LDM: Large Tensorial SDF Model for Textured Mesh Generation

Rengan Xie **Zhejiang University** State Key Lab of CAD&CG

rgxie@zju.edu.cn

Kai Huang Institute of Computing Technology, CAS ky.huang.null@gmail.com

> Oi Ye **Zhejiang University** qi.ye@zju.edu.cn

Wenting Zheng Zhejiang University State Key Lab of CAD&CG wtzheng@cad.zju.edu.cn

Yizheng Chen

Zhejiang Lab chenyizheng1230126.com

Qi Wang **Zhejiang University** wanina1995@gmail.com

Yuchi Huo **Zhejiang University** State Key Lab of CAD&CG Zhejiang Lab eehyc0@gmail.com

chenvis@zju.edu.cn

Wei Chen

Zhejiang University

Abstract

Previous efforts have managed to generate production-ready 3D assets from text or images. However, these methods primarily employ NeRF or 3D Gaussian representations, which are not adept at producing smooth, high-quality geometries required by modern rendering pipelines. In this paper, we propose LDM, a novel feed-forward framework capable of generating high-fidelity, illumination-decoupled textured mesh from a single image or text prompts. We firstly utilize a multi-view diffusion model to generate sparse multi-view inputs from single images or text prompts, and then a transformer-based model is trained to predict a tensorial SDF field from these sparse multi-view image inputs. Finally, we employ a gradient-based mesh optimization layer to refine this model, enabling it to produce an SDF field from which high-quality textured meshes can be extracted. Extensive experiments demonstrate that our method can generate diverse, high-quality 3D mesh assets with corresponding decomposed RGB textures within seconds. The project code is available at https://github.com/rgxie/LDM.

Keywords: 3D Generation model, Diffusion, Intrinsic decomposition, Relighting.

1. Introduction

Generating 3D content is vital for diverse applications and tasks. Recent 3D generation works leverage large 2D content generation diffusion models [27] to generate multiview consistent images based on certain conditions, and elevate the 2D prior knowledge within diffusion models to 3d space through score distillation sampling (SDS) [25, 14] or reconstruction methods [18, 19]. Despite the impressive results, these require additional optimization time during the forward generation of 3D assets, often taking tens of minutes.

LRM [6], on the other hand, propose an end-to-end 3D Large Reconstruction Model that predicts the 3D model of an object from a single input image in about 5 seconds, requiring no optimization time during the forward process. LRM utilizes a Transformer to encode the input image as a condition and predicts a NeRF of the object in tri-plane representation. Instant3d [11] enhances LRM by using multiview diffusion to generate multiple viewpoints from a single image. These views, encoded by an image encoder, serve as constraints to predict the target NeRF of the object, enabling the generation of higher-quality 3D assets from various perspectives. However, constrained by the overhead of volume rendering, LRM-based models can only be optimized under lower-resolution image patches. To achieve greater efficiency and higher resolution supervision, LGM and TMGS [33, 44] utilize 3DGS [10] as the representation of objects, employing UNet and Transformer architectures to predict and generate the Gaussian fields of objects. However, both the NeRF and 3DGS representations struggle with issues of unsmooth geometries and suboptimal geometric quality.

SDF representation has been proven capable of reconstructing high-quality. However, introducing SDF to the learning based 3D generation pipeline encounters the challenges of high memory computation and unstable convergence. Zero12345 and Zero12345++ [16, 15] introduce 3D CNNs or 3D diffusion to predict SDF volumes from image



Relighting

Material editing

Figure 1. Given a text prompt or a single image, our framework can generate corresponding high-quality 3D assets within seconds, including illumination-decoupled texture maps, facilitating integration into various applications, such as relighting and material editing.

inputs. Nevertheless, the quality of the output is constrained by the resolution of the volume grid, and both 3D diffusion and 3D CNNs incur significant computational overhead. As observed in CRM [37], predicting the correct geometry using pure RGB images is extremely hard, and it adds extra geometry information into the inputs of the pipeline to help the convergence of geometry generation with the SDF representation.

In addition to the geometry component, generating lighting-decomposed texture maps is crucial for editing and relighting 3D assets in downstream applications. However, existing learning-based 3D generation methods cannot provide lighting-decomposed texture maps, their textures embed shading effects such as shadows and highlights.

To generate 3D assets with high-quality meshes and illumination-decoupled textures within the end-to-end learning based framework, we propose LDM, a novel 3D generation pipeline with tensorial SDF representation and decoupled color field, as shown in Figure 1. Specifically, to improve the object surface quality and reduce memory requirements, we combine an SDF representation for the geometry with a low-rank tensorial representation in TensoRF[3]. Furthermore, to adapt to the light decoupling task in the generative approach, we introduce a shading model that separates the final output into an albedo color and a shading color.

To predict a decoupled representation, SDF values, albedo color, and shading color at any point in space are decoded from the tensorial representation and supervised by reference albedo and color images using volume rendering. However, the method used to convert SDF to a density field in reconstruction works, such as VolSDF [41], is not applicable to our generative approach. In these works, a parameter β is required to control the sharpness of the conversion and must be adjusted based on the convergence status of the SDF for each scene, while a large β means smoother geometry. Unlike reconstruction tasks that focus on a single object or scene, the generative approach needs to handle a variety of geometries, some smooth and others highly detailed. To address this issue, we introduce an adaptive β adjusting schedule, making β an optimizable parameter to allow β to be adjusted directly by the gradient descent from the final losses.

To achieve finer texture and geometry details, a differentiable iso-surface extraction layer like Flexicube [29] can be integrated to enable high-resolution image supervision following the tensorial SDF representation. Specifically, directly training to predict the SDF field using the Flexicube [29] rendering layer fails to converge due to sparse view supervision for each object and gradient descent impacting only the vicinity of the SDF grid vertices. To address this, a preliminary volume rendering training phase is essential. This phase samples points throughout the entire space to stabilize training and provide an initial SDF, allowing the Flexicube renderer to refine local quality. Extensive experiments show that our method is capable of generating diverse, high-quality 3D mesh assets along with decomposed RGB textures in just seconds.

In summary, our main contributions are:

• We propose the first feed-forward framework capable of generating high-quality meshes with illuminationdecoupled RGB textures from text or a single image input in just a few seconds.

- We introduce tensorial representation into the generation of objects, accurately representing illuminationdecoupled SDF fields, and enhancing convergence speed. An adaptive conversion of SDF to density strategy is proposed to enable the convergence of the novel representation in our feed-forward 3D generation framework.
- We integrate a gradient-based mesh optimization layer to train our framework, enabling it to produce an SDF field from which high-quality triangular meshes and illumination-decoupled textures can be extracted.

2. Related Work

Diffusion Models for Multi-view **Synthesis** Zero123 [18] demonstrates that large diffusion models have learned rich 3D priors about the visual world, even though they are only trained on 2D images. Furthermore, they introduce the first refined diffusion model capable of generating plausible images of objects from any viewpoint. A series of subsequent works [30, 31, 19, 20] have focused on enhancing the 3D consistency and resolution of multiview images generated by diffusion models. Zero123++ [30] proposes that generating consistent multi-view images hinges on accurately modeling their joint distribution and integrates six images in a 3×2 layout into a single frame for simultaneous multi-view generation during one diffusion process. MVDream [31] and Wonder3D [20] both recommend implementing multi-view attentions to enable feature sharing during the multi-view generation process. Free3D [43] introduces Plücker Embeddings of pixels in images' corresponding rays and a Ray Conditional Normalization layer during the diffusion generation process, which helps to produce multi-view images with more accurate perspectives. Although these methods have generated impressive multi-view images, achieving 3D consistency among them remains a challenge.

Lifting 2D Diffusion for 3D Generation Recent breakthroughs in multi-view diffusion model [27] have rapidly sparked interest [25, 18] in reconstructing 3D content generation from various conditions, such as Text and images, known as 2D-lifting, include distillation-based approaches and reconstructed-based approaches. *Distillation-based approaches*, exemplified by Dreamfusion and Magic3D [25, 14], leverage 2D diffusion models for optimizing 3D representations through score distillation sampling (SDS). *The reconstructed-based approaches* are notably represented by Zero123 and SyncDreamer [18, 19], which involves generating images with multi-view consistency using diffusion models, followed by employing volume rendering techniques to optimize 3D representations, such as neural radiance fields (NeRF) [22]or signed distance functions (SDF) [34]. These methods have achieved impressive results, but there are still some issues. The distillation approach tends to struggle with multi-faceted problems and requires a longer time to refine a single shape. While Zero123 [18] is capable of producing images from any viewpoint using a single input image and its relative pose, it faces challenges in maintaining consistency across multiview images, leading to 3D representations with blurred geometric details. Moreover, these methods require additional optimization time during the forward generation of 3D assets.

Feed-forward 3D Generative Models Currently, there are some feed-forward methods capable of generating 3D assets directly from text and image inputs without optimization. LRM [6] introduced the first end-to-end 3D Large Reconstruction Model, predicting a NeRF of the object from a single image in just 5 seconds. Since LRM only uses a single image as the condition for the transformer, the resulting object's back side exhibits sparse coloring, inconsistent with the front. Instant3D [11] improves LRM by generating multiple viewpoints from a single image through multi-view diffusion. The encoded views act as constraints to predict the object's NeRF, enhancing the quality of 3D assets from different perspectives. 3DGS [10] has emerged as a new attempt at object representation because it can quickly produce high-quality, high-resolution images from any viewpoint through splatting and rasterization. LGM [33]encodes attributes of 3D Gaussians from multiple views into splatter images [32] and uses a UNet to predict these images from RGB inputs.TMGS [44] uses a transformer to predict tri-plane attribute fields, enriching point cloud models of objects with additional properties to achieve a 3DGS representation. Nevertheless, both 3DGS and NeRF representations encounter difficulties with issues of uneven geometries and suboptimal geometric quality. SDF representation has been proven capable of reconstructing high-quality geometries Zero12345 and Zero12345++ [16, 15] use 3D CNNs or 3D diffusion to predict SDF volumes from images, but the output quality is limited by the volume grid's resolution. CRM [37] uses diffusion to predict six Canonical Coordinates Maps from a single input and optimizes a UNet to predict a tri-plane from these CCMS, which includes an SDF field. CRM integrates a differentiable iso-surface extraction layer [29] to optimize the generated SDF. Concurrent work InstantMesh [38] utilizes an LRM-based model to predict 3D assets represented as NeRF from multi-view images produced by multi-view diffusion. These NeRFs are subsequently refined into a Signed Distance Field (SDF). These efforts have produced impressive textured meshes; however, their textures embed shading effects such as shadows and highlights, making these meshes unsuitable for relighting and material editing. This limitation hinders their use in downstream applications.

3. Method

We propose the LDM, a framework as shown in Figure 2 that takes a single image or a text prompt as input and generates a corresponding triangular mesh equipped with illumination-decoupled textures. We first employ a multiview diffusion model to generate multiple viewpoint images of the target object, conditioned on either an input text prompt or image, as detailed in Section 3.1. In Section 3.2, we discuss the use of tensorial SDF representation for generating objects, which offers robust expressive power. Next, we introduce our two-stage training strategy. In Section 3.3, we introduce how to train a transformer-based model to generate tensorial SDF representations from sparse multiple views using volume rendering in the first stage. Finally, in the second stage detailed in Section 3.4, we introduce a gradient-based mesh optimization layer, Flexicube [29], to refine the generative model and enhance the quality of the extracted textures.

3.1. Conditional Multi-view Generation

Since Zero123 [18], numerous diffusion models capable of producing conditional multi-view images have emerged, such as MVdream, Imagedream, SyncDreamer, and Zero123Plus [31, 35, 19, 30]. We utilize MVDream for text inputs and ImageDream for image inputs. Both models are engineered to generate multi-view images from four orthogonal azimuths at a fixed elevation, which will serve as input conditions for the subsequent generation pipeline. Compared to other methods, MVDream and ImageDream employ joint training with synthetic data from the Objaverse dataset [4] and real data from the large-scale text-to-image (t2i) dataset, LAION5B [28]. Thus, We choose these models in our framework, aiming to boost our generalization capabilities for realistic object generation.

Although the implementation of a cross-view attention mechanism in those models has improved the consistency of the generated multi-view images, yielding impressive results, subtle 3D inconsistencies may still occur. To simulate this perturbation, inspired by previous methods [33, 6], we randomly apply a grid distortion to the input views and introduce a camera jitter to the camera poses during training. By deliberately introducing these inconsistencies during training, our framework is better equipped to handle the 3D inconsistencies encountered in images generated through multi-view diffusion.

3.2. Tensorial SDF and Decoupled Color Field

Recently, many studies have proposed various implicit neural representations for 3d objects or scenes, combined with differentiable rendering techniques, achieving impressive reconstruction results. NeRF [22] employs an MLP to encode the radiance field of a scene object and uses volume rendering to transform these volume densities and colors within the radiance field into an image from the target viewpoint. Unlike NeRF, TensoRF [3] represents the scene's radiance field with a 4D tensor. This tensor describes a 3D voxel grid, each voxel containing multi-channel features. TensoRF further decomposes this 4D tensor into several compact, low-rank components, greatly reducing memory requirements and improving rendering quality. NeRF-based methods often suffer from the issue of uneven geometries in reconstructed object surfaces. To address this, works such as VolSDF and NeuS [39, 34] propose replacing density with SDF in implicit representations to enhance the geometric quality of the reconstructions, resulting in smoother surfaces. Furthermore, TensoSDF [12], which combines tensorial representation with SDF, has proven effective in encoding various material features, such as albedo and metallic properties, in inverse rendering tasks. Inspired by these studies, we propose employing a tensorial SDF representation to depict both the illumination-decoupled color field and SDF field of the objects generated in this work.

Tensorial representation. We utilize the Vector-Matrix factorization method proposed by TensoRF as our tensorial representation. Specifically, our tensorial representation is formulated as follows:

$$V_p = V_k^X \circ M_k^{YZ} \oplus V_k^Y \circ M_k^{ZX} \oplus V_k^Z \circ M_k^{XY}, \quad (1)$$

where V_p denoted the feature vector of position p. In addition, V_k^m and $M_k^{\tilde{m}}$ represent the k-th vector and matrix factors of their corresponding spatial axes m, and \tilde{m} is the two axes orthogonal to m (e.g., $\tilde{X} = YZ$). \circ and \oplus denote the element-wise multiplication and concatenation operations. Moreover, unlike TensoRF, which uses two separate tensor fields to encode geometry and appearance, we leverage a single shared tensorial field for both. This approach enhances the correlation between geometry and appearance, as mentioned in [12].

Model SDF and decoupled color field. Our shading model differs from traditional PBR-based shading formulations and is more akin to intrinsic image decomposition works [13, 1, 8], where the color is divided into a reflectance component (albedo) and a shading component that captures illumination information in image space. This approximation has been proven to be a fast and effective approach. Using PBR-based shading models requires modeling and predicting the ambient lighting of the input condition images when generating 3D assets, which could lead to unstable model training and difficulty in convergence. Additionally, the condition images are generated by a diffusion model, they cannot ensure consistent lighting across multi-



Figure 2. **The overview of our framework.** When given an image or text prompt condition, we first utilize a diffusion model to generate multiple viewpoint images. These images are then encoded into image feature tokens using the DINO2 image encoder. Subsequently, these tokens are fed into a transform-based tensorial object reconstructor, resulting in a tensorial SDF representation. The tensorial SDF representation can be further rendered using volume rendering or the Flexicube render layer to produce images or extract meshes.

ple images, which further complicates the direct prediction of lighting. Adopting a simplified shading method that is closer to intrinsic image decomposition is a trade-off approach for use in feed-forward generative models.

Therefore, we decompose the final rendering color into two parts. The first part is the albedo color, which reflects the inherent color of the object and remains constant for each object. The second part is the shading color, which is determined by the varying lighting conditions the object receives and its different material properties. Given the decomposed appearance field and SDF field, we introduce multiple MLP decoders to decode the features extracted from the tensorial representation. This approach enables predicting various material properties and SDF from compact features. Formally,

$$s = \Theta_s(V_p, p), (c_a, c_r) = \Theta_c(V_p, p), \tag{2}$$

where Θ_s , Θ_c represent the MLPs used to decode the SDF value, albedo color, and shading color. Equipped with the albedo color c_a and shading colors c_s , the final rendering color c of point p can be calculated as,

$$c = c_a \cdot c_r. \tag{3}$$

Once the albedo map with decoupled illumination is obtained, artists can enhance the 3D assets with PBR materials such as metallic and roughness, tailoring them to specific use cases and making them ready for application. In new scenes, the shading color can be calculated based on different materials and lighting conditions, following the rendering equation [9].

Adaptive conversion of SDF to density. Currently, we can obtain the SDF of any point on the object, but

SDF cannot be directly used for volume rendering. Follow VolSDF [39], we model the conversion process from SDF value *s* to the corresponding density σ by the following formula:

$$\sigma = \begin{cases} \frac{1}{\beta} \left(1 - \frac{1}{2} exp\left(\frac{s}{\beta}\right) \right) & \text{if } s < 0, \\ \frac{1}{2\beta} exp\left(-\frac{s}{\beta}\right) \right) & \text{if } s \ge 0, \end{cases}$$
(4)

where $\beta > 0$ is an important parameter that controls the sharpness of the conversion from SDF to density. A large beta will result in the density field obtained from the SDF conversion being too smooth, causing the loss of geometric details. Conversely, a small beta will lead to geometric fragmentation.



Figure 3. Comparing model training performance across different Beta schedules.

To determine the appropriate beta value, several schemes are conducted. The first scheme is using the empirical mean of beta values from the reconstructed models. We reconstructed 200 scenes from multi-view images using VolSDF [41] and collected the optimized beta values from the 200 VolSDF models. The mean of these values sets the beta value to train our network. However, starting with a too-small beta made the network ultimately fail to converge. The second scheme is a linear scheduling mechanism for beta to increase training stability, following the approach of BakeSDF [40]. We initialize beta with a fixed value and gradually decrease it as the training epochs progressed, as $\beta = \beta_0 \left(1 + \frac{\beta_0 - \beta_1}{\beta_1} t^{-1}\right)$, where t goes from 0 to 1 during training, $\beta_0 = 0.1$, and β_1 for the three ray points hierarchical sampling is 0.015, 0.003, and 0.001, respectively. Unfortunately, this also failed. We attribute it to that the manually set beta decay strategy might not align with the convergence rate of our generation network. Therefore, we introduce an adaptive beta adjust schedule and allow the beta value to be adjusted directly by the gradient descent process instead of a manual schedule. Specifically, we initially set beta to a relatively large value of 0.1 and made it an optimizable parameter. We train it together with the image encoder and tensorial object reconstructor within the framework. We found that this approach allows beta to gradually decrease as the model acquired more knowledge, eventually converging with the model. As seen in Figure 3, this technique effectively guides the model toward convergence.

3.3. Feed-forward Large Reconstruction Model

In this section, we describe how a transformer-based model reconstructs a tensorial representation from sparse view images generated by multi-view diffusion. Reconstructing 3D representations from sparse view inputs has always been a challenge in the community, with traditional methods struggling to incorporate and apply prior 3D knowledge. In recent years, diffusion models trained on vast datasets have demonstrated a robust understanding of prior knowledge, capable of generating images under various controlled conditions. This has sparked interest in exploring large models for 3D generation. Inspired by Instant3D and LRM [11, 6], we propose a new transformer-based model architecture that includes an image encoder and a tensorial object reconstructor to predict a tensorial SDF field conditioned on multi-view image features

Specifically, we use a vision transformer, DINO2 [2], to encode feature tokens F_I from multi-view images as shown in Figure 2. To ensure that the resulting feature tokens incorporate multi-view camera information, we modulate the camera information as described in LRM [6] and inject it into the image encoder using AdaLN [7, 24]. Furthermore, we arrange the tensor vectors V_k^m and $M_k^{\tilde{m}}$ into learnable tokens and feed these tokens into a feature decoder composed of a sequence of transformer layers, resulting in a series of tensor tokens. It should be noted that the information from the conditioned image feature F_I is connected to the cross-attention layer within these transformer layers. Finally, we reshape and upsample the tensor tokens into the final tensorial SDF representation. Then, all the transformer architecture can be trained in an end-to-end manner using image reconstruction loss at novel views images rendered from the tensorial SDF representation using volume rendering [22]. More details of the network structure can be found in the supplementary material. The training loss in this stage is defined as,

$$\mathcal{L}_{1} = \mathcal{L}_{mse}(I_{rgb}, I_{rgb}^{GT}) + \mathcal{L}_{mse}(I_{\alpha}, I_{\alpha}^{GT}) + \mathcal{L}_{mse}(I_{h}, I_{h}^{GT}) + \lambda_{vgg}\mathcal{L}_{lpips}(I_{rgb}, I_{rgb}^{GT}),$$
(5)

where I_{rgb} , I_{α} , I_{h} respectively represent the predicted final color image, albedo image, and mask image. Following LRM [6], we apply mean square error loss and VGG-based LPIPS loss [42] to the predicted images and corresponding reference images. During training, we set $\lambda_{vgg} = 2$.

3.4. Lifting SDF to Fine Mesh

Till now, we have introduced a transformer structure that enables the generation of a tensorial SDF representation from sparse images. We can directly extract geometric surfaces from the SDF representation using Marching Cubes [21] after training. However, the strategy has the issue that given the high computational cost of volume rendering, only a patch of the image is rendered during training, thus being unable to utilize full-resolution supervisory images fully. In addition, a gradient-based mesh optimization layer after an SDF field can improve the silhouette quality as described in Nvdiffrec [23].

FlexiCubes [29] is a state-of-the-art differentiable isosurface extraction layer. Employing the differentiable method to extract meshes from the tensorial SDF representation, rasterizing, and rendering at high resolution for supervised learning is beneficial for achieving finer details. One strategy is to use the Flexicube render layer from the outset to replace the computationally expensive volume rendering. However, due to the design of the discrete SDF grid, under sparse view supervision, each gradient descent step affects only the vicinity of the SDF grid vertex. Therefore, directly using Flexicube from the start during training cannot get the network to converge. CRM [37] has similar observations: predicting the correct geometry directly with Flexicube is extremely difficult using pure RGB images. In contrast, volume rendering sampling points throughout the entire space and gradients affect the entire space.

Two-stage training from global to local. Therefore, we combine the advantages of both training with volume rendering and Flexicube and propose a two-stage training strategy. In the first stage, we use the more stable volume rendering for the learning of the global features of the network, allowing the network to achieve global convergence.



Figure 4. Qualitative comparison with baselines shows that our method produces high-quality 3D assets with smooth geometry and clear textures, which align well with the input image.

In the second stage, we initialize the training with the model from the first stage and use Flexicube to optimize the local features, achieving higher resolution textures. During the second stage, to predict a set of weights in each grid cell and the deformation of each grid vertex following Flexicube [29], we additionally introduced two MLP decoders to decode the tensor features from the tensorial SDF:

$$d = \Theta_d(V_p, p), w = \Theta_d(V_p, p), \tag{6}$$

where Θ_d , Θ_w represent the MLPs used to decode the grid point deformation value and weights of the grid cell, which are needed for the FlexiCube rendering pipeline, respectively. The model is then trained under supervision using the following loss:

$$\mathcal{L}_{2} = \mathcal{L}_{1} + \lambda_{d} \left\| I_{d} - I_{d}^{gt} \right\|_{1} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \tag{7}$$

where I_d represent the rendered depth map, I_d^{gt} are the corresponding reference images. The L_{reg} is the regularization term following FlexiCubes [29]. During training, we set $\lambda_d = 0.5$, $\lambda_{reg} = 0.005$.

4. Experiment

4.1. Implementation Details

Training Datasets We train the model in GObjaverse from RichDreamer [26], which is rendered using the TIDE renderer on Objaverse [4] and includes G-buffer rendering

data such as albedo, RGB, depth, and normal map images from multiple views in resolution of 512×512 . Following LGM [33], we utilize a filtered subset, which excludes lowquality 3D models, resulting in a final set of around 80K 3D objects. Specifically, this dataset includes 36 random views of a centered object, with elevations ranging from -5° to 30° , and two additional views for the top and bottom. Although the camera poses for images produced by multi-view diffusion are predetermined, during training, a random subset of 8 images is selected. 4 images are used as inputs to the model, and the remaining 4 are employed for novel view supervision. This strategy not only makes the model robust to inconsistencies in inference inputs but also ensures that the image encoder is sensitive to camera poses. For inference, images with fixed camera poses generated in the first stage are fed to the tensorial SDF reconstruction.

Table 1. The features of different methods, our framework provides comprehensive feature support.

| Method | Tto3d | Ito3d | Rep. | Illum. dec. |
|-----------|--------------|--------------|----------|--------------|
| LRM | × | \checkmark | NeRF | × |
| LGM | \checkmark | \checkmark | 3DGS | × |
| CRM | × | \checkmark | SDF | × |
| Instand3d | \checkmark | \checkmark | NeRF | × |
| Ours | \checkmark | \checkmark | TensoSDF | \checkmark |

Training details We train our model in two stages. In the first stage, we train using a volume rendering pipeline



Figure 5. The effect of illumination decoupled texture. We perform relighting in new scenes for both illumination-decomposed textures and non-decomposed textures. The 3D assets without illumination decomposition display incorrect shadows in the new scenes.



Figure 6. Comparing model training performance across different object representations, Tensorial SDF emerges as the outstanding performer.

and loss \mathcal{L}_1 defined as Equation (5). Specifically, we render random patches in size 128×128 cropped from the original 512×512 resolution images during training, which conserves GPU memory while increasing local resolution. However, we observed that initiating training with too large

a scaling factor can impede model convergence. Therefore, we progressively increase the original resolution linearly with epochs, from 192 to 512. Additionally, we increase the likelihood of selecting crops covered by the foreground mask. This stage takes 4 days with batch size 32. We used the Adam optimizer with a learning rate of 4e-4, weight decay of 0.05, and betas of (0.9, 0.95). The learning rate is cosine annealed to 0 during the training. In the second stage, we train the model with Flexicube pipeline with full resolution using the loss defined as Equation (7). This stage takes 2 days with batch size 32. We use the same optimization settings as in the first stage, merely adjusting the learning rate to 1e-5. All training is conducted on 16 NVIDIA A6000 48GB GPUs.

4.2. Comparison

We compare our method with the previous state-of-theart feed-forward generation methods, including One-2-3-45 [17], LRM [6], CRM [37] and LGM [33]. We compare the different features of various methods in Table 1. We

Table 2. Quantitative results demonstrate that the color and geometric quality of the 3D assets generated by our method outperforms other methods.

| | PSNR ↑ | SSIM \uparrow | LPIPS \downarrow | $\text{CD}\downarrow$ | IoU↑ |
|------------|--------|-----------------|--------------------|-----------------------|--------|
| One-2-3-45 | 18.93 | 0.779 | 0.166 | 0.0614 | 0.4126 |
| CRM | 21.59 | 0.864 | 0.159 | 0.0335 | 0.4213 |
| OpenLRM | 20.29 | 0.829 | 0.186 | 0.0482 | 0.3731 |
| LGM | 20.05 | 0.798 | 0.176 | 0.0417 | 0.4331 |
| Ours | 22.52 | 0.873 | 0.143 | 0.0241 | 0.4361 |

evaluate our method on the Google Scanned Object (GSO) dataset [5]. To quantitatively evaluate the image quality synthesized by our approach, we adopt three standard metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM) [36], and Learned Perceptual Image Patch Similarity (LPIPS) [42]. Regarding geometric quality, we report the Chamfer Distance (CD) and Volume Intersection over Union (IoU) as metrics. The quantitative results, as shown in Table 2, demonstrate that our method outperforms others in both color and geometry. Additionally, we show a qualitative comparison of various methods in the Figure 4. As shown in the figure, these methods have all achieved good results, but some of them struggle with certain detailed aspects. One2345 [17] is limited by the discrete SDF volume grid representation, resulting in blurred textures and geometry. LRM [6], lacking multi-view images as constraints, produces 3D models with inconsistent back and front colors and suffers from multi-face Janus problem [31]. LGM [33], using 3DGS as its representation, tends to produce blurred details in areas with dense geometry, such as leaf clusters. CRM [37] exhibits geometric discontinuities in the frog's hands and the flower's petioles. However, our method successfully generates 3D assets that are well-aligned with the input conditions, featuring smooth and intact geometric structures.

4.3. Ablation Study

The effect of illumination decoupled texture. As shown in Figure 5, We generate 3d asserts using LDM, applying both illumination-decoupled albedo texture and original RGB texture. Then, we relight the generated results in new scenes. It can be observed that the area highlighted by the red arrow contains excessive shadows, which is incorrect. Without proper illumination decomposition, the shadows baked into the RGB texture from the original lighting combine with the shadows of the new scene, leading to inaccurate results. Therefore, the generation of illuminationdecoupled 3D assets plays a crucial role in the usability of downstream applications.

Comparisons across different object representation. We conducted ablation studies on various object representations to validate their expressiveness and training conver-



Input Tensorial SDF* Triplane SDF Tensorial SDF* Triplane SDF

Figure 7. Qualitative comparison results of Tensorial SDF* and Triplane SDF. Note, for a fair comparison, the Tensorial SDF* used here is a scaled small model aligned with Triplane SDF, not the full model. Quantitative effect of the Flexlayer, evaluate the performance of the fine-tuned model with varying numbers of input views. More discussion can be find in Section 4.3.

| Method | PSNR↑ | SSIM↑ | LPIPS↓ |
|-------------------|-------|-------|--------|
| w/ Flex | 23.21 | 0.869 | 0.131 |
| w/o Flex | 22.76 | 0.853 | 0.145 |
| w/ view finetune | 22.89 | 0.856 | 0.149 |
| w/o view finetune | 21.63 | 0.786 | 0.157 |

gence speed. Note that due to the substantial training costs associated with our final model, the subsequent analytical experiments employ a significantly reduced version of the LDM model. Inspired by previous works [6, 11], we reduced the training dataset to a subset containing 10k objects to accelerate model convergence. We conducted training without image cropping using volume rendering, rendering the images at a fixed resolution of 128×128 . Each model was trained on 8 NVIDIA A6000 GPUs over 40 epochs. The quantitative results, as shown in Figure 6, indicate that compared to the representation of Triplane SDF, the Tensorial SDF representation achieves better quality under the same number of epochs of training and converges faster. Furthermore, it can be observed that constructing objects as Tensorial SDF rather than Tensorial NeRF also aids in convergence. We believe this may be because the SDF representation introduces stronger smoothness constraints, reducing geometric disintegration. Additionally, qualitative comparisons of results, as illustrated in Figure 7, demonstrate that Tensorial SDF exhibits superior geometric details with smoother surfaces (e.g., the car hood).

The effect of Flexicubes layer We conducted ablation experiments to verify the impact of introducing the Flexicube layer for model fine-tuning. As shown in Figure 8 and Table 3, training with the Flexicubes layer effectively enhances the texture clarity, thanks to the utilization of higher resolution during training. In addition, introducing depth constraints leads to smoother geometric surfaces.

Comparisons across different diffusion methods As most text and image-conditioned multi-view diffusion models generate 4 different views, our framework is trained un-



Figure 8. Ablation study shows that the Flexicubes layer helps improve the texture details of generated objects.

der the assumption of having 4 input views. We evaluated the performance of our framework when switching between multi-view images generated by different diffusion models. Specifically, we try to switch the input of our model to the image predicted by popular multi-view diffusion models Zero123plus [30]. The experimental results, as shown in Figure 10 and Table 3, indicate that when we directly employ a model trained with 4 views to reconstruct 6 views as input, it does not perform well and produces some misaligned results. However, after fine-tuning our model with 6 views input for only 10k iterations, it is able to generate reliable results. This experiment demonstrates that our framework exhibits good generalization across different multiview diffusion results.

5. Application

As our method generates 3D assets with illuminationdecomposed texture maps, it easily supports applications such as relighting and material editing. As shown in Figure 9, we relight the 3D assets in different new scenes, producing convincing synthesized images. Additionally, we apply PBR material textures to the 3D assets and edit their metallicity, roughness, and albedo properties, with the resulting images shown in Figure 11.

6. Limitations

While our model can generate high-quality illuminationdecoupled 3D assets, there are still some limitations. Like other transformer-based approaches [11, 6], the size of the tensorial SDF tokens produced by our model is capped at 64x64, constraining the resolution of the final 3D assets. In addition, our illumination-decoupled module is not designed to handle complex materials, such as translucent surfaces, due to our simplified rendering composition. Finally, the quality of our reconstructed 3D models is influenced by the inconsistencies in the multi-view diffusion. We reserve these challenges for future exploration. Further discussion can be found in the appendix Section C.

7. Conclusion

This paper introduces a novel feed-forward framework capable of reconstructing 3D meshes with illuminationdecoupled textures from a single image input or text prompt in 10 seconds. Our method utilizes a conditional multiview diffusion model to generate consistent four-view images and lifts them to 3D using a transformer-based large reconstruction model. We propose to generate objects represented as tensorial SDF field, which is more expressive and can accelerate model convergence compared to previous triplane representations. In addition, we decompose the rendering color into albedo color and shading color, which enables the generated 3D assets to be easily relit and edited for material properties. We believe that the ability to produce relightable 3D assets is very important for downstream applications.

Acknowledgement

This work is partially supported by Zhejiang Province "Jianbing" Research and Development Project (2023C01042), National Key R&D Program of China (No. 2024YDLN0011), NSFC (No. 62441205), and NSFC (No. U22B2034). We would like to thank all reviewers for their detailed feedback.

Appendices

A. Detail for Network architecture

We use the DINOv2-ViT-B/14 [2] as our image encoder, a transformer-based model, which has 12 layers and the hidden dimension of the transformer is 768. Regarding camera features, we flatten the extrinsic parameters of each view into 16 dimensions and encode them into a 1024dimensional vector using a 2-layer MLP. After encoding, each image yields 257 image feature tokens, including the [CLS] token. All tokens from these views are concatenated together to obtain a total of N×257 tokens, where N is the number of input images. These image feature tokens serve as condition features in the subsequent generation process.

Next, we use a transformer-based tensorial object reconstructor to predict a tensorial SDF representation from a sequence of learnable tokens with a size of $(3\times32\times32+3\times32)\times1024$, where $3\times32\times32$ and 3×32 are the number token numbers align with V_k^m and $M_k^{\tilde{m}}$ represent the tensor matrix factors and vector of their corresponding spatial axes. And 1024 is the hidden dimension of the transformer decoder. After being decoded by the transformer, conducted through cross attention from image features, we obtain the same number of tensor tokens. Next, following LRM[6], we use a de-convolution layer



Figure 9. Relighting results of 3d assets generated by our framework.



Figure 10. blation study for transferring the model to a different number of inputs views shows that good results can be achieved with simple fine-tuning.

and an MLP layer to upscale tensor matrix factors from $3 \times (32 \times 32) \times 1024$ to $3 \times (64 \times 64) \times 40$, and tensor vector from $3 \times 32 \times 1024$ to $3 \times 64 \times 40$.

Finally, these tensor tokens are reshaped into a tensorial SDF representation, from which we can compute the feature vector of any point with a dimension of 120. Then, for the first stage of volume rendering training, we use a 4-layer MLP with 64 hidden dimensions to decode a 1-dim SDF value from this feature. We use another 4-layer MLP with 64 hidden dimensions to decode a 6-dim color value

from this feature, where the first three dimensions are considered the albedo color and the last three dimensions are considered the shading color. Finally, for the second stage of Flexcubes layer [29] training, we use a 2-layer MLP with 64 hidden dimensions to decode an 8-dim weights value from this feature. We use another 2-layer MLP with 64 hidden dimensions to decode a 1-dim deformation value from this feature.

B. More results for 3D generation task

As shown in Figure 12, we present more 3D results generated from text prompts or single image inputs. Our method is capable of producing good results for both realworld images and unreal images.

C. Failure Cases and Limitations

As discussed in the main text, our framework can produce convincing results from a single image or text prompt input, but there are still certain instances where it fails, as shown in Figure 13. In the left side of Figure 13, we input an image of a potted plant to predict its 3D assets. The overall result meets expectations, with the plant's leaves and base being well generated, but the thin stems fail to generate properly. There are multiple reasons for this issue. On the one hand, the size of the tensorial SDF tokens produced by our model is capped at 64x64, which may result



Figure 11. We applied different material textures to the generated 3D assets and modified the albedo color to produce some interesting results, such as a metallic hamburger.

in the loss of fine geometric details. On the other hand, the multi-view images generated by the multi-view diffusion may be inconsistent, causing thin stems to misalign across different views, leading to generation failures. On the right side of Figure 13, we can see that due to misalignment in the multi-view images, the text on the predicted fire extinguisher appears blurred. The misaligned highlights also caused errors in the material prediction, with the highlights being incorrectly predicted as part of the albedo. Additionally, the prediction of the shading color for the metallic and lacquer materials on the fire extinguisher's surface is inaccurate. While part of the red metallic surface is correctly predicted in the shading color, some of the black text is mistakenly predicted as part of the shading color as well.

References

- C. Careaga and Y. Aksoy. Intrinsic image decomposition via ordinal shading. ACM Transactions on Graphics, 43(1):1– 24, 2023. 4
- [2] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 6, 10
- [3] A. Chen, Z. Xu, A. Geiger, J. Yu, and H. Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, pages 333–350. Springer, 2022. 2, 4
- [4] M. Deitke, R. Liu, M. Wallingford, H. Ngo, O. Michel, A. Kusupati, A. Fan, C. Laforte, V. Voleti, S. Y. Gadre, et al. Objaverse-xl: A universe of 10m+ 3d objects. *Advances in Neural Information Processing Systems*, 36, 2024. 4, 7
- [5] L. Downs, A. Francis, N. Koenig, B. Kinman, R. Hickman, K. Reymann, T. B. McHugh, and V. Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. In 2022 International Conference on Robotics and Automation (ICRA), pages 2553–2560. IEEE, 2022. 9
- [6] Y. Hong, K. Zhang, J. Gu, S. Bi, Y. Zhou, D. Liu, F. Liu, K. Sunkavalli, T. Bui, and H. Tan. Lrm: Large recon-

struction model for single image to 3d. *arXiv preprint arXiv:2311.04400*, 2023. **1**, **3**, **4**, **6**, **8**, **9**, **10**

- [7] X. Huang and S. Belongie. Arbitrary style transfer in realtime with adaptive instance normalization. In *Proceedings of the IEEE international conference on computer vision*, pages 1501–1510, 2017. 6
- [8] M. Janner, J. Wu, T. D. Kulkarni, I. Yildirim, and J. Tenenbaum. Self-supervised intrinsic image decomposition. *Advances in neural information processing systems*, 30, 2017.
- [9] J. T. Kajiya. The rendering equation. In Proceedings of the 13th annual conference on Computer graphics and interactive techniques, pages 143–150, 1986. 5
- [10] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023. 1, 3
- [11] J. Li, H. Tan, K. Zhang, Z. Xu, F. Luan, Y. Xu, Y. Hong, K. Sunkavalli, G. Shakhnarovich, and S. Bi. Instant3d: Fast text-to-3d with sparse-view generation and large reconstruction model. *arXiv preprint arXiv:2311.06214*, 2023. 1, 3, 6, 9, 10
- [12] J. Li, L. Wang, L. Zhang, and B. Wang. Tensosdf: Roughness-aware tensorial representation for robust geometry and material reconstruction. arXiv preprint arXiv:2402.02771, 2024. 4
- [13] Z. Li and N. Snavely. Learning intrinsic image decomposition from watching the world. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9039–9048, 2018. 4
- [14] C.-H. Lin, J. Gao, L. Tang, T. Takikawa, X. Zeng, X. Huang, K. Kreis, S. Fidler, M.-Y. Liu, and T.-Y. Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 300–309, 2023. 1, 3
- [15] M. Liu, R. Shi, L. Chen, Z. Zhang, C. Xu, X. Wei, H. Chen, C. Zeng, J. Gu, and H. Su. One-2-3-45++: Fast single image to 3d objects with consistent multi-view generation and 3d diffusion. arXiv preprint arXiv:2311.07885, 2023. 1, 3
- [16] M. Liu, C. Xu, H. Jin, L. Chen, M. Varma T, Z. Xu, and H. Su. One-2-3-45: Any single image to 3d mesh in 45 sec-



Figure 12. More predicted results for text-to-3D generation task.

onds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3

- [17] M. Liu, C. Xu, H. Jin, L. Chen, Z. Xu, H. Su, 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. 8, 9
- [18] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF International Con-*

ference on Computer Vision, pages 9298–9309, 2023. 1, 3, 4

- [19] Y. Liu, C. Lin, Z. Zeng, X. Long, L. Liu, T. Komura, and W. Wang. Syncdreamer: Generating multiviewconsistent images from a single-view image. *arXiv preprint arXiv:2309.03453*, 2023. 1, 3, 4
- [20] X. Long, Y.-C. Guo, C. Lin, Y. Liu, Z. Dou, L. Liu, Y. Ma, S.-H. Zhang, M. Habermann, C. Theobalt, et al. Wonder3d: Single image to 3d using cross-domain diffusion. arXiv



Figure 13. Some failure cases in the predictions reveal the limitations of our method.

preprint arXiv:2310.15008, 2023. 3

- W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353. 1998. 6
- [22] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European conference on computer vision*, pages 405–421. Springer, 2020. 3, 4, 6
- [23] J. Munkberg, J. Hasselgren, T. Shen, J. Gao, W. Chen, A. Evans, T. Müller, and S. Fidler. Extracting triangular 3d models, materials, and lighting from images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8280–8290, 2022. 6
- [24] W. Peebles and S. Xie. Scalable diffusion models with transformers. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 4195–4205, 2023. 6
- [25] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988, 2022. 1, 3
- [26] L. Qiu, G. Chen, X. Gu, Q. zuo, M. Xu, Y. Wu, W. Yuan, Z. Dong, L. Bo, and X. Han. Richdreamer: A generalizable normal-depth diffusion model for detail richness in text-to-3d. arXiv preprint arXiv:2311.16918, 2023. 7
- [27] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 10684– 10695, 2022. 1, 3
- [28] C. Schuhmann, R. Beaumont, R. Vencu, C. Gordon, R. Wightman, M. Cherti, T. Coombes, A. Katta, C. Mullis, M. Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. 4

- [29] T. Shen, J. Munkberg, J. Hasselgren, K. Yin, Z. Wang, W. Chen, Z. Gojcic, S. Fidler, N. Sharp, and J. Gao. Flexible isosurface extraction for gradient-based mesh optimization. *ACM Transactions on Graphics (TOG)*, 42(4):1–16, 2023. 2, 3, 4, 6, 7, 11
- [30] R. Shi, H. Chen, Z. Zhang, M. Liu, C. Xu, X. Wei, L. Chen, C. Zeng, and H. Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 3, 4, 10
- [31] Y. Shi, P. Wang, J. Ye, M. Long, K. Li, and X. Yang. Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512, 2023. 3, 4, 9
- [32] S. Szymanowicz, C. Rupprecht, and A. Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. arXiv preprint arXiv:2312.13150, 2023. 3
- [33] J. Tang, Z. Chen, X. Chen, T. Wang, G. Zeng, and Z. Liu. Lgm: Large multi-view gaussian model for high-resolution 3d content creation. *arXiv preprint arXiv:2402.05054*, 2024. 1, 3, 4, 7, 8, 9
- [34] P. Wang, L. Liu, Y. Liu, C. Theobalt, T. Komura, and W. Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. 3, 4
- [35] P. Wang and Y. Shi. Imagedream: Image-prompt multi-view diffusion for 3d generation. arXiv preprint arXiv:2312.02201, 2023. 4
- [36] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 9
- [37] Z. Wang, Y. Wang, Y. Chen, C. Xiang, S. Chen, D. Yu, C. Li, H. Su, and J. Zhu. Crm: Single image to 3d textured mesh with convolutional reconstruction model. *arXiv* preprint arXiv:2403.05034, 2024. 2, 3, 6, 8, 9

- [38] J. Xu, W. Cheng, Y. Gao, X. Wang, S. Gao, and Y. Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 3
- [39] L. Yariv, J. Gu, Y. Kasten, and Y. Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 4, 5
- [40] L. Yariv, P. Hedman, C. Reiser, D. Verbin, P. P. Srinivasan, R. Szeliski, J. T. Barron, and B. Mildenhall. Bakedsdf: Meshing neural sdfs for real-time view synthesis. In ACM SIGGRAPH 2023 Conference Proceedings, pages 1–9, 2023.
- [41] L. Yariv, Y. Kasten, D. Moran, M. Galun, M. Atzmon, B. Ronen, and Y. Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33:2492–2502, 2020. 2, 6
- [42] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
 6, 9
- [43] C. Zheng and A. Vedaldi. Free3d: Consistent novel view synthesis without 3d representation. arXiv preprint arXiv:2312.04551, 2023. 3
- [44] Z.-X. Zou, Z. Yu, Y.-C. Guo, Y. Li, D. Liang, Y.-P. Cao, and S.-H. Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. arXiv preprint arXiv:2312.09147, 2023. 1, 3