of frames in an X-Ray varies depending on the complexity of the 3D objects. The dotted yellow lines indicate the ray or sequence direction.

3. We showcase the state-of-the-art performance of our X-Ray in 3D generation quality, setting a new benchmark for Image-to-3D modeling.

2 Related Work

2.1 Representation for 3D Models

Handling 3D data is much more complex and resource-intensive than dealing with 1D (*e.g.* text and voice) and 2D (*e.g.* image) data. This complexity makes it imperative to find effective ways to organize, process, and infer 3D information. Traditional methods for representing 3D data include meshes, which are good for creating detailed visuals but hard to be generalized; point clouds, which are simple and useful for capturing real-world scenes but lack consistent and dense structure in 3D creation; Although, 3D Gaussian Splatting [21] smooths point cloud data into continuous surface but requires additional an initial point cloud as shape, making them less flexible for 3D synthesis; voxels which are excellent for detailed volumetric data but require much computing resources. Multi-Plane Images [35, 51] try to extent the depth concept to multi-layer with a fixed distance, but they can only describe the visible surface toward camera.

Recent advancements in 3D representation have primarily focused on point-level details and implicit functions, such as Occupancy [32], Signed Distance Fields (SDF) [55, 11], Triplanes [10, 16], and Neural representations [36, 21]. These methods have significantly enhanced modeling and rendering capabilities. Occupancy models map the location of any 3D point to its probability of being inside or outside an object, offering a probabilistic approach to shape definition. SDFs [55, 19] refine this concept by quantifying the nearest signed surface distance from any given point, improving the precision of surface representations. Triplanes [3] employ intersecting 2D planes to provide a more efficient route to 3D representation, albeit with some detail loss. NeRFs and Gaussian Splatting [21, 36] produce remarkably realistic renderings from a limited number of viewpoints but require extensive computational effort. Despite these advancements, implicit function-based models often face challenges in extracting full and high-resolution 3D features, hindering high-quality generalization. Empirically, representations that focus on surface is more efficient, and representations with grid representation are more easily to be generalized [37]. Consequently, capturing all surface attributes and organizing in a dense but lightweight data structure renders our X-Ray an accurate, efficient and generalized representation. Noted that our X-Ray is similar to Depth peeling method [9], which is designed for rendering transparent surfaces, while X-Ray transform any 3D object in video format. Besides, our main contribution is using video diffusion as generator to generate objects.

2.2 Generative Models for 3D Generation

Recent 3D generative models can be primarily categorized into two types: diffusion-based [30, 37, 38] and rendering-based [14, 13, 49, 41, 34, 54, 33, 47]. Diffusion-based models fall under direct 3D supervision, while rendering-based models belong to indirect 2D image supervision. Diffusion-based generative models have emerged as powerful tools for 3D generation, leveraging stochastic diffusion processes to gradually transition from noise to structured 3D objects. These models, such as DPM [30], DiT-3D [37], and Point-E [38], have demonstrated remarkable ability in generating high-quality 3D point clouds. They operate by iteratively refining a random noise distribution into a coherent structure that resembles the target 3D shape, capturing complex geometries and surface details with high fidelity. The strength of these models lies in their capacity to model the distribution of 3D points in a continuous space, allowing for the generation of 3D objects with nuanced variations and detailed textures. However, both point-based networks [28, 37, 17] and voxel-based networks [37] is limited by generating high-resolution objects. Another group of methods [41, 34, 54, 33, 47] adopt Score Distillation Sampling (SDS) as prior to train a NeRF [36] or 3D Gaussian Splatting [21] for 3D generation. However, it is not efficient for optimizing a number of different views over a short period.

On the other hand, rendering-based generative models focus on the visible aspect of 3D generation, transforming abstract 3D representations into detailed and photorealistic images or videos. Models such as LRM [14], Open-LRM [13], LGM [48], DMV3D [57], and TripoSR [49] employ advanced rendering techniques to achieve this. However, rendering-based models are optimized only for the visible surfaces of objects, making it difficult to synthesize the invisible or internal surfaces.

In response to these challenges, our approach utilizes a video diffusion model as the foundation for developing 3D X-Ray. This strategy benefits from the strengths of existing video diffusion models while innovatively addressing the limitations of rendering-based techniques, offering a more comprehensive solution for 3D generation that is sensitive to both visible and hidden parts of objects.

3 Our X-Ray Representation

In this section, we detail our X-Ray representation, which includes both encoding and decoding process to facilitate the conversion between 3D mesh formats and our X-Ray representation. Specifically, the encoding process converts a 3D mesh into our proposed X-Ray format, and the decoding process converts our X-Ray back into a 3D mesh.

3.1 Encoding

Given a 3D object under an arbitrary camera view, we apply the ray casting algorithm to encode a 3D object mesh into the proposed X-Ray representation. The ray casting algorithm plays a crucial role in both computer graphics and computational geometry, where it is used for scene rendering, visibility determination, and addressing geometric queries. Specifically, a ray is emitted from camera center into the environment and all the interactions of this ray with target 3D objects are captured sequentially. For each ray $r \in \mathcal{R}$ within the field of view that intersects with L sequential faces in the mesh, we record their 3D attributes which include depth (distance to camera center) $\mathbf{D}_r = (\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_L) \in \mathbb{R}^{L \times 1}$, normal $\mathbf{N}_r = (\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_L) \in \mathbb{R}^{L \times 3}$, and color $\mathbf{C}_r = (\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_L) \in \mathbb{R}^{L \times 3}$. To indicate surface presence, we denote Hit $\mathbf{H}_r = (\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_L) \in \mathbb{R}^{L \times 1}$ to indicate whether there is a surface. Since L is usually very small (Sec. 5.3), we denote the efficient and thin grid voxels $\mathbf{X} \in \mathbb{R}^{H \times W \times L \times 8}$ as the object representation, where the ray \mathbf{X}_{ij} with image coordinate $[i, j]$ can be represented as:

$$\mathbf{X}_{ij} = \mathbf{X}_r = (\mathbf{H}_r, \mathbf{D}_r, \mathbf{N}_r, \mathbf{C}_r) \in \mathbb{R}^{L \times 8}. \quad (1)$$

Note that $\mathbf{X}[i, j, k] = 0$ when there is no surface for the k^{th} layer at the image ray coordinate $[i, j]$. Examples of X-Ray encoding are shown in Fig. 2. Through the encoding process, we can transform any mesh into a sequential representation with varying lengths, same as a video with different numbers of frames. Finally, we transpose the voxels as **X-Ray** $\mathbf{X} \in \mathbb{R}^{L \times 8 \times H \times W}$, resembling a video format with L frames, each with a resolution of $H \times W$ and 8 feature channels.

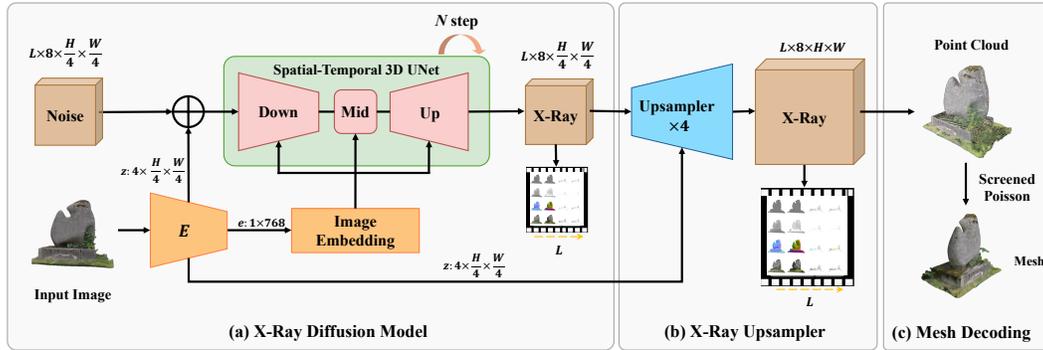


Figure 3: Overview of our proposed generative pipeline for the X-Ray 3D representation. There are three main components: (a) The X-Ray diffusion model, which generates a low-resolution X-Ray from an image input. (b) The upsampler, which enlarges the low-resolution X-Ray into $4\times$ high resolution. (c) The mesh decoding model, which decodes the high-resolution X-Ray into a point cloud with color and normal, and then converts it into the final generated mesh.

3.2 Decoding

The decoding process converts the X-Ray representation back into a 3D mesh. To achieve this, we first convert the video format of X-Ray into a point cloud and subsequently apply the Screened Poisson algorithm [20] to transform the point cloud into a 3D mesh.

X-Ray \rightarrow Point Cloud. Given an X-Ray, we first compute the 3D object's point cloud $\mathbf{P}_r = \{\mathbf{P}_x, \mathbf{P}_n, \mathbf{P}_c\}$ for Ray r , including location \mathbf{P}_x , color \mathbf{P}_c , and normal \mathbf{P}_n defined by the equation:

$$\mathbf{P}_x = \mathbf{r}_o + \mathbf{D}_r \cdot \mathbf{r}_d, \quad \mathbf{P}_n = \mathbf{N}_r, \quad \mathbf{P}_c = \mathbf{C}_r, \quad \text{when } \mathbf{H}_r = 1. \quad (2)$$

\mathbf{r}_o and \mathbf{r}_d denote the origin and direction of the camera ray, respectively. Furthermore, \mathbf{D}_r , \mathbf{N}_r , \mathbf{C}_r and \mathbf{H}_r are the depth, surface normal, color and hit attributes of the ray defined in Eq. 1. Upon processing all camera rays, we obtain a comprehensive point cloud $\mathbf{P} = \{\mathbf{P}_r\}_{r \in \mathcal{R}}$ representation that includes location, normal, and color attributes of the 3D object.

Point Cloud \rightarrow Mesh. The Screened Poisson algorithm [20] for converting point clouds with color and normal into 3D colored meshes is a classic method that leverages the mathematical principles of the Poisson equation. The core idea involves solving a variation of the Poisson equation to interpolate a smooth surface that fits the input point cloud. The Poisson equation is a partial differential equation of the form: $\nabla^2 \phi = f$, where ∇^2 denotes the Laplace operator (which represents the divergence of the gradient of a function), ϕ is the potential field to be solved, and f is a scalar function representing the divergence of the vector field derived from the input point cloud. In the context of point cloud to 3D mesh conversion, the algorithm first employs the given normal to define a vector field that suggests the orientation of the surface at each point. The divergence of this vector field serves as the function f the Poisson equation.

Encoding-Decoding Intrinsic Error. The encoding-decoding process will introduce an intrinsic and minor reconstruction error that varies with the number of layer L and the frame resolution (H, W) . To explore this, we conduct an experiment in Sec. 5.3 aimed at analyzing these variables to identify their optimal values. Our goal is to achieve a balance where all pertinent information is preserved while maintaining a lightweight model.

4 X-Ray for 3D Generation

Our primary objective of introducing a new 3D representation model is to facilitate the generation of 3D objects from single images. The challenge lies in accurately predicting the characteristics that are not immediately visible on the first surface when only a single image is available. To overcome this challenge, we utilize diffusion and upsampling models for X-Ray synthesis. Given that our proposed X-Ray is in video format, we leverage advanced Video Diffusion models as our backbone. To exploit this structure for high-resolution X-Ray synthesis, we incorporate principles from advanced video

diffusion models as our foundational framework. Notable models in this domain include Stable Video Diffusion (SVD) [2], VideoFusion [31], and the state-of-the-art Sora. To efficiently train the diffusion model, we begin by training a low-resolution X-Ray diffusion model that generates X-Ray from a single image. Subsequently, we employ an upsampler to enhance these synthesized X-Rays to high resolution. This two-step approach ensures a more manageable and efficient training process, gradually improving the quality of the output.

Framework Overview. Fig. 3 presents an overview of our generative model using X-Ray representation (cf. Sec. 3.1) for 3D generation. Our X-Ray diffusion model (cf. Sec. 4.1) operates at the core of the framework, transforming random Gaussian noise into a low-resolution X-Ray representation conditioned by an input image. These low-resolution X-Rays are then enhanced to high resolution through the application of a 3D Spatial-Temporal Upsampler (cf. Sec. 4.2). Finally, the high-resolution X-Rays are decoded into 3D meshes using a combination of point cloud transformation and the Screened Poisson algorithm (cf. Sec. 3.2).

4.1 X-Ray Diffusion Model

Diffusion models [43] are generative models that transform a random noise distribution into a data distribution through a reverse process, counteracting a forward process that incrementally adds Gaussian noise to the data. The forward process is a Markov chain described by $x_t = \sqrt{\alpha_t}x_{t-1} + \sqrt{1 - \alpha_t}\epsilon$, where x_t represents the data at step t , α_t controls the noise level, and $\epsilon \sim \mathcal{N}(0, I)$ is sampled noise. The reverse process, aimed at reconstructing the original data from noise, is modeled by a neural network predicting the noise added at each step or directly denoising the data, following $x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t^2}} \epsilon_\theta(x_t, t) \right)$, with $\epsilon_\theta(x_t, t)$ being the predicted noise. Training involves optimizing the network to minimize the difference between the original and reconstructed data, effectively learning to invert the noise addition process given by:

$$\mathcal{L}_{dm} = \mathbb{E}_{x, \epsilon \sim \mathcal{N}(0,1), t} [\|\epsilon - \theta(x_t, t)\|^2], \quad (3)$$

where t is uniformly sampled from the set $\{1, \dots, T\}$.

Diffusion Model for X-Ray. A prevalent technique in diffusion models is the utilization of latent spaces with a VQ-VAE [7] to perform the initial data transformation to compress the data. This method poses a significant challenge for our X-Ray representation as it requires the development of a VQ-VAE model from scratch due to the absence of a suitable off-the-shelf latent model for X-Ray, which would consequently increase our training burden. Another promising approach for efficiently training high-resolution generators is the cascaded synthesis pipeline. This method, exemplified by works such as Imagen [44], DeepFloyd IF [1], and Stable Cascaded [40], involves progressively training the diffusion model or upsampling network from lower to higher resolutions. Given our limited computing resources, we opted to implement this cascaded upsampling strategy. This technique facilitates a more gradual and controlled enhancement of X-Ray quality, providing a more flexible and efficient alternative to traditional latent space diffusion models.

Specifically, we use the Spatial-Temporal 3D U-Net network from Stable Video Diffusion [2] for our diffusion model to generate low-resolution X-Rays, with modifications to the input and output channels. As shown in Fig. 3, the input image is sent to the encoder \mathbf{E} , producing an image latent z via image VAE [7] and an embedding e via a ViT [5]. z is concatenated with the X-Ray latent as input, and e interacts with the 3D U-Net through cross attention [2] to finally output the denoised latent. This model employs spatial-temporal attention mechanisms to alternately extract features from 2D frame spaces and 1D surface layer sequences, enhancing its ability to process and interpret the different layers of the X-Ray. This approach allows for nuanced handling of the temporal information inherent in sequential X-Ray data, crucial for achieving high-quality diffusion results.

4.2 X-Ray Upsampler

The X-Ray upsampler focuses on enhancing previously generated X-Rays to a higher resolution. We considered two potential methods: point cloud up-sampling and video up-sampling. Encoding low-resolution X-Rays into a point cloud with color and normal information is straightforward (see Sec 3.1). However, point cloud up-sampling often increases only the number of points without effectively enhancing attributes like texture and color due to its unstructured nature. To improve

efficiency and consistency, we adopted a video up-sampling approach using a spatial-temporal VAE decoder from Stable Video Diffusion (SVD) [2]. We concatenate the previous image latent z with the low-resolution X-Ray as input and output the high-resolution X-Ray. The model upsamples previously synthesized low-resolution X-Ray frames fourfold while preserving the original layer number L . It applies attention at both the 2D surface frame and 1D surface layer levels, enhancing frame resolution and overall quality. This makes the X-Ray diffusion model, followed by the upsampler, a more integrated and effective solution.

Loss. The loss function for the Upsampler differs notably from that of the diffusion model. While the diffusion model loss typically addresses volumetric or textural aspects, the Upsampler loss concentrates specifically on the surface area accuracy, reflecting the critical importance of maintaining high fidelity in the enhanced images. The loss function we use for the Upsampler is given by:

$$\mathcal{L}_{up} = \|\mathbf{X}_{gt}[\mathbf{H}_{gt}] - \mathbf{X}_{up}[\mathbf{H}_{gt}]\|^2 + \|\mathbf{H}_{gt} - \mathbf{H}_{up}\|^2, \quad (4)$$

where $\mathbf{X}_{gt}[\mathbf{H}_{gt}]$ represents the ground-truth high-resolution X-Ray at hit surface and $\mathbf{X}_{up}[\mathbf{H}_{gt}]$ denotes the Upsampler’s output at hit surface, and \mathbf{H}_{gt} , \mathbf{H}_{up} denotes the ground-truth and upsampled Hit, respectively. The loss is computed as the squared Euclidean distance between these two matrices, quantifying the pixel-wise discrepancy in surface details. This metric effectively ensures that the upsampling process preserves essential surface features, thereby optimizing the quality and utility of the resulting high-resolution X-Ray.

5 Experiments

5.1 Dataset and Implementation

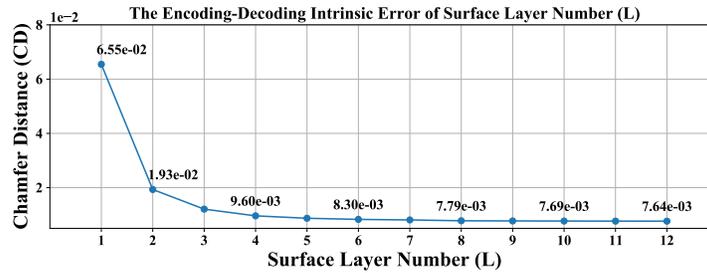
Datasets. We train our X-Ray pipeline using a subset of the Objaverse dataset [4], which have been removed entries with missing textures and inadequate prompts as outlined in [48]. This subset consists of more than 60,000 3D objects. For each object, we select 8 random camera views, covering azimuth angles from -180 to 180 degrees and elevation angles from 0 to 45 degrees with camera distance to object center fixed at 1.2. The images are then rendered using Blender Software, and the corresponding X-Rays are generated through the ray casting algorithm provided by the trimesh library [50]. Through these processes, we create a dataset of approximately 480,000 paired images and X-Rays to train the generative model. For the evaluation datasets, we adopt two commonly adopted datasets: Google Scanned Objects [6] and OmniObject3D [56], to assess generative performance via single-view reconstruction tasks.

Metrics. Recent 3D generative models [26, 18, 27, 14, 49] lack unified reconstruction evaluation metrics due to the challenge of determining an object’s size and orientation from a single image [46]. To ensure fair comparisons with state-of-the-art methods, we align all methods before evaluation. The predicted and ground-truth 3D objects will be normalized to a range of -0.5 to 0.5 along all three axes and face forward the same $-z$ axis. We then align them using the Iterative Closest Points (ICP) algorithm and calculate Chamfer Distance (CD) in $L1$ norm and F-Score (FS) at threshold 0.1 (FS@0.1) for reconstruction.

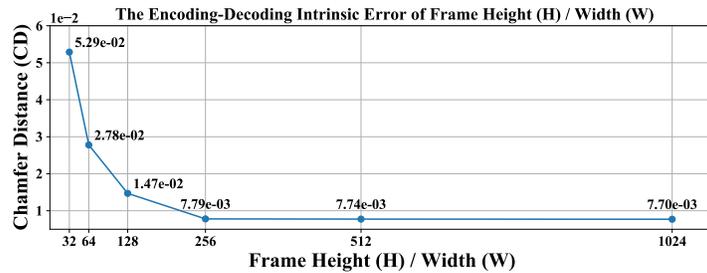
Implementation Details. Our X-Ray diffusion model is based on the Spatial-Temporal 3D UNet architecture used in Stable Video Diffusion (SVD) [2], modified to synthesize 8 channels: 1 hit channel, 1 depth channel, 3 color channels, and 3 normal channels, compared to the original 4 channels. During training, we maintain a learning rate of 0.0001 using the AdamW optimizer. Since different X-Rays have varying numbers of layers, we pad or truncate them to a uniform 8 layers for efficient batching and training. Each layer’s frame has dimensions of 64×64 . For the upsampler, each layer’s output remains at 8 channels, but the resolution of each frame is increased to 256×256 to enhance detail and clarity in the upscaled X-Ray. The entire training pipeline is conducted on 8 NVIDIA A100 GPU servers for two weeks. During inference, the 3D generation process takes approximately 7 seconds: about 1 second for the diffusion model, 1 second for the upsampler, and 5 seconds for mesh decoding. As for GPU usage during inference, the GPU memory required is 4.8 GB for X-Ray diffusion model and 2.5 GB for X-Ray Upsampler.

Table 1: Comparison with other 3D representations in Efficiency.

Metric	3D Grid	Multi-View Depths (8 views)	MPI (8 planes)	Point Cloud (200,000 points)	X-Ray (8 layers)
Memory (↓)	67.09 MB	1.57 MB	1.57 MB	0.90 MB	0.62 MB
Encoding Method	Voxlization	Rendering	Slicing & Rendering	Sampling	Ray Casting
Encoding Time (↓)	0.105 s	0.045 s	0.049 s	0.013 s	0.016 s
Decoding Method	Poisson	Fusion & Poisson	Poisson	Poisson	Poisson
Decoding Time (↓)	~5 s	~10 s	~5 s	~5 s	~5 s
CD(↓)	7.7e-3	1.1e-2	8.9e-3	7.2e-3	7.8e-03



(a) The intrinsic error of layer number (L), when $H = W = 256$



(b) The intrinsic error of frame height (H) or width (W) when $L = 8$

Figure 4: The encoding-decoding intrinsic error of different frame resolutions and number of layers.

5.2 Efficient Comparison with Different 3D Representation

We compared the efficiency of different representations using 500 3D meshes from random selected models. The results showed that both point cloud and X-Ray were highly efficient, with lower memory, faster encoding & decoding times. However, the X-Ray had the advantage of being reorganizable as a video format for diffusion models, leading to better performance.

5.3 Analysis of Encoding-Decoding Intrinsic Error

Due to the finite number of layer L and resolution H, W of X-Ray, a slight intrinsic error is inevitable during the encoding of 3D meshes into X-Ray format and the subsequent decoding back into 3D meshes. To quantitatively assess this error, we conducted an experiment to evaluate intrinsic error via Chamfer Distance (CD) ($L1$ norm) between the original (ground-truth) mesh and the encoding-decoding mesh across various resolutions and layers. In the experiments, we set the frame resolution through a series of predefined values: 32, 64, 128, 256, 512, and 1,024 and vary the number of layers from 1 to 12. Fig. 4a indicates that the intrinsic error decreases as the number of layers in our X-Ray representation increases, becoming convergent after 8 layers. Similarity as illustrated in Fig. 4b, the intrinsic error decreases with increasing resolution and stabilizes after 256. Therefore, for a balance between accuracy and efficiency, we use a resolution $8 \times 256 \times 256$ with $H = W = 256$ and $L = 8$ in our experiments. Compared with the dense volume achieving a $256 \times 256 \times 256$ resolution for voxel-based methods, our X-Ray representation is significantly efficient for focusing only on surfaces and reducing the data volume by 96.88%.

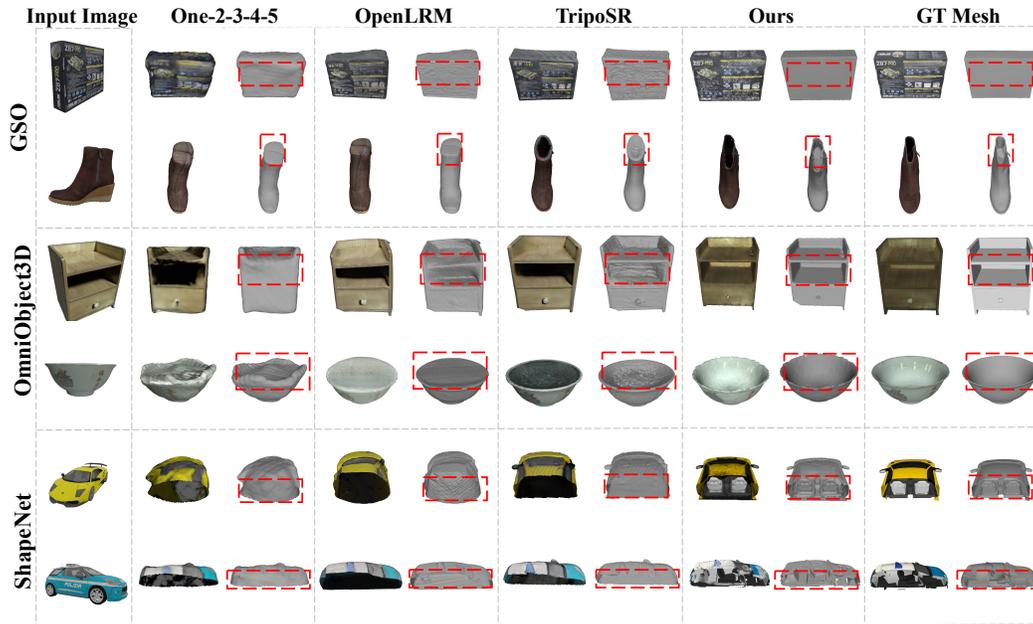


Figure 5: Quantitative Comparison in Image-to-3D Generation.

5.4 Quantitative Comparison

For image-to-3D mesh generation on the GSO [6] and OmniObject3D [56] datasets, we build a new benchmark that randomly selects 500 image-mesh pairs from each dataset. We re-ran the baselines using their released source code and used normalized Chamfer distance and F-score for a fair comparison. The results, summarized in Tab. 2, show that our X-Ray method achieves significant improvements over all previous rendering-based state-of-the-art methods [25, 18, 8, 13, 49] on both datasets, with a relative **33%** improvement ($0.084 \rightarrow 0.056$) in Chamfer distance and also a significant improvement in FS@0.1 ($0.878 \rightarrow 0.973$, where the maximum is 1) on GSO dataset and similar performance on OmniObject3D dataset. This demonstrates the superiority of our approach over rendering-based methods.

Table 2: Quantitative reconstruction comparison on the GSO [6] and OmniObject3D [56] datasets

Datasets	Metrics	One-2-3-45 [25]	ZeroShape [18]	TGS [8]	OpenLRM [13]	TripoSR [49]	X-Ray (Ours)
GSO [6]	CD ↓	0.175	0.136	0.096	0.143	0.084	0.056
	FS@0.1 ↑	0.465	0.627	0.803	0.621	0.878	0.973
OmniObject3D [56]	CD ↓	0.187	0.138	0.091	0.148	0.080	0.054
	FS@0.1 ↑	0.490	0.619	0.822	0.664	0.892	0.972

5.5 Qualitative Comparison

A qualitative comparison effectively demonstrates the advantages of our proposed X-Ray generation method. Using the three datasets mentioned, we selected single images as input and generated 3D meshes without and with textures, as shown in Fig. 5. Our proposed method has three key advantages: 1. It can decompose shape and appearance to accurately reconstruct flat surfaces (Rows 1 and 3); 2. It can detect the sealing of containers (Rows 2 and 4); 3. It can generate the internal structure of objects within their outer surfaces (Rows 5 and 6). These obvious advantages highlight the our effectiveness.

5.6 Failure Cases.

The failure cases highlight the limitations of the current generative model based on X-Ray representation when the number of frame layers is very large. As illustrated in Fig. 6, given an input image containing a complex object, such as an exquisite hamburger, the number of frame layers L of the encoded or generated X-Ray tends to exceed the maximum length of 8. Consequently, any surface

behind layer 8 will be omitted, resulting in missing parts of the reconstructed mesh. The solution is to increase the value of L so that the X-Ray can represent more surfaces. However, this would also increase computing resource requirements. We will reconsider the sparsity of deeper surface frames and propose a more efficient generative model to overcome these failure cases.

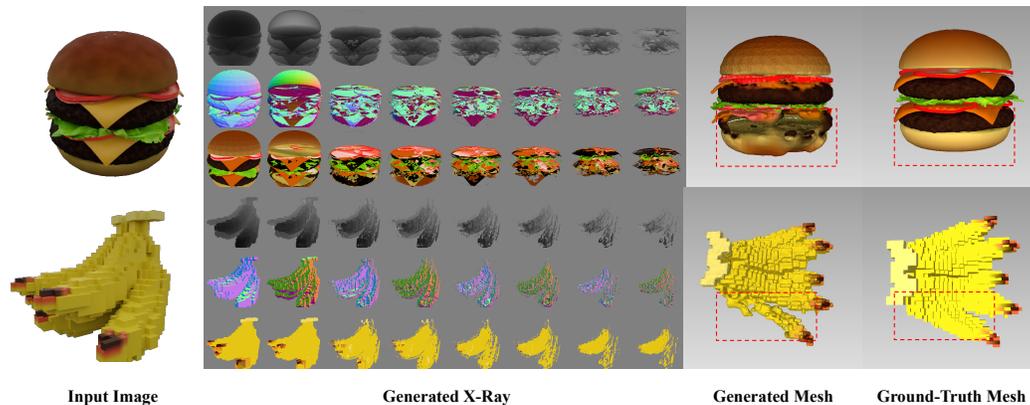


Figure 6: Failure cases. The generated meshes will miss behind parts because of the limited number of frame layers.

6 Conclusion

In this work, we introduced a novel X-Ray representation for 3D objects that encompasses both visible and hidden surfaces within the camera’s field of view, unlike recent rendering-based methods that typically focus only on the visible surface. We demonstrated the effectiveness of the X-Ray approach in single-view 3D generation tasks. Our generative model shows that the underlying generator for X-Ray shares foundational similarities with existing video diffusion models, allowing us to leverage their advantages. Experimental results highlight the outstanding performance of our method.

Limitations. Our generator uses the Stable Video Diffusion (SVD) pipeline to produce high-quality X-Ray. However, the X-Rays generated consist of an uncertain number of sequential layers, with the posterior layers tending to be more and more sparse, which can be redundant. Additionally, the generated mesh lacks enough smoothness and has missing part when X-Ray is truncated. We plan to explore advanced network architectures, such as Large Language Model, that can better handle the complexities of X-Ray data, including layer sparsity and sequential format. Additionally, we aim to investigate further applications of the X-Ray representation to broaden its utility and impact in 3D modeling and beyond.

Acknowledgement

This research / project is supported by the National Research Foundation (NRF) Singapore, under its NRF-Investigatorship Programme (Award ID. NRF-NRFI09-0008).

References

- [1] DeepFloyd Lab at StabilityAI. DeepFloyd IF: a novel state-of-the-art open-source text-to-image model with a high degree of photorealism and language understanding. <https://www.deepfloyd.ai/deepfloyd-if>, 2023. Retrieved on 2023-11-08.
- [2] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. Stable video diffusion: Scaling latent video diffusion models to large datasets. *CoRR*, 2023.
- [3] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022.

- [4] 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. *arXiv preprint arXiv:2212.08051*, 2022.
- [5] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021.
- [6] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas Barlow McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In *2022 International Conference on Robotics and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022*, 2022.
- [7] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis, 2020.
- [8] Zi-Xin Zou et. al. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. *arXiv preprint arXiv:2312.09147*, 2023.
- [9] Cass W. Everitt. Interactive order-independent transparency. In *NVIDIA*, 2001.
- [10] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3d: A generative model of high quality 3d textured shapes learned from images. In *Advances In Neural Information Processing Systems*, 2022.
- [11] Wenheng Ge, Tao Hu, Haoyu Zhao, Shu Liu, and Ying-Cong Chen. Ref-neus: Ambiguity-reduced neural implicit surface learning for multi-view reconstruction with reflection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4251–4260, 2023.
- [12] JunYoung Gwak, Christopher B Choy, and Silvio Savarese. Generative sparse detection networks for 3d single-shot object detection. In *European conference on computer vision*, 2020.
- [13] Zexin He and Tengfei Wang. Openlrm: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM>, 2023.
- [14] 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. *CoRR*, abs/2311.04400, 2023.
- [15] Tao Hu, Shu Liu, Yilun Chen, Tiancheng Shen, and Jiaya Jia. Efficientnerf efficient neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12902–12911, 2022.
- [16] Tao Hu, Xiaogang Xu, Ruihang Chu, and Jiaya Jia. Trivol: Point cloud rendering via triple volumes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20732–20741, 2023.
- [17] Tao Hu, Xiaogang Xu, Shu Liu, and Jiaya Jia. Point2pix: Photo-realistic point cloud rendering via neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8349–8358, 2023.
- [18] Zixuan Huang, Stefan Stojanov, Anh Thai, Varun Jampani, and James M Rehg. Zeroshape: Regression-based zero-shot shape reconstruction. *arXiv preprint arXiv:2312.14198*, 2023.
- [19] Ka-Hei Hui, Aditya Sanghi, Arianna Rampini, Kamal Rahimi Malekshan, Zhengzhe Liu, Hooman Shayani, and Chi-Wing Fu. Make-a-shape: a ten-million-scale 3d shape model. *CoRR*, 2024.
- [20] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Trans. Graph.*, 2013.
- [21] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [23] Xiaoyu Li, Qi Zhang, Di Kang, Weihao Cheng, Yiming Gao, Jingbo Zhang, Zhihao Liang, Jing Liao, Yan-Pei Cao, and Ying Shan. Advances in 3d generation: A survey. *CoRR*, abs/2401.17807, 2024.

- [24] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020.
- [25] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36, 2024.
- [26] 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, 2023.
- [27] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. Syncdreamer: Generating multiview-consistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023.
- [28] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel CNN for efficient 3d deep learning. In *NeurIPS*, 2019.
- [29] Zhen Liu, Yao Feng, Yuliang Xiu, Weiyang Liu, Liam Paull, Michael J. Black, and Bernhard Schölkopf. Ghost on the shell: An expressive representation of general 3d shapes. In *ICLR*, 2024.
- [30] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *CVPR*, 2021.
- [31] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [32] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019.
- [33] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. *arXiv preprint arXiv:2211.07600*, 2022.
- [34] Oscar Michel, Roi Bar-On, Richard Liu, Sagie Benaim, and Rana Hanocka. Text2mesh: Text-driven neural stylization for meshes. *arXiv preprint arXiv:2112.03221*, 2021.
- [35] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019.
- [36] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [37] Shentong Mo, Enze Xie, Yue Wu, Junsong Chen, Matthias Nießner, and Zhenguo Li. Fast training of diffusion transformer with extreme masking for 3d point clouds generation. *arXiv preprint arXiv:2312.07231*, 2023.
- [38] Alex Nichol, Heewoo Jun, Prafulla Dhariwal, Pamela Mishkin, and Mark Chen. Point-e: A system for generating 3d point clouds from complex prompts. *CoRR*, abs/2212.08751, 2022.
- [39] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023.
- [40] Pablo Pernias, Dominic Rampas, Mats L. Richter, Christopher J. Pal, and Marc Aubreville. Wuerstchen: An efficient architecture for large-scale text-to-image diffusion models, 2023.
- [41] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022.
- [42] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017.
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 2022.
- [44] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L. Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022.

- [45] Sara Fridovich-Keil and Alex Yu, Matthew Tancik, Qinlong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, 2022.
- [46] Kyle Sargent, Zizhang Li, Tanmay Shah, Charles Herrmann, Hong-Xing Yu, Yunzhi Zhang, Eric Ryan Chan, Dmitry Lagun, Li Fei-Fei, Deqing Sun, and Jiajun Wu. ZeroNVS: Zero-shot 360-degree view synthesis from a single real image. *arXiv preprint arXiv:2310.17994*, 2023.
- [47] Jiayang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023.
- [48] 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. *arXiv preprint arXiv:2402.05054*, 2024.
- [49] Dmitry Tochilkin, David Pankratz, Zexiang Liu, Zixuan Huang, , Adam Letts, Yangguang Li, Ding Liang, Christian Laforte, Varun Jampani, and Yan-Pei Cao. Triposr: Fast 3d object reconstruction from a single image. *arXiv preprint arXiv:2403.02151*, 2024.
- [50] Trimesh. Trimesh [computer software]. <https://github.com/mikedh/trimesh>, 2019.
- [51] Richard Tucker and Noah Snavely. Single-view view synthesis with multiplane images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [52] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [53] Vikram Voleti, Chun-Han Yao, Mark Boss, Adam Letts, David Pankratz, Dmitry Tochilkin, Christian Laforte, Robin Rombach, and Varun Jampani. SV3D: novel multi-view synthesis and 3d generation from a single image using latent video diffusion. *CoRR*, abs/2403.12008, 2024.
- [54] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A. Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. *arXiv preprint arXiv:2212.00774*, 2022.
- [55] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021.
- [56] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Liang Pan Jiawei Ren, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, Dahua Lin, and Ziwei Liu. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. In *CVPR*, 2023.
- [57] Yinghao Xu, Hao Tan, Fujun Luan, Sai Bi, Peng Wang, Jiahao Li, Zifan Shi, Kalyan Sunkavalli, Gordon Wetzstein, Zexiang Xu, and Kai Zhang. DMV3D: denoising multi-view diffusion using 3d large reconstruction model. *CoRR*, abs/2311.09217, 2023.
- [58] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *ICCV*, 2021.

A Appendix

The Appendix section contains detailed information on the network structure (including the X-Ray diffusion model and upsampler), ablation studies, additional experimental quantitative and qualitative results, failure cases, and key source code.

A.1 Network Details

X-Ray Diffusion Model As outlined in our study, we utilized the 3D Spatial-Temporal UNet as the foundational architecture for our diffusion model. For image-to-3D generation, the input includes 4-channel image latent and 8-channel noise X-Ray and the output has 8 channel to denoise X-Ray. Initially, we loaded all network parameters and fine-tuned the model on our training dataset. Despite these efforts, the final performance fell short of expectations. We hypothesized that this underperformance stemmed from the limited size of our training dataset, which was insufficient for such a parameter-rich model, as well as a substantial domain gap between the original video data and our X-ray data.

To address these issues, we reduced the network size to 10% of its original configuration and trained the model from scratch using a significantly larger batch size. This approach proved highly effective, as the final performance exceeded the current state-of-the-art methods by a considerable margin. Consequently, our findings suggest that a smaller, more focused network trained from scratch can be more effective than a larger pre-trained model when faced with limited and domain-specific datasets. This approach not only enhances performance but also highlights the importance of customizing the model size and training strategy to the specific characteristics of the data. Our results indicate that careful consideration of the network architecture and training regimen is crucial for optimizing performance, particularly in specialized applications such as X-ray generation. This study provides valuable insights into model adaptation and training strategies that can be applied to other domains facing similar challenges.

The following JSON file contains the configuration details for our X-Ray Spatial-Temporal UNet.

```
1 {
2   "class_name": "UNetSpatioTemporalConditionModel",
3   "addition_time_embed_dim": 256,
4   "block_out_channels": [
5     64,
6     128,
7     256,
8     256
9   ],
10  "cross_attention_dim": 1024,
11  "down_block_types": [
12    "CrossAttnDownBlockSpatioTemporal",
13    "CrossAttnDownBlockSpatioTemporal",
14    "CrossAttnDownBlockSpatioTemporal",
15    "DownBlockSpatioTemporal"
16  ],
17  "in_channels": 12,
18  "latent_channels": 8,
19  "layers_per_block": 1,
20  "num_attention_heads": [
21    4,
22    8,
23    16,
24    16
25  ],
26  "num_frames": 8,
27  "out_channels": 8,
28  "projection_class_embeddings_input_dim": 768,
29  "sample_size": 64,
30  "transformer_layers_per_block": 1,
31  "up_block_types": [
32    "UpBlockSpatioTemporal",
```

```

33     "CrossAttnUpBlockSpatioTemporal" ,
34     "CrossAttnUpBlockSpatioTemporal" ,
35     "CrossAttnUpBlockSpatioTemporal"
36 ]
37 }

```

A.2 X-Ray Upsampler

The X-Ray diffusion model enables the generation of low-resolution 3D objects. To enhance resolution and improve performance, our X-Ray upsampler increases the frame resolution by a factor of 4. For image-to-3D generation, we concatenate the 4-channel image latent representation with the 8-channel low-resolution X-Ray, producing an 8-channel high-resolution X-Ray.

This process begins with the diffusion model creating a coarse, low-resolution 3D representation that captures the essential structure of the object. The upsampler then refines this representation, significantly improving the resolution and adding finer details that contribute to the realism of the 3D model. By combining the latent image features with the initial low-resolution X-Ray data, we ensure that the final high-resolution output retains the context and nuances of the original image while also incorporating the detailed structural information provided by the X-Ray data.

The following JSON file contains the configuration details for our X-Ray Upsampler.

```

1 {
2   "_class_name": "AutoencoderKLTemporalDecoder",
3   "block_out_channels": [
4     128,
5     256,
6     512
7   ],
8   "down_block_types": [
9     "DownEncoderBlock2D",
10    "DownEncoderBlock2D",
11    "DownEncoderBlock2D"
12  ],
13  "force_upcast": false,
14  "in_channels": 4,
15  "latent_channels": 12,
16  "layers_per_block": 2,
17  "out_channels": 8,
18  "sample_size": 768,
19  "scaling_factor": 1.0
20 }

```

A.3 Ablation Studies

The Effect of Diffusion Model

As described in Sec. A.1, we train the diffusion model with different configurations, including varying model sizes and how to conduct initialization. Specifically, we evaluate three models:

1. Finetuned Original UNet
2. Randomly Initialized Original UNet
3. Randomly Initialized UNet with 10% Parameters

The experimental evaluation results on the GSO [6] dataset are illustrated in Tab. 3. These results allow us to analyze the impact of different initialization and scaling strategies on the performance of the diffusion model. By comparing the outcomes, we can identify the trade-offs between model size, training time, and overall accuracy. We observed that finetuning the model did not introduce significant improvement because of the domain gap between video and X-Ray data. Also, it is not necessary to adopt a large diffusion model for Objaverse dataset [4] for significantly reducing the batch size. Thus, we adopt the randomly initialized UNet with 10% parameters as our diffusion

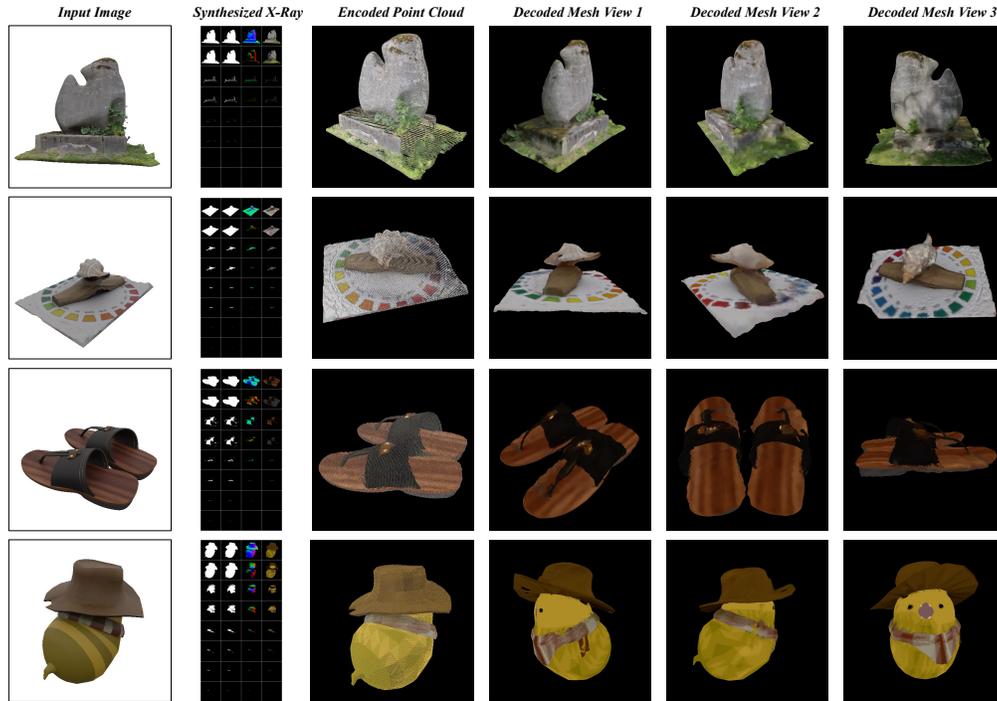


Figure 7: Visualization of Image-to-3D Generation from X-Ray.

model. These findings provide insights into the optimal configuration for diffusion models in 3D related tasks.

Table 3: Quantitative reconstruction comparison in different diffusion model configurations

Model Configurations	CD ↓	FS@0.1 ↑	Training Time (days) ↓	Inference Time (seconds)	Batch Size ↑	Model Size (GB) ↓
Finetuned Original UNet	0.095	0.812	~ 14	~ 18	2	6.1
Randomly Initialized Original UNet	0.099	0.806	~ 14	~ 18	2	6.1
Randomly Initialized UNet with 10% Parameters	0.056	0.973	7	~ 7	24	0.6

The Effect of Hit H

The original surface attributes only contain depth D , normal N , and color C . We add an additional Hit H attribute to indicate whether there is a surface. In this ablation study, we demonstrate the necessity of including the Hit H attribute. We conduct training experiments with and without the Hit H attribute and evaluate on the GSO [6] dataset. The results, shown in Tab. 4, indicate that including the Hit H attribute can improve performance. The reason might be that the X-Ray is sparse and requires an indicator to balance the generative process. The UNet model and Upsampler with the Hit H attribute achieves better CD and FS@0.1 scores, demonstrating its importance in accurate 3D generation.

Table 4: Quantitative evaluation of the effect of Hit H attribute on the GSO [6] dataset

Diffusion Model	CD ↓	FS@0.1 ↑	Upsampler	CD ↓	FS@0.1 ↑
wo. Hit H	0.074	0.901	wo. Hit H	0.060	0.956
w/ Hit H	0.068	0.934	w/ Hit H	0.056	0.973

A.4 More Visualization

We shown more Visualization results of single-view image generation in Fig. 7. Also, we extend the task to generate 3D object from text prompt. Text-to-3D mesh generation can also be achieved via using text-to-image, object segmentation, and image-to-3D processes. We utilize established diffusion models that are already proficient in image synthesis from textual descriptions instead of

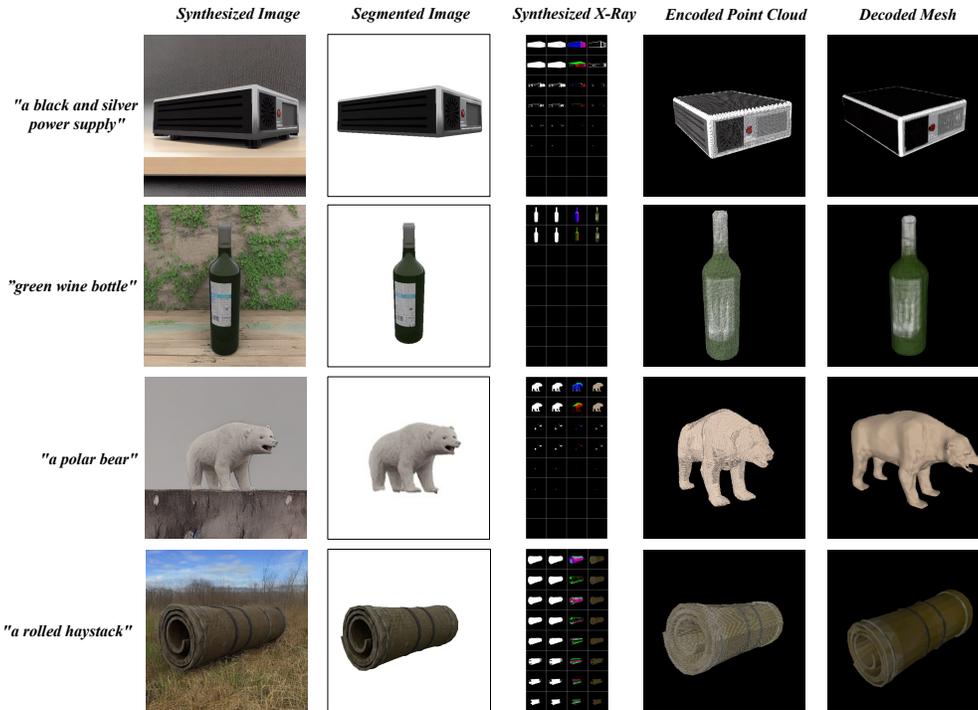


Figure 8: Visualization of Text-to-3D Generation from X-Ray.

developing a new text-conditioned diffusion model. One Model such as Stable Diffusion [43], Stable Cascaded [40], or DiT [39] are employed to generate images based on the input text. Following this, we apply an image segmentation tool, specifically the Segment Anything Model [22] to eliminate the background. This streamlined method avoids the complexities of training a new model from scratch instead of making use of sophisticated pre-trained models to handle the text-to-image translation, thereby simplifying the process of generating 3D meshes from textual inputs. The output results of Image-to-3D and Text-to-3D are illustrated in Fig. 8.

A.5 Key Source Code

Ray Casting: The source code of ray casting the mesh to get the point cloud.

```

1 """
2 The source code of ray casting the mesh to get the point cloud.
3 """
4
5 from trimesh.ray.ray_pyembree import RayMeshIntersector
6
7 def ray_cast_mesh(mesh, rays_origins, ray_directions):
8     intersector = RayMeshIntersector(mesh)
9     index_triangles, index_ray, point_cloud = intersector.
intersects_id(
10         ray_origins=rays_origins,
11         ray_directions=ray_directions,
12         multiple_hits=True,
13         return_locations=True)
14     return index_triangles, index_ray, point_cloud

```

X-Ray to Point Cloud: The source code of transferring X-Ray to Point Clouds with normals and colors.

```

1 """
2 The source code of transferring X-Ray to Point Clouds with normals and
3 colors.

```

```

3 """
4
5 import numpy as np
6 import open3d as o3d
7
8 def get_rays(directions, c2w):
9     # Rotate ray directions from camera coordinate to the world
10    coordinate
11    rays_d = directions @ c2w[:3, :3].T # (H, W, 3)
12    rays_d = rays_d / (np.linalg.norm(rays_d, axis=-1, keepdims=True)
13    + 1e-8)
14
15    # The origin of all rays is the camera origin in world coordinate
16    rays_o = np.broadcast_to(c2w[:3, 3], rays_d.shape) # (H, W, 3)
17
18    return rays_o, rays_d
19
20 def X_Ray_to_Point_Cloud(XRay):
21    """
22    Converts X-Ray data to a point cloud with normals and colors.
23    """
24    XDepths, XNormals, XColors, XHits = XRay[:, 0:1], XRay[:, 1:4],
25    XRay[:, 4:7], XRay[:, 7:8]
26
27    camera_angle_x = 0.8575560450553894
28    image_width = XDepths.shape[-1]
29    image_height = XDepths.shape[-2]
30    fx = 0.5 * image_width / np.tan(0.5 * camera_angle_x)
31
32    rays_screen_coords = np.mgrid[0:image_height, 0:image_width].
33    reshape(
34        2, image_height * image_width).T # [h, w, 2]
35
36    grid = rays_screen_coords.reshape(image_height, image_width, 2)
37
38    cx = image_width / 2.0
39    cy = image_height / 2.0
40
41    i, j = grid[..., 1], grid[..., 0]
42
43    directions = np.stack([(i - cx) / fx, -(j - cy) / fx, -np.
44    ones_like(i)], -1) # (H, W, 3)
45
46    c2w = np.eye(4).astype(np.float32)
47
48    rays_origins, ray_directions = get_rays(directions, c2w)
49    rays_origins = rays_origins[None].repeat(XDepths.shape[0], 0)
50    ray_directions = ray_directions[None].repeat(XDepths.shape[0], 0)
51
52    XDepths = XDepths.transpose(0, 2, 3, 1)
53    XNormals = XNormals.transpose(0, 2, 3, 1)
54    XColors = XColors.transpose(0, 2, 3, 1)
55
56    rays_origins = rays_origins[XHits]
57    ray_directions = ray_directions[XHits]
58    XDepths = XDepths[XHits]
59    normals = XNormals[XHits]
60    colors = XColors[XHits]
61    xyz = rays_origins + ray_directions * XDepths
62
63    # convert to open3d point cloud
64    xyz = xyz.reshape(-1, 3)
65    normals = normals.reshape(-1, 3)
66    colors = colors.reshape(-1, 3)

```

```

63 pcd = o3d.geometry.PointCloud()
64 pcd.points = o3d.utility.Vector3dVector(xyz)
65 pcd.normals = o3d.utility.Vector3dVector(normals)
66 pcd.colors = o3d.utility.Vector3dVector(colors)
67
68 return pcd

```

Point Cloud to Mesh: The source code of transferring the predicted point cloud to mesh using Poisson Surface Reconstruction.

```

1 """
2 The source code of transferring the predicted point cloud to mesh
  using
3 Poisson Surface Reconstruction.
4 """
5
6 import open3d as o3d
7
8 # Load point cloud
9 pcd = o3d.io.read_point_cloud("path_to_your_point_cloud.ply")
10
11 def poisson_surface_reconstruction(pcd):
12     """
13     Converts a point cloud to a mesh using Poisson Surface
14     Reconstruction.
15     """
16     # Ensure the point cloud has normals
17     if not pcd.has_normals():
18         pcd.estimate_normals()
19
20     # Perform Screened Poisson Surface Reconstruction
21     mesh, densities = o3d.geometry.TriangleMesh.
22     create_from_point_cloud_poisson(
23         pcd, depth=9, width=0, scale=1.1, linear_fit=False
24     )
25
26     # Optionally crop the mesh using the density values to remove low-
27     # density areas
28     # You can adjust the threshold based on your requirements
29     density_threshold = 0.01
30     vertices_to_remove = densities < density_threshold
31     mesh.remove_vertices_by_mask(vertices_to_remove)
32
33     # Assign colors to the mesh
34     mesh.vertex_colors = o3d.utility.Vector3dVector(pcd.colors)
35     return mesh

```

```

1 """
2 The source code of evaluating the predicted mesh using Chamfer
  Distance and F-Score.
3 """
4
5 import numpy as np
6 from scipy.spatial import cKDTree
7
8 def chamfer_distance_and_f_score(P, Q, threshold=0.1):
9     """
10     Calculates the Chamfer Distance and F-Score between two point
11     clouds.
12     """
13     kdtree_P = cKDTree(P)
14     kdtree_Q = cKDTree(Q)
15
16     dist_P_to_Q, _ = kdtree_P.query(Q)
17     dist_Q_to_P, _ = kdtree_Q.query(P)

```

```

17
18     chamfer_dist = np.mean(dist_P_to_Q) + np.mean(dist_Q_to_P)
19
20     precision = np.mean(dist_P_to_Q < threshold)
21     recall = np.mean(dist_Q_to_P < threshold)
22
23     if precision + recall > 0:
24         f_score = 2 * (precision * recall) / (precision + recall)
25     else:
26         f_score = 0.0
27
28     return chamfer_dist, f_score

```

Normalized Metrics: The source code of normalize and align the predicted point cloud to the gt point cloud before evaluation.

```

1 """
2 The source code of normalize and align the predicted point cloud to
3 the
4 gt point cloud before evaluation.
5 """
6 import numpy as np
7 import open3d as o3d
8
9 def ray_cast_mesh(mesh, rays_origins, ray_directions):
10     """
11     Performs ray casting on a mesh to obtain a point cloud.
12     """
13     intersector = RayMeshIntersector(mesh)
14     index_triangles, index_ray, point_cloud = intersector.
15     intersects_id(
16         ray_origins=rays_origins,
17         ray_directions=ray_directions,
18         multiple_hits=True,
19         return_locations=True)
20     return index_triangles, index_ray, point_cloud

```

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper discusses the limitations of the work performed by the authors in the conclusion section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The theory assumptions in this paper have already been proved in previous works (Diffusion Model and Screened Poisson Algorithm).

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: The training and evaluation source code is uploaded so that the reviewer can check the details and re-implement our results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The training and evaluation source code is uploaded so that the reviewer can check the details and re-implement our results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The paper specify all the training and test details.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The paper adopts the mean metric values for 3 times testing for a fair comparison.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In the Experimental section Implementation section, the paper introduces the detailed implementation details, including the computing resources, training and evaluation details. Also, the source code is uploaded for checking.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The paper is sure to preserve anonymity.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: This paper contributes to fundamental 3D representation research and extends beyond just 3D generation tasks. In the future, it has the potential to enhance other research fields and benefit human society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The paper cites the original paper that produced the code package or dataset, and follows the licenses for existing assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The paper introduce the details of the dataset to be released.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowd-sourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.