## Tex4D: Zero-shot 4D Scene Texturing with Video Diffusion Models

Jingzhi Bao<sup>1</sup> Xueting Li<sup>2</sup> Ming-Hsuan Yang<sup>3</sup> <sup>1</sup>CUHK-Shenzhen <sup>2</sup>NVIDIA <sup>3</sup>UC Merced https://tex4d.github.io



Figure 1. Tex4D Application. Our synthesized dynamic textures can be easily integrated into graphics pipelines.

#### Abstract

3D meshes are extensively employed in movies, games, AR, and VR for their efficiency in animation and minimal memory footprint, leading to the creation of a vast number of mesh sequences. However, creating dynamic textures for these mesh sequences to model the appearance transformations remains laborintensive for professional artists. In this work, we present **Tex4D**, a zero-shot approach that creates multi-view and temporally consistent dynamic mesh textures by integrating the inherent 3D geometry knowledge with the expressiveness of video diffusion models. Given an untextured mesh sequence and a text prompt as inputs, our method enhances multi-view consistency by synchronizing the diffusion process across different views through latent aggregation in the UV space. To ensure temporal consistency, such as lighting changes, wrinkles, and appearance transformations, we leverage prior knowledge from a conditional video generation model for texture synthesis. Straightforwardly combining the video diffusion model and the UV texture aggregation leads to blurry results. We analyze the underlying causes and propose a simple yet effective modification to the DDIM sampling process to address this issue. Additionally, we introduce a reference latent texture to strengthen the correlation between frames during the denoising process. To the best of our knowledge, Tex4D is the first method specifically designed for 4D scene texturing. Extensive experiments demonstrate its superiority in producing multi-view and multi-frame consistent dynamic textures for mesh sequences.



Figure 2. Given an untextured mesh sequence and a text prompt as inputs (Top), **Tex4D** generates multi-view, dynamic textures. Below, we show renderings of the textured meshes from three views and four timestamps.

## **1** Introduction

3D meshes are widely used in modeling, computer-aided design (CAD), animation, and computer graphics due to their low memory footprint and efficiency in animation. Visual artists, game designers, and movie creators build numerous animated mesh sequences for visual applications. However, creating vivid videos involves complex post-processing steps, such as creating dynamic textures for appearance transformations, as shown in Fig. 1. These steps are labor-intensive and require specialized expertise by artists.

On the other hand, recent advancements in generative models have democratized content creation and demonstrated impressive performance in image and video synthesis. For instance, video generation models [21, 16, 29, 20, 70, 81, 23, 69, 77, 66, 11, 12] trained on large-scale video datasets [2, 48] allow users to create realistic video clips from various inputs such as text prompts, images, or geometric conditions. However, these text-to-video generation models, which are trained solely on 2D data, often struggle with spatial consistency when applied to multi-view image generation [54, 50, 33, 61, 35, 49, 27, 53, 58] or 3D object texturing [7, 34, 45, 24].

To address these limitations, two main approaches have been developed. One approach [45, 10, 7] focuses on resolving multi-view inconsistency in static 3D object texturing by synchronizing multi-view image diffusion processes. While these methods produce multi-view consistent textures for static 3D objects, they do not address the challenge of generating dynamically changing textures for mesh sequences. Another approach [18, 32, 39] aims to generate video clips based on the rendering (e.g., depth, normal or UV maps) of an untextured mesh sequence. To encourage temporal consistency, these methods modify the attention mechanism in 2D diffusion models and utilize inherent correspondences in a mesh sequence to facilitate feature synchronization between frames. Although these techniques can be adapted for multi-view image generation by treating camera pose movement as temporal motion, they usually produce inconsistent 3D texturing due to insufficient exploitation of 3D geometry priors.

In this paper, we introduce a novel task: 4D scene texturing. Given an animated untextured 3D mesh sequence and a text prompt, our goal is to generate dynamic textures that are both temporally and multi-view consistent. We aim to texture 4D scenes while capturing temporal variations, such as lighting and appearance changes, to produce vivid visual results—a key requirement in downstream tasks like character generation. Unlike existing works, we fully leverage 3D geometry knowledge from the mesh sequence to enforce multi-view consistency. Specifically, we develop a method that synchronizes the diffusion process from different views through latent aggregation in the UV space. To ensure temporal consistency, we employ prior knowledge from a conditional video generation model for texture sequence synthesis and introduce a reference latent texture to enhance frame-to-frame correlations during the denoising process. However, naively integrating the UV texture

aggregation into the video diffusion process causes the variance shift problem, leading to blurry results. To resolve this issue, we propose an effective modification to the DDIM [52] sampling process by rewriting the equation. Our method is computationally efficient thanks to its zero-shot nature. The textured mesh sequence can be rendered from any camera view, thus supporting various applications in content creation. Our key contributions are:

- We present **Tex4D**, a zero-shot pipeline for generating high-fidelity dynamic textures that are temporally and multi-view consistent, utilizing video diffusion models and mesh sequence controls.
- To leverage priors from existing video diffusion models, we develop an effective modification to the DDIM sampling process to address the variance shift issue caused by multi-view texture aggregation and design a background learning module.
- We introduce a reference UV blending mechanism to establish correlations during the denoising steps, addressing self-occlusions, and synchronizing the diffusion process in invisible regions.
- Our method is not only computationally efficient, but also demonstrates comparable if not superior performance to various state-of-the-art baselines.

#### 2 Related Work

**Video Stylization and Editing.** Video diffusion models have shown remarkable performance in the field of video generation. These models learn motions and dynamics from large-scale video datasets using 3D-UNet to create high-quality, realistic, and temporally coherent videos. Although these approaches show compelling results, the generated videos lack fine-grained control, inhibiting their application in stylization and editing. To solve this issue, inspired by ControlNet [74], SparseCtrl [18] trains a sparse encoder from scratch using frame masks and sparse conditioning images as input to guide the video diffusion model. CTRL-Adapter [32] proposes a trainable intermediate adapter to connect the features between ControlNet and video diffusion models.

Meanwhile, [56] observed that the spatial features of T2I models play an influential role in determining the structure and appearance, Text2Video-Zero [25] uses a frame-warping method to animate the foreground object by T2I models and [63, 8, 42] propose utilizing self-attention injection and cross-frame attention to generate stylized and temporally consistent video using DDIM inversion [52]. Subsequently, numerous works [78, 6, 68, 17, 15] generate temporally consistent videos utilizing T2I diffusion models by spatial latent alignment without training. However, the synthesized videos usually show flickerings due to the empirical correspondences, such as feature embedding distances and UV maps, which are insufficient to express the continuous content in the latent space. Another line of work [51, 3, 5, 67, 19] is to train additional modules on large-scale video datasets to construct feature mappings, for example, Text2LIVE [3] applies test-time training with the CLIP loss, and MagicAnimate [67] introduced an appearance encoder to retain intricate clothes details.

**Texture Synthesis.** With the rapid development of foundation models, researchers have focused on applying their generation capability and adaptability to simplify the process of designing textures and reduce the expertise required. To incorporate the result 3D content with prior knowledge, earlier works [26, 38, 13] jointly optimize the meshes and textures from scratch with the simple semantic loss from the pre-trained CLIP [43] to encourage the 3D alignment between the generated results and the semantic priors. However, the results show apparent artifacts and distortion because the semantic feature cannot provide fine-grained supervision during the generation of 3D content.

DreamFusion [41] and similar models [31, 60, 40, 37, 9] distill the learned 2D diffusion priors from the pre-trained diffusion models [46] to synthesize the 3D content by Score Distillation Sampling (SDS). These methods render 2D projections of the 3D asset parameters and compare them against reference images, iteratively refining the 3D asset parameters to minimize the discrepancy of the target distribution of 3D shapes learned by the diffusion model. Although these approaches enable people without expertise to generate detailed 3D content by textual prompt, their results are typically over-saturated and over-smoothed, hindering their application in actual cases. Another line of optimization-based methods [71, 73, 4] turned to fuse 3D shape information, such as vertex positions, depth maps, and normal maps, with the pre-trained diffusion model by training separate modules on 3D datasets. Still, they require a specific UV layout process to achieve plausible results.

Recently, TexFusion [7] and numerous zero-shot methods [34, 45, 24] have shown significant success in generating globally consistent textures without additional 3D datasets. Based on depth-aware diffusion models, they sequentially inpaint the latents in the UV domain to ensure the spatial

consistency of latents observed across different views. Then, they decode the latents from multiple views and finally synthesize the RGB texture through back projection.

However, these methods generate static 3D assets and overlook temporal changes in visual presentations, such as videos. To our knowledge, this is the first approach to synthesize multi-view dynamic textures for mesh sequences, enabling appearance transformations.

#### **3** Preliminaries

**Video Diffusion Prior.** In this paper, we adopt CTRL-Adapter [32] as our prior model to provide dynamic information. CTRL-Adapter aims to adapt a pre-trained text-to-video diffusion model to conditions for various types of images, such as depth or normal map sequences. The key idea behind CTRL-Adapter is to leverage a pre-trained ControlNet [74] and to align its latents with those of the video diffusion model through a learnable mapping module. Intuitively, the video diffusion model generates temporally consistent video frames that capture dynamic elements like character motions and lighting, while the ControlNet further enhances this capability by allowing the model to condition on geometric information, such as depth and normal map sequences. This makes CTRL-Adapter particularly effective in providing a temporally consistent texture prior to our 4D scene texturing task. Specifically, we leverage the depth-conditioned CTRL-Adapter model. Given a sequence of depth images denoted as  $\{D_1, ..., D_K\}$  and a text prompt  $\mathcal{P}$ , CTRL-Adapter (denoted as  $\mathcal{C}$ ) synthesizes a frame sequence F by  $F = \mathcal{C}(\{D_1, ..., D_K\}, \mathcal{P})$ .

**DDIM Sampling.** DDIM [52] is a widely used sampling method in diffusion models due to its superior efficiency and deterministic nature compared to DDPM [22]. To enhance numerical stability and prevent temporal color shifts in video diffusion, numerous models [77, 21] employ a learning-based sampling technique known as v-prediction [47]. At each denoising step, the sampling process for the latents (denoted as  $z_t$ ) can be described as follows:

$$z_{t-1} = \sqrt{\alpha_{t-1}} \cdot \hat{z}_0(z_t) + \sqrt{1 - \alpha_{t-1}} \cdot \epsilon_\theta(z_t),$$
  
$$\hat{z}_0(z_t) = \frac{z_t - \sqrt{1 - \alpha_t} \cdot \epsilon_\theta}{\sqrt{\alpha_t}}, \quad \epsilon_\theta(z_t) = \epsilon_\theta,$$
  
(1)

where  $\alpha_t$  is the noise variance at time step t,  $\epsilon_{\theta}$  is the estimated noise from the U-Net denoising module, which is expected to follow  $\mathcal{N}(0, \mathcal{I})$ , and  $\hat{z}_0(z_t)$  denotes the predicted original sample (i.e., the latents at timestep 0). After the v-parameterization, the predicted original sample  $\hat{z}_0(z_t)$  and the predicted epsilon  $\epsilon_{\theta}(z_t)$  are computed as follows:

$$\hat{\boldsymbol{z}}_0(\boldsymbol{z}_t) = \sqrt{\alpha_t} \cdot \boldsymbol{z}_t - \sqrt{1 - \alpha_t} \cdot \boldsymbol{\epsilon}_{\theta}, \quad \boldsymbol{\epsilon}_{\theta}(\boldsymbol{z}_t) = \sqrt{\alpha_t} \cdot \boldsymbol{\epsilon}_{\theta} + \sqrt{1 - \alpha_t} \cdot \boldsymbol{z}_t.$$
(2)

We leverage an enhanced DDIM sampling process in video diffusion models, along with a multi-view consistent texture aggregation mechanism to synthesize 4D textures.

## 4 Method

Given an untextured mesh animation and a text prompt, our goal is to generate a multi-view and multi-frame consistent texture sequence for each mesh that aligns with both the text description and motion cues while capturing the dynamics from video diffusion models.

To optimize computational efficiency, we uniformly sample K key frames from the video and synthesize textures for these keyframes. Textures for the remaining frames are then generated by interpolating the key frame textures. Formally, given K animated meshes at the keyframes  $(\{M_1, ..., M_K\})$ , along with a text description  $\mathcal{P}$ , our method produces temporally and spatially consistent UV maps denoted as  $\{UV_1, ..., UV_K\}$ , in a zero-shot manner.

Previous texture generation methods [45, 10, 7] typically inpaint and update textures sequentially using pre-defined camera views in an incremental manner. However, these approaches rely on view-dependent depth conditions and lack global spatial consistency, often resulting in visible discontinuities in the assembled texture map. This issue arises from error accumulation during the autoregressive view update process, as noted by [4]. To resolve these issues, rather than processing each view independently, recent methods [34] propose to generate multi-view textures simultaneously through diffusion. In this work, we similarly leverage the UV space as an intermediate representation to ensure multi-view consistency.



Figure 3. **Overview.** Given a mesh sequence and a text prompt as inputs, Tex4D generates a UV-parameterized texture sequence that is both globally and temporally consistent. At each diffusion step, latent views are aggregated into UV space, followed by multi-view latent texture diffusion to ensure global consistency. To maintain temporal coherence and address self-occlusions, a Reference UV Blending module is applied at each step. Finally, the latent textures are back-projected and decoded to produce RGB textures for each frame.

#### 4.1 Overview

As shown in Fig. 3, given a sequence of K meshes, we start by rendering the mesh at V predefined, uniformly sampled camera poses to obtain multi-view depth images (denoted as  $\{D_{1,1}, ..., D_{1,K}, D_{2,1}, ..., D_{V,K}\}$ ), which serve as the geometric conditions. To generate textures for each mesh, we initialize  $V \times K$  noise images sampled from a Normal distribution (denoted as  $\{z_{b}^{1,1}, ..., z_{b}^{1,K}, z^{2,1}, ..., z^{V,K}\}$ ). Additionally, we initialize an extra noise map sequence  $\{z_{b}^{1}, ..., z_{b}^{K}\}$  for the backgrounds learning. This noise map corresponds to the texture of a plane mesh that is composited with the foreground object at each diffusion step (See Sec. 4.3). Next, for each view  $v \in \{1, ..., V\}$ , we apply the video diffusion model [32] discussed in Sec. 3 to simultaneously denoise all latents and obtain multi-frame consistent images as  $\{I^{1,v}, ..., I^{K,v}\} = C(\{D_{1,v}, ..., D_{K,v}\}, \mathcal{P})$ , where  $\mathcal{P}$  is the provided text prompt. Finally, we un-project and aggregate all denoised multi-view images for each mesh to formulate temporally consistent UV textures.

Applying the video diffusion model independently to each camera view often results in multi-view inconsistencies. Inspired by [34, 24, 75], we aggregate the multi-view latents of each mesh in the UV space to merge observations across different views at each denoising step, and then render latent from the latent texture to ensure multi-view consistency. Furthermore, we composite the rendered foreground latents with the background latents at each diffusion step (discussed in Sec. 4.2), which is essential to exploit prior in the video diffusion model (see Fig. 15). Nonetheless, such a simple aggregation method introduces blurriness in the final results. In Sec. 4.3, we analyze the underlying causes and propose a simple yet effective method to enhance the denoising process. Additionally, we create a reference UV to handle self-occlusions and further improve temporal consistency in Sec. 4.4.

#### 4.2 Multi-view Latents Aggregation in the UV Space

We describe the aggregation of multi-view latents in the UV space. For frame  $k \in \{1, ..., K\}$ , we aggregate the multi-view latents  $\{z^{1,k}, \ldots, z^{V,k}\}$  in the UV space by:

$$\mathcal{T}^{k}\left(\boldsymbol{z}^{k}\right) = \frac{\sum_{v=1}^{V} \mathcal{R}^{-1}(\boldsymbol{z}^{v,k}, c_{v}) \odot \cos\left(\theta^{v}\right)^{\alpha}}{\sum_{v=1}^{V} \cos\left(\theta^{v}\right)^{\alpha}},\tag{3}$$

where  $\mathcal{R}^{-1}$  represents the inverse rendering operator that un-projects the latents to the UV space, thus  $\mathcal{R}^{-1}(\boldsymbol{z}^{v,k}, \boldsymbol{c}_{v})$  produces a partial latent UV texture from view  $v, \cos(\theta^{v})$  is the cosine map buffered by the geometry shader, recording the cosine value between the view direction and the surface normal for each pixel,  $\alpha$  is a scaling factor, and  $\boldsymbol{c}_{v}$  denotes one of the predefined cameras. After multi-view latents aggregation, we obtain multi-view consistent latents by rendering the aggregated UV latent map using  $\tilde{\boldsymbol{z}}^{v,k} = \mathcal{R}(\mathcal{T}^{k}; \boldsymbol{c}_{v})$ , where  $\mathcal{R}$  is the rendering operation.

#### 4.3 Multi-frame Consistent Texture Generation

The aggregation process discussed above yields multi-view consistent latents  $\{\tilde{z}^{v,k}\}$  for the denoising steps. However, this simple aggregation and projection strategy leads to a blurry appearance, as shown in Fig. 8(b). This issue arises primarily because the aggregation process depicted in Eq. 3 derails the DDIM denoise process. Specifically, the estimated noise  $\epsilon_{\theta}(z_t)$  for each step in Eq. 1 is expected to follow  $\mathcal{N}(0,\mathcal{I})$ , but Eq. 3 indicates that after aggregating multi-view latents, the expected norm of variance of the noise distribution would be less than  $\mathcal{I}$ . We denote this as the "variance shift" issue caused by the texture aggregation.

To resolve this issue, we rewrite the estimated noise  $\epsilon_{\theta}$  as the combination of the t-step latent  $z_t$  and the estimated latent  $\hat{z}_0(z_t)$  at step 0. The v-parameterized predicted epsilon  $\epsilon_{\theta}(z_t)$  in Eq. 2 is can be equivalently expressed as:

$$\begin{aligned} \boldsymbol{\epsilon}_{\boldsymbol{\theta}} &= \left(\sqrt{\alpha_t} \cdot \boldsymbol{z}_t - \hat{\boldsymbol{z}}_0(\boldsymbol{z}_t)\right) / \sqrt{1 - \alpha_t} \\ \boldsymbol{\epsilon}_{\boldsymbol{\theta}}(\boldsymbol{z}_t) &= \sqrt{\alpha_t} \cdot \boldsymbol{\epsilon}_{\boldsymbol{\theta}} + \sqrt{1 - \alpha_t} \cdot \boldsymbol{z}_t \\ &= \sqrt{\frac{\alpha_t}{1 - \alpha_t}} \cdot \left(\sqrt{\alpha_t} \boldsymbol{z}_t - \hat{\boldsymbol{z}}_0(\boldsymbol{z}_t)\right) + \sqrt{1 - \alpha_t} \cdot \boldsymbol{z}_t. \end{aligned}$$
(4)

In practice, we carry out this denoising technique in the UV space. Specifically, we first compute the original texture map (i.e., texture map at step 0, denoted as  $\hat{\mathcal{T}}_0$ ) by aggregating the predicted original multi-view image latents through Eq. 3. The noisy latent texture map at time step t (denoted as  $\mathcal{T}_t$ ) can be similarly computed. We denoise in step t by:

$$\mathcal{T}_{t-1} = \sqrt{\alpha_{t-1}} \cdot \hat{\mathcal{T}}_0 + \sqrt{1 - \alpha_{t-1}} \cdot \left(\sqrt{\frac{\alpha_t}{1 - \alpha_t}} \cdot (\sqrt{\alpha_t}\mathcal{T}_t - \hat{\mathcal{T}}_0) + \sqrt{1 - \alpha_t} \cdot \mathcal{T}_t\right).$$
(5)

Through experimentation, we observe that background optimization plays a crucial role in fully exploiting the prior within the video diffusion model. As shown in Fig. 8(c), using a simple white background produces blurry results. This may stem from a mismatch between the white-background images and the training dataset, which likely contains fewer such examples, affecting the denoising process. To resolve this issue, we compute the final latents as the combination of the foreground latent  $\tilde{z}_{t-1}$  projected from the aggregated UV latents and the residual background latent  $z_{b,t-1}$  denoised by diffusion models. Specifically, we composite the estimated latents in the t - 1 step as follows:

$$\boldsymbol{z}_{t-1} = \tilde{\boldsymbol{z}}_{t-1} \odot \boldsymbol{\mathcal{M}}_{\text{fg}} + \boldsymbol{z}_{b,t-1} \odot (1 - \boldsymbol{\mathcal{M}}_{\text{fg}}), \quad \tilde{\boldsymbol{z}}_{t-1}, \boldsymbol{\mathcal{M}}_{\text{fg}} = \mathcal{R}\left(\mathcal{T}_{t-1}; \boldsymbol{c}_{v}\right), \tag{6}$$

where  $\mathcal{M}_{fg}$  represents the foreground mask of the mesh, and  $\mathcal{R}$  is the rendering operation.

To summarize, our diffusion process starts with  $K \times (V + 1)$  randomly initialized noise maps sampled (i.e.,  $\{\boldsymbol{z}_{T}^{1,k}, \ldots, \boldsymbol{z}_{T}^{V,k}\}$ , for foreground,  $\{\boldsymbol{z}_{b}^{1}, \ldots, \boldsymbol{z}_{b}^{K}\}$  for background) and denoise them into images simultaneously. At each denoising step t with the key frame k, we derive the estimated noises  $\{\boldsymbol{\epsilon}_{t-1}^{1,k}, \ldots, \boldsymbol{\epsilon}_{t-1}^{V,k}\}$  using the video diffusion model and calculate the estimated original latent  $\{\hat{\boldsymbol{z}}_{0}^{1,k}, \ldots, \hat{\boldsymbol{z}}_{0}^{V,k}\}$  by Eq. 1. Then, we use Eq. 3 to aggregate the latents onto UV space. Next, we utilize Eq. 5 to take the diffusion step in the UV space, and render the synchronized latents  $\{\tilde{\boldsymbol{z}}_{t-1}^{1,k}, \ldots, \tilde{\boldsymbol{z}}_{t-1}^{V,k}\}$  from latent UVs  $\{\mathcal{T}_{t-1}^{1}, \ldots, \mathcal{T}_{t-1}^{K}\}$  to ensure multi-view consistency. Finally, we composite the denoised latent with the latents at step t - 1 according to foreground masks by Eq. 6.

#### 4.4 Reference UV Blending

While the video diffusion model ensures temporal consistency for latents from each view, consistency can sometimes diminish after aggregation in the texture domain. This issue primarily stems from the view-dependent nature of the depth conditions and the limited resolution of latents, which can

a sketch of bot in a sandy beach, Van-Gogh style." "a cute alien in neon lights, cyberpunk style."



Figure 4. **Qualitative Results.** Our method generates multi-view consistent dynamic textures with a diverse set of styles and prompts. Zoom in to view the details. More results are provided in the supplementary material.



Figure 5. **Qualitative Comparison with Text-to-4D Methods.** Our methods generates multi-view consistent compared with Text-to-4D methods as our approach fully utilizes the geometry information of the meshes. lead to distortions when features from different camera angles are combined onto the UV texture. Additionally, self-occlusion during animation often results in information loss in invisible regions.

To address these challenges, we propose a reference UV map to enhance correlations between latent textures across frames. Specifically, the reference UV map is constructed by sequentially combining latent textures over time, with each new texture filling only the empty texels of the reference UV map. Each texture is blended using the reference UV  $T_{UV}$  with a mask  $M_{UV}$  that labels the visible region:

$$\mathcal{T}_{t}^{k} = \left( (1-\lambda) \cdot \mathcal{T}_{t}^{k} + \lambda \cdot \mathcal{T}_{\mathcal{U}\mathcal{V}} \right) \odot \mathcal{M}_{\mathcal{U}\mathcal{V}}^{k} + \mathcal{T}_{\mathcal{U}\mathcal{V}} \odot \left( 1 - \mathcal{M}_{\mathcal{U}\mathcal{V}}^{k} \right)$$
(7)

where  $\lambda$  is the blending weight for the reference UV in the visible region, while the invisible region is simply replaced with the reference texture. We empirically set the blending weight to 0.2.

#### **5** Experiments

**Datasets.** We source our datasets from two primary repositories: human motion diffusion outputs and the Mixamo and Sketchfab websites. We employ the text-to-motion diffusion model [55] to compare our approach with LatentMan [15]. For comparison with Generative Rendering [6], we obtain animated characters from the Mixamo. Specifically, we first use Blender [14] to extract meshes, joints, skinning weights, and animation data from the FBX files. Then, we apply linear blend skinning to animate the meshes. We utilize XATLAS to parameterize the mesh and unwrap the UVs.

**Baselines.** To our knowledge, no existing studies tackle dynamic texture generation. We adopt eight recent methods, rendering the input based on their configurations to establish baselines, including Text-to-4D methods, video stylization methods and video generation methods with various control mechanisms. SV4D [65] and L4GM [44] are text-to-4D methods, both taking a front view as the input. PnP-Diffusion [56] is an image stylization method that guides the generation with DDIM features. We extend the method on a frame-by-frame basis for comparison, aligning with previous work [17]. Built upon cross-frame attention, Text2Video-Zero [25] guides the video by warping latents to enhance video dynamics implicitly. We leverage its official extension for comparison. TokenFlow [17], Generative Rendering [6], and LatentMan [15] establish latent correspondences through nearest neighbor and DensePose features. Gen-1 [16] is a video-to-video model that transforms the untextured mesh renders into stylized outputs. Given the lack of the source code for Generative Rendering, we utilize the experimental results presented in their video demos for qualitative comparison. Additionally, we compare our method with the texture generation method Text2Tex [10].

**Evaluation Metrics.** Quantitatively evaluating multi-view consistency and temporal coherence remains challenging. We conduct a user study to assess overall performance, including appearance quality, spatio-temporal consistency, and prompt fidelity based on human preferences. Additionally, we measure videl-level multi-view temporal coherence using Fréchet Video Distance (FVD) [57] following [28, 65], along with a CLIP-based Consistency Score following [34].

Table 1. **Quantitative evaluation**. We present metric values and a comparison of the percentage and statistics of user preference for our approach against other methods. Our method shows the best spatio-temporal consistency as measured by the FVD [57] and Consistency Score [34]. Users consistently favored Tex4D over all baselines.



"a Stormtrooper is Swimming"

Figure 6. **Qualitative comparisons of multi-view video generation.** We compare our method against PnPdiffusion [56], TokenFlow [17], Text2Video-Zero [25], Generative Rendering [6] (from their video demo), and Gen-1 [16]. We generate videos in the front view and the side view (blue box) on Mixamo dataset. Our method generates vivid videos that align with the textual prompts while preserving spatial consistency.

#### 5.1 Qualitative Results

We present qualitative evaluation in Fig. 5 and Fig. 6. Text-to-4D methods fail to render plausible results as decoupling geometry and appearance simultaneously is more challenging. Generative Rendering, TokenFlow, and Text2Video-Zero, which rely on T2I diffusion models with cross-frame attention mechanisms, exhibit noticeable flickering compared to other methods. This issue stems partly from the empirical and implicit correspondence mapping used to enforce interframe latent consistency, as the correspondences in the latent space



Figure 7. **Comparison with Text2Tex.** Texture generation method [10] shows empty texels in invisible regions and fails to model dynamics across time.

may not precisely align with those in the RGB space. In contrast, our approach interpolates the frames between key frame textures in RGB space, eliminating artifacts caused by latent manipulation. PnP-Diffusion edits frames independently and generates detailed and sophisticated appearances but suffers from poor spatio-temporal consistency due to the loss of temporal correlations in the latent space. While Gen-1 produces high-quality videos, it fails to maintain multi-view consistency.

Furthermore, we present multi-view results showcasing a variety of styles and prompts in Fig. 4. Our method, driven by video diffusion models, effectively accounts for the styles and captures temporal variations over time. As shown in Fig. 7, Tex4D effectively handles the invisible regions compared with the traditional texture generation method Text2Tex [10], which also fails to model dynamics.

#### 5.2 Quantitative Evaluation

To quantitatively assess the effectiveness of our proposed method, we follow prior research [15, 17, 16] and conduct a comprehensive A/B user study. Our study involved 67 participants who provided a total of 1104 valid responses based on six different scenes drawn from six previous works, with each scene producing videos from two different views. During each evaluation, participants were



Figure 8. Comparisons of multi-view denoising algorithms and ablation on background priors. (a) The simple multi-view diffusion algorithms by Eq. 1, and (b) aggregation of  $z_{t-1}$  result in a blurry appearance compared to (d) our results. Without background priors, our approach fails to generate plausible video clips.

presented with rendered meshes and depth conditions viewed from two angles, serving as motion references. They were shown a pair of videos: one generated by our approach and the other from a baseline method. Participants were asked to select the method that exhibited superior performance in three criteria: 1) appearance quality, 2) spatial and temporal consistency, and 3) fidelity to the prompts. Table 1 summarizes the preference percentage and statistics (ranging from 1 to 5) of our method over other methods. Our method significantly surpasses other methods by a large margin. In addition, our method achieves the best FVD and Consistency Score that demonstrates better multi-view consistency in generated video clips.

#### 5.3 Ablation Study

Ablation for texture aggregation. In Fig. 8 (a) and (b), we present two alternative texture aggregation methods. In Fig. 8 (a), we un-project  $\hat{z}_0(z_t)$  and  $\epsilon_{\theta}(z_t)$  into UV space for aggregation. In Fig. 8 (b), we map  $z_{t-1}$  to the UV space. Both these two approaches show inferior results compared to our method, which verifies the effectiveness of the proposed texture aggregation algorithm.



Ablation for UV blending module. In Sec. 4.4, we propose a reference UV blending schema to resolve the temporal inconsistency caused by latent aggregation. To validate the effectiveness of this mechanism (See Sec. 4.4), we conduct an ablation study by disable

Figure 9. Ablation study on the reference UV blending module. Without this module, the generated textures lose consistency over time.

(See Sec. 4.4), we conduct an ablation study by disabling the reference UV blending module (setting  $\lambda$  to 0). As shown in Fig. 9, without the UV blending module, our method generates textures with noticeable distortions on the Joker's face over time.

**Ablation for background priors.** Sec. 4.3 discusses the importance of including a plausible background prior. To verify the effectiveness of this design, we replace the learnable background latents with an all-white background while keeping all other parts unchanged. Fig. 8 (c) illustrates that this ablation experiment produces significantly blurrier textures compared to our full method, highlighting the importance of background learning.

## 6 Conclusions

In this paper, we present Tex4D, a zero-shot approach that generates multi-view, multi-frame consistent dynamic textures for untextured, animated mesh sequences based on a text prompt. By incorporating texture aggregation in the UV space within the diffusion process of a conditional video diffusion model, we ensure both temporal and spatial coherence in the generated textures. To leverage priors from existing video diffusion models, we develop an effective modification to the DDIM sampling process to address the variance shift issue caused by multi-view texture aggregation and design a background learning module. Additionally, we enhance temporal consistency by introducing a reference UV map and developing a dynamic background learning framework to produce fully textured 4D scenes. Extensive experiments show that our method can synthesize realistic and consistent 4D textures, demonstrating its superiority against state-of-the-art baselines.

#### References

- [1] J. Ansel, E. Yang, H. He, N. Gimelshein, A. Jain, M. Voznesensky, B. Bao, P. Bell, D. Berard, E. Burovski, G. Chauhan, A. Chourdia, W. Constable, A. Desmaison, Z. DeVito, E. Ellison, W. Feng, J. Gong, M. Gschwind, B. Hirsh, S. Huang, K. Kalambarkar, L. Kirsch, M. Lazos, M. Lezcano, Y. Liang, J. Liang, Y. Lu, C. Luk, B. Maher, Y. Pan, C. Puhrsch, M. Reso, M. Saroufim, M. Y. Siraichi, H. Suk, M. Suo, P. Tillet, E. Wang, X. Wang, W. Wen, S. Zhang, X. Zhao, K. Zhou, R. Zou, A. Mathews, G. Chanan, P. Wu, and S. Chintala. PyTorch 2: Faster Machine Learning Through Dynamic Python Bytecode Transformation and Graph Compilation. In 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '24). ACM, Apr. 2024.
- [2] M. Bain, A. Nagrani, G. Varol, and A. Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021.
- [3] O. Bar-Tal, D. Ofri-Amar, R. Fridman, Y. Kasten, and T. Dekel. Text2live: Text-driven layered image and video editing. In ECCV, pages 707–723, 2022.
- [4] R. Bensadoun, Y. Kleiman, I. Azuri, O. Harosh, A. Vedaldi, N. Neverova, and O. Gafni. Meta 3d texturegen: Fast and consistent texture generation for 3d objects. arXiv preprint arXiv:2407.02430, 2024.
- [5] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S. W. Kim, S. Fidler, and K. Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023.
- [6] S. Cai, D. Ceylan, M. Gadelha, C.-H. Huang, T. Wang, and G. Wetzstein. Generative rendering: Controllable 4d-guided video generation with 2d diffusion models. In CVPR, 2024.
- [7] T. Cao, K. Kreis, S. Fidler, N. Sharp, and K. Yin. Texfusion: Synthesizing 3d textures with text-guided image diffusion models. In *ICCV*, 2023.
- [8] D. Ceylan, C.-H. Huang, and N. J. Mitra. Pix2video: Video editing using image diffusion. In ICCV, 2023.
- [9] D. Z. Chen, H. Li, H.-Y. Lee, S. Tulyakov, and M. Nießner. Scenetex: High-quality texture synthesis for indoor scenes via diffusion priors. arXiv preprint arXiv:2311.17261, 2023.
- [10] D. Z. Chen, Y. Siddiqui, H.-Y. Lee, S. Tulyakov, and M. Nießner. Text2tex: Text-driven texture synthesis via diffusion models. arXiv preprint arXiv:2303.11396, 2023.
- [11] H. Chen, M. Xia, Y. He, Y. Zhang, X. Cun, S. Yang, J. Xing, Y. Liu, Q. Chen, X. Wang, C. Weng, and Y. Shan. Videocrafter1: Open diffusion models for high-quality video generation. arXiv preprint arXiv:2310.19512, 2023.
- [12] H. Chen, Y. Zhang, X. Cun, M. Xia, X. Wang, C. Weng, and Y. Shan. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. arXiv preprint arXiv:2401.09047, 2024.
- [13] Y. Chen, R. Chen, J. Lei, Y. Zhang, and K. Jia. Tango: Text-driven photorealistic and robust 3d stylization via lighting decomposition. *NeurIPS*, 35:30923–30936, 2022.
- [14] B. O. Community. *Blender a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2024.
- [15] A. Eldesokey and P. Wonka. Latentman: Generating consistent animated characters using image diffusion models. In CVPR, pages 7510–7519, 2024.
- [16] P. Esser, J. Chiu, P. Atighehchian, J. Granskog, and A. Germanidis. Structure and content-guided video synthesis with diffusion models. In *ICCV*, pages 7346–7356, 2023.
- [17] M. Geyer, O. Bar-Tal, S. Bagon, and T. Dekel. Tokenflow: Consistent diffusion features for consistent video editing. arXiv preprint arXiv:2307.10373, 2023.
- [18] Y. Guo, C. Yang, A. Rao, M. Agrawala, D. Lin, and B. Dai. Sparsectrl: Adding sparse controls to text-to-video diffusion models. arXiv preprint arXiv:2311.16933, 2023.
- [19] Y. Guo, C. Yang, A. Rao, Z. Liang, Y. Wang, Y. Qiao, M. Agrawala, D. Lin, and B. Dai. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. arXiv preprint arxiv:2307.04725, 2023.
- [20] Y. He, T. Yang, Y. Zhang, Y. Shan, and Q. Chen. Latent video diffusion models for high-fidelity long video generation. arXiv preprint arXiv:2211.13221, 2022.

- [21] J. Ho, W. Chan, C. Saharia, J. Whang, R. Gao, A. Gritsenko, D. P. Kingma, B. Poole, M. Norouzi, D. J. Fleet, et al. Imagen video: High definition video generation with diffusion models. arXiv preprint arXiv:2210.02303, 2022.
- [22] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33:6840–6851, 2020.
- [23] W. Hong, M. Ding, W. Zheng, X. Liu, and J. Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. arXiv preprint arXiv:2205.15868, 2022.
- [24] D. Huo, Z. Guo, X. Zuo, Z. Shi, J. Lu, P. Dai, S. Xu, L. Cheng, and Y.-H. Yang. Texgen: Text-guided 3d texture generation with multi-view sampling and resampling. *ECCV*, 2024.
- [25] L. Khachatryan, A. Movsisyan, V. Tadevosyan, R. Henschel, Z. Wang, S. Navasardyan, and H. Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. arXiv preprint arXiv:2303.13439, 2023.
- [26] N. M. Khalid, T. Xie, E. Belilovsky, and P. Tiberiu. Clip-mesh: Generating textured meshes from text using pretrained image-text models. *SIGGRAPH Aisa*, December 2022.
- [27] J.-g. Kwak, E. Dong, Y. Jin, H. Ko, S. Mahajan, and K. M. Yi. Vivid-1-to-3: Novel view synthesis with video diffusion models. arXiv preprint arXiv:2312.01305, 2023.
- [28] B. Li, C. Zheng, W. Zhu, J. Mai, B. Zhang, P. Wonka, and B. Ghanem. Vivid-zoo: Multi-view video generation with diffusion model. arXiv preprint arXiv:2406.08659, 2024.
- [29] X. Li, W. Chu, Y. Wu, W. Yuan, F. Liu, Q. Zhang, F. Li, H. Feng, E. Ding, and J. Wang. Videogen: A reference-guided latent diffusion approach for high definition text-to-video generation. arXiv preprint arXiv:2309.00398, 2023.
- [30] R. Liang, Z. Gojcic, H. Ling, J. Munkberg, J. Hasselgren, Z.-H. Lin, J. Gao, A. Keller, N. Vijaykumar, S. Fidler, et al. Diffusionrenderer: Neural inverse and forward rendering with video diffusion models. *arXiv preprint arXiv:2501.18590*, 2025.
- [31] 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 CVPR, 2023.
- [32] H. Lin, J. Cho, A. Zala, and M. Bansal. Ctrl-adapter: An efficient and versatile framework for adapting diverse controls to any diffusion model. arXiv preprint arXiv:2404.09967, 2024.
- [33] 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.
- [34] Y. Liu, M. Xie, H. Liu, and T.-T. Wong. Text-guided texturing by synchronized multi-view diffusion. arXiv preprint arXiv:2311.12891, 2023.
- [35] 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 preprint arXiv:2310.15008, 2023.
- [36] Y. Lu, Y. Liang, L. Zhu, and Y. Yang. Freelong: Training-free long video generation with spectralblend temporal attention. arXiv preprint arXiv:2407.19918, 2024.
- [37] G. Metzer, E. Richardson, O. Patashnik, R. Giryes, and D. Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. arXiv preprint arXiv:2211.07600, 2022.
- [38] O. Michel, R. Bar-On, R. Liu, S. Benaim, and R. Hanocka. Text2mesh: Text-driven neural stylization for meshes. arXiv preprint arXiv:2112.03221, 2021.
- [39] B. Peng, J. Wang, Y. Zhang, W. Li, M.-C. Yang, and J. Jia. Controlnext: Powerful and efficient control for image and video generation. arXiv preprint arXiv:2408.06070, 2024.
- [40] R. Po and G. Wetzstein. Compositional 3d scene generation using locally conditioned diffusion. In 2024 International Conference on 3D Vision (3DV), pages 651–663. IEEE, 2024.
- [41] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. arXiv, 2022.
- [42] C. Qi, X. Cun, Y. Zhang, C. Lei, X. Wang, Y. Shan, and Q. Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. arXiv preprint arXiv:2303.09535, 2023.

- [43] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021.
- [44] J. Ren, C. Xie, A. Mirzaei, K. Kreis, Z. Liu, A. Torralba, S. Fidler, S. W. Kim, H. Ling, et al. L4gm: Large 4d gaussian reconstruction model. *Advances in Neural Information Processing Systems*, 37:56828–56858, 2024.
- [45] E. Richardson, G. Metzer, Y. Alaluf, R. Giryes, and D. Cohen-Or. Texture: Text-guided texturing of 3d shapes. In SIGGRAPH, pages 1–11, 2023.
- [46] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. arXiv preprint arXiv:2112.10752, 2021.
- [47] T. Salimans and J. Ho. Progressive distillation for fast sampling of diffusion models. arXiv preprint arXiv:2202.00512, 2022.
- [48] C. Schuhmann, R. Vencu, R. Beaumont, R. Kaczmarczyk, C. Mullis, A. Katta, T. Coombes, J. Jitsev, and A. Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. arXiv preprint arXiv:2111.02114, 2021.
- [49] 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.
- [50] Y. Shi, P. Wang, J. Ye, L. Mai, K. Li, and X. Yang. Mvdream: Multi-view diffusion for 3d generation. arXiv preprint arXiv:2308.16512, 2023.
- [51] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, et al. Make-a-video: Text-to-video generation without text-video data. arXiv preprint arXiv:2209.14792, 2022.
- [52] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502, 2020.
- [53] S. Tang, J. Chen, D. Wang, C. Tang, F. Zhang, Y. Fan, V. Chandra, Y. Furukawa, and R. Ranjan. Mvdiffusion++: A dense high-resolution multi-view diffusion model for single or sparse-view 3d object reconstruction. arXiv preprint arXiv:2402.12712, 2024.
- [54] S. Tang, F. Zhang, J. Chen, P. Wang, and Y. Furukawa. Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion. arXiv, 2023.
- [55] G. Tevet, S. Raab, B. Gordon, Y. Shafir, D. Cohen-or, and A. H. Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023.
- [56] N. Tumanyan, M. Geyer, S. Bagon, and T. Dekel. Plug-and-play diffusion features for text-driven imageto-image translation. In CVPR, pages 1921–1930, June 2023.
- [57] T. Unterthiner, S. Van Steenkiste, K. Kurach, R. Marinier, M. Michalski, and S. Gelly. Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717, 2018.
- [58] V. Voleti, C.-H. Yao, M. Boss, A. Letts, D. Pankratz, D. Tochilkin, C. Laforte, R. Rombach, and V. Jampani. SV3D: Novel multi-view synthesis and 3D generation from a single image using latent video diffusion. In ECCV, 2024.
- [59] P. von Platen, S. Patil, A. Lozhkov, P. Cuenca, N. Lambert, K. Rasul, M. Davaadorj, and T. Wolf. Diffusers: State-of-the-art diffusion models. https://github.com/huggingface/diffusers, 2022.
- [60] Z. Wang, C. Lu, Y. Wang, F. Bao, C. Li, H. Su, and J. Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023.
- [61] H. Weng, T. Yang, J. Wang, Y. Li, T. Zhang, C. Chen, and L. Zhang. Consistent123: Improve consistency for one image to 3d object synthesis. arXiv preprint arXiv:2310.08092, 2023.
- [62] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics.
- [63] J. Z. Wu, Y. Ge, X. Wang, S. W. Lei, Y. Gu, Y. Shi, W. Hsu, Y. Shan, X. Qie, and M. Z. Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, pages 7623–7633, 2023.

- [64] R. Wu, R. Gao, B. Poole, A. Trevithick, C. Zheng, J. T. Barron, and A. Holynski. Cat4d: Create anything in 4d with multi-view video diffusion models. arXiv preprint arXiv:2411.18613, 2024.
- [65] Y. Xie, C.-H. Yao, V. Voleti, H. Jiang, and V. Jampani. SV4D: Dynamic 3d content generation with multi-frame and multi-view consistency. arXiv preprint arXiv:2407.17470, 2024.
- [66] J. Xing, M. Xia, Y. Zhang, H. Chen, X. Wang, T.-T. Wong, and Y. Shan. Dynamicrafter: Animating open-domain images with video diffusion priors. arXiv preprint arXiv:2310.12190, 2023.
- [67] Z. Xu, J. Zhang, J. H. Liew, H. Yan, J.-W. Liu, C. Zhang, J. Feng, and M. Z. Shou. Magicanimate: Temporally consistent human image animation using diffusion model. In CVPR, 2024.
- [68] S. Yang, Y. Zhou, Z. Liu, , and C. C. Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Aisa*, 2023.
- [69] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024.
- [70] S. Yu, K. Sohn, S. Kim, and J. Shin. Video probabilistic diffusion models in projected latent space. In CVPR, pages 18456–18466, 2023.
- [71] X. Yu, P. Dai, W. Li, L. Ma, Z. Liu, and X. Qi. Texture generation on 3d meshes with point-uv diffusion. In *ICCV*, pages 4206–4216, 2023.
- [72] Y. Yuan, Y. Guo, C. Wang, W. Zhang, H. Xu, and L. Zhang. Freqprior: Improving video diffusion models with frequency filtering gaussian noise. arXiv preprint arXiv:2502.03496, 2025.
- [73] X. Zeng, X. Chen, Z. Qi, W. Liu, Z. Zhao, Z. Wang, B. Fu, Y. Liu, and G. Yu. Paint3d: Paint anything 3d with lighting-less texture diffusion models. In CVPR, pages 4252–4262, 2024.
- [74] L. Zhang, A. Rao, and M. Agrawala. Adding conditional control to text-to-image diffusion models. In CVPR, pages 3836–3847, 2023.
- [75] L. Zhang, Z. Wang, Q. Zhang, Q. Qiu, A. Pang, H. Jiang, W. Yang, L. Xu, and J. Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. ACM Transactions on Graphics (TOG), 2024.
- [76] S. Zhang, J. Wang, Y. Zhang, K. Zhao, H. Yuan, Z. Qin, X. Wang, D. Zhao, and J. Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. arXiv preprint arXiv:2311.04145, 2023.
- [77] S. Zhang, J. Wang, Y. Zhang, K. Zhao, H. Yuan, Z. Qing, X. Wang, D. Zhao, and J. Zhou. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. arXiv preprint arXiv:2311.04145, 2023.
- [78] Y. Zhang, Y. Wei, D. Jiang, X. Zhang, W. Zuo, and Q. Tian. Controlvideo: Training-free controllable text-to-video generation. arXiv preprint arXiv:2305.13077, 2023.
- [79] Y. Zhao, C.-C. Lin, K. Lin, Z. Yan, L. Li, Z. Yang, J. Wang, G. H. Lee, and L. Wang. Genxd: Generating any 3d and 4d scenes. *arXiv preprint arXiv:2411.02319*, 2024.
- [80] W. Zheng, J. Teng, Z. Yang, W. Wang, J. Chen, X. Gu, Y. Dong, M. Ding, and J. Tang. Cogview3: Finer and faster text-to-image generation via relay diffusion. arXiv preprint arXiv:2403.05121, 2024.
- [81] D. Zhou, W. Wang, H. Yan, W. Lv, Y. Zhu, and J. Feng. Magicvideo: Efficient video generation with latent diffusion models. arXiv preprint arXiv:2211.11018, 2022.

# Appendix

A	More Implementation Details							
	A.1	Implementation Details	15					
	A.2	Denoising Algorithm of Our Method	15					
	A.3	V-Prediction	15					
B	More Qualitative Results							
	<b>B</b> .1	Graphics Application and Video Demo	16					
	<b>B</b> .2	Multi-view Results	16					
	<b>B.3</b>	Texture Results	19					
С	Mor	e Ablation Results	19					
D	More Method Comparisons							
	<b>D</b> .1	Comparison with Depth-Conditioned Video Diffusion Models	20					
	D.2	Comparison with Textured Mesh Animations	21					
	D.3	Comparison with Text-to-4D Methods	22					
E	User	User Study						
F	Broader Discussion							
	F.1	Robustness of Different UV Mappings	24					
	F.2	Highly Structured Texture Generation	24					
	F.3	Runtime Breakdown	25					
	F.4	Long Texture Sequence Generation	25					
	F.5	Analysis of Appearance Details	25					
	F.6	Societal Impacts	26					
G	Limitations							
	<b>G</b> .1	Panoramic Background Modeling	26					
	G.2	Computation Time	26					
н	Lice	Licenses 20						

## **A** More Implementation Details

#### A.1 Implementation Details

We utilize the CTRL-Adapter [32], trained on the video diffusion model I2VGen-XL [77], as the backbone for generation, with the denoising steps set to T = 50. Initially, we center the untextured mesh sequence and pre-define six different viewpoints around the Y-axis in the XZ-plane, uniformly sampled in spherical coordinates, along with an additional top view with an elevation angle of zero and an azimuth angle of  $30^{\circ}$ . For latent initialization, we first sample Gaussian noise on the latent textures and then render 2D latent samples for each view to improve the coherence of the generated outputs. During denoising, we upscale the latent resolution to  $96 \times 96$  to reduce aliasing. We empirically set the blending coefficient to 0.2 during our experiment. It takes approximately 30 minutes to generate a video with 24 keyframes taken on an RTX A6000 Ada GPU. We decode the denoised latents in keyframes to RGB images, and then un-project and aggregate these images to transform the latent UV maps to RGB textures as previous works [34, 7, 24]. The resolutions of latent texture and RGB texture are  $1024 \times 1024$ ,  $1536 \times 1536$  respectively. Finally, we interpolate the textures of the keyframes at intervals of 3 to synthesize the final video clips.

**Background Optimization** For each frame, we use a latent texture as the background shared across multiple views. The first frame is initialized from the image provided by the user as the CTRL-Adapter input. Then, we jointly optimize the background and foreground as described in Sec. 4.3. In our early attempts, we also used a cube or a hemicube mesh to model a free-view background, but found that the backgrounds were not accurate within uncaptured regions by cameras, such as the bottom. Note that our main focus is not a solution for free-view background generation, which is an open research problem explored by concurrent works like Cat4D [64] and GenXD [79].

**Texture Interpolation Details** We compute RGB textures for non-key frames by interpolating RGB textures of nearby keyframes. Specifically, when the denoising stage finishes, we have RGB textures for keyframes, i.e.,  $T_{rgb,t} = \frac{t-k \cdot n}{n} \cdot T_{rgb,k} + \frac{(k+1) \cdot n-t}{n} \cdot T_{rgb,k+1}$ , where t is the t-th frame in the final video, k is the nearest prior keyframe, and n is the keyframe interval.

#### A.2 Denoising Algorithm of Our Method

We present the complete workflow of our method in Algorithm 1. For clarity, we omit the notation for the latent variables  $z_b$  representing the background plane texture, as they follow the same scheme as the foreground latents. The reference UV map  $T_{UV}$  is constructed by progressively combining latent textures over time, with each new texture filling only the unoccupied texels in the reference UV map. We denote this process as "Combine" in the following workflow.

#### A.3 V-Prediction

Tex4D is a zero-shot approach built on a pre-trained conditional video diffusion model, where v-prediction is a technique commonly used in video diffusion models (e.g., I2VGen-XL [77], Imagen [21], CogVideoX [23, 69], CogView3 [80]) to accelerate the training and prevent temporal color shifts. In our method, we utilize CTRL-Adapter [32], a conditional video diffusion model that guides video by depth maps trained on the DDIM v-prediction mechanism. Hence, we use v-prediction to ensure the proper functioning of the conditional video diffusion model.

## **B** More Qualitative Results

In this section, we provide more qualitative results of Tex4D to deliver a comprehensive application and analysis of our method. First, we demonstrate the novel application of our method in collaboration with technical artists as in Sec. B.1. Our synthesized dynamic texture sequence can be seamlessly integrated into Blender's shader editor and we provide the rendered video demo in the supplementary material. Then, we analyze the multi-view rendering results of Tex4D with a diverse set of prompts and styles in Sec. B.2. Finally, we visualize the generated texture sequence from three different cases to demonstrate that our method can directly bake lighting variations and appearances into textures without the need for post-processing by technical artists in creating vivid videos.

#### Algorithm 1 Tex4D

**Input:** UV maps  $UV = \{UV_1, ..., UV_K\}$ ; depth maps  $\mathcal{D} = \{D_{1,1}, ..., D_{1,V}, D_{2,1}, ..., D_{K,V}\}$ ; text prompt  $\mathcal{P}$ ; CTRL-Adapter model  $\mathcal{C}$ ; rendering operation  $\mathcal{R}$ ; unproject operation  $\mathcal{R}^{-1}$ ; cameras c; T diffusion steps;  $\mathcal{T}$  latent textures (including foreground and background);  $\lambda$  blending weight; k keyframes

$$\begin{split} &\mathcal{T}_{T} \sim \mathcal{N}(\mathbf{0},\mathcal{I}) & \text{// Sample noise in UV space} \\ &\tilde{z}_{T}, \mathcal{M}_{\mathrm{fg}} = \mathcal{R}(\mathcal{T}_{T}; \mathbf{c}) \\ &z_{b,T} \sim \mathcal{N}(\mathbf{0},\mathcal{I}) & \text{// Composite latents} \\ &z = z_{T} = \tilde{z}_{T} \odot \mathcal{M}_{\mathrm{fg}} + z_{b,T} \odot (1 - \mathcal{M}_{\mathrm{fg}}) & \text{// Composite latents} \\ & \text{For } t = T, \dots, 1 \text{ do} & \text{// Composite latents} \\ &\tilde{z}_{b,t-1} \leftarrow \mathcal{C}(z_{b,t}; \mathcal{D}, \mathcal{P}) & \text{// Estimate noise from } \mathcal{C} \\ &\tilde{z}_{0}(z_{t}) = \sqrt{\alpha_{t}} \cdot z_{t} - \sqrt{1 - \alpha_{t}} \cdot \epsilon_{\theta} & \text{// Estimate noise from } \mathcal{C} \\ &\tilde{z}_{0}(z_{t}) = \sqrt{\alpha_{t}} \cdot z_{t} - \sqrt{1 - \alpha_{t}} \cdot \epsilon_{\theta} & \text{// Bake textures by Eq. 3} \\ & \mathcal{T}_{\mathcal{U}\mathcal{V}} = \text{Combine}(\hat{\mathcal{T}}_{0}; \mathcal{M}_{\mathcal{U}\mathcal{V}}) & \text{// Bake textures by Eq. 3} \\ & \mathcal{T}_{\mathcal{U}\mathcal{V}} = \text{Combine}(\hat{\mathcal{T}}_{0}; \mathcal{M}_{\mathcal{U}\mathcal{V}}) & \text{For } k \text{ in } 1, \dots, K \text{ do} \\ & \mathcal{T}_{t-1}^{k} = \sqrt{\alpha_{t-1}} \cdot \hat{\mathcal{T}}_{0}^{k} + \sqrt{1 - \alpha_{t-1}} \left( \sqrt{\frac{\alpha_{t}}{1 - \alpha_{t}}} \cdot \left( \sqrt{\alpha_{t}} \mathcal{T}_{t}^{k} - \hat{\mathcal{T}}_{0}^{k} \right) + \sqrt{1 - \alpha_{t}} \cdot \mathcal{T}_{t}^{k} \right) // \text{Denoise Eq. 5} \\ & \mathcal{T}_{t-1}^{k} = \left( (1 - \lambda) \cdot \mathcal{T}_{t-1}^{k} + \lambda \cdot \mathcal{T}_{\mathcal{U}\mathcal{V}} \right) \odot \mathcal{M}_{\mathcal{U}\mathcal{V}}^{k} + \mathcal{T}_{\mathcal{U}\mathcal{V}} \odot \left( 1 - \mathcal{M}_{\mathcal{U}\mathcal{V}}^{k} \right) & // \text{Blend textures by Eq. 7} \\ & \tilde{z}_{t-1}, \mathcal{M}_{\text{fg}} = \mathcal{R} \left( \mathcal{T}_{t-1}; c, \mathcal{U}\mathcal{V} \right) \\ & z_{t-1} = \tilde{z}_{t-1} \odot \mathcal{M}_{\text{fg}} + z_{b,t-1} \odot \left( 1 - \mathcal{M}_{\text{fg}} \right) & // \text{Composite latents by Eq. 6} \\ & z = z_{t-1} \\ \end{array}$$



Figure 10. **Tex4D Applications.** Our synthesized dynamic textures can be easily integrated into graphics pipelines. We utilize the shader editor in Blender to animate textures with image sequence nodes. The dynamic textures help technical artists render vivid videos without additional lighting and mesh controls. We provide a video demo to further demonstrate the effectiveness of Tex4D.

#### **B.1** Graphics Application and Video Demo

As shown in Fig. 10, Tex4D demonstrates its utility in the graphics pipeline by integrating dynamic texture sequences into Blender for rendering. This integration enables seamless visualization of animated textures directly on 3D models, showcasing Tex4D's capability to handle complex visual dynamics in real-world applications. We highly recommend the reviewers watch our supplementary videos for details.

#### **B.2** Multi-view Results

In this section, we discuss the multi-view results of our methods and other methods. As shown in Fig. 13, our method generates vivid videos that align with the textual prompts while preserving spatial consistency. PnP-Diffusion [56], TokenFlow [17] and Text2Video-Zero [25] generate videos that are not aligned with the text prompt, due to the implicit correspondence used in the multi-frame



"the Joker dances, comic style"



"the terminator dancing in the milky way



"a rusty robot dances in ruins"



"a sketch of bot dancing in a sandy beach, Van-Gogh style."

Figure 11. **More qualitative results.** We present the results of Tex4D with brief prompts, demonstrating the ability of Tex4D to generate multi-view consistent textures.



"a ghost flashed a magical light, causing dramatic shifts in lighting."



"a dingy, magic King Boo, flashing a weird light, static background."



"a sprite of fiery plums tilts its head, in full color."



"a spirit in neon tilts its head, cyberpunk style."

Figure 12. **More qualitative results on non-human character animations.** We present the results of Tex4D with prompts emphasizing the dynamics, demonstrating the ability of Tex4D to capture the dynamics from video diffusion models.



"a cyberpunk walks"

Figure 13. **Qualitative comparisons of multi-view video generation.** We compare our method against PnPdiffusion [56], TokenFlow [17], Text2Video-Zero [25], Generative Rendering [15], and Gen-1 [16]. We generate videos in the front view and the side view (blue box) on Human Diffusion Models. Our method generates vivid videos that align with the textual prompts while preserving spatial consistency.

attention. LatentMan [15] and Gen-1 [16] generate vivid videos, but the multi-view consistency of the characters is not well-preserved.

In Fig. 11, we present additional characters generated by Tex4D, showcasing the method's effectiveness and its ability to generalize across a diverse array of styles and prompts. We also evaluate Tex4D on non-human character animations in Fig. 12, demonstrating its robust generalization capabilities across various types of mesh sequences. In each case, we provide two different views to show that our method can ensure multi-view consistency.

To emphasize the temporal changes in the generated textures, we also design some prompts, for example, 'flashed a magical light', 'dramatic shifts in lighting', 'cyberpunk style' in our experiments. As shown in Fig. 12, the results of 'ghost', 'King Boo' and 'Snowman' validate the effectiveness of our method in generating different level of temporal changes by a variety of textual prompts, while maintaining the consistency both spatially and temporally. Additionally, we provide a supplementary video that includes baseline comparisons and multi-view results for all examples.

#### **B.3** Texture Results

In this section, we present the texture sequences, which are the intermediate results of our pipeline. Our method utilizes XATLAS to unwrap the UV maps from meshes without human labor. XATLAS is a widely used library for mesh parameterization and UV unwrapping, commonly integrated into popular tools and engines, facilitating efficient texture mapping in 3D graphics applications. As shown in Fig. 14, our method seamlessly bakes temporal changes, including lighting variations, wrinkles, and appearance transformations, directly into the textures, removing the need for manual post-production by technical artists.

## C More Ablation Results

Ablation on Background To show the effects of various background latent initialization strategies, we provide additional examples, including the approach used in the texture synthesis method [34] and a background that contrasts sharply with the foreground object, as shown in Fig. 15. In detail, SyncMVD [33] encodes the backgrounds with alternative random solid color images. For the high-contrast background experiment, we use the latents obtained from the DDIM inversion of highly contrasted foreground and background to initialize the latents.



"a mossy stone monster is dancing in a mysterious forest."

Figure 14. **Visualization of generated textures for mesh sequences.** Our method effectively incorporates temporal changes, such as lighting variations and appearance transformations, directly into the textures, eliminating the need for post-production by technical artists.



(c) Ours w/o background priors (foreground composited with our background)

Figure 15. More ablation study on the background priors. We present three ablations, including the approach used in the texture synthesis method SyncMVD [34], a background that contrasts sharply with the foreground, and without background priors.

## **D** More Method Comparisons

#### D.1 Comparison with Depth-Conditioned Video Diffusion Models

While depth-conditioned video diffusion models effectively generate visually compelling results from a single viewpoint, they often struggle to maintain consistent multi-view representations of a single object, such as a character, across different perspectives. To highlight this limitation, we present multi-view results from the depth-conditioned video diffusion model in Fig. 17. The primary cause of this issue is that depth conditions are inherently view-dependent and lack global geometry information, in contrast to UV maps, which provide global information about the 3D space, enabling a unique mapping for each 3D point across all views.



Figure 16. **Results of textured mesh animation.** We present the visual results of Text2Tex [10] with our backgrounds. Text2Tex fails to capture temporal variations between frames and results in empty texels in invisible regions.

#### **D.2** Comparison with Textured Mesh Animations

In this section, we highlight the differences between our method and traditional approaches, demonstrating the effectiveness of 4D texturing in capturing temporal variations (e.g., lighting and appearance transformation) within mesh sequences to produce vivid visual results. Traditional methods typically involve texturing a base mesh (often called a clay mesh) and animating it using skinning techniques. This approach lacks the capacity to represent dynamic visual transformations and the animated sequence is then refined by technical artists who control scene lighting or simulate cloth dynamics to achieve the final visual presentation. This process is labor-intensive and demands specialized expertise in cinematic production and technical engines.

Instead, Tex4D introduces the first solution to dynamic texture generation by leveraging the expressiveness of video diffusion models. This is a fundamental task in filmmaking and game design to model character appearance transformations.

Compared with the static texture generation method, our method presents an alternative by directly integrating complex temporal changes into mesh sequences. As shown in Fig. 4, 11 and 12, our approach effectively captures temporal effects such as dynamic lighting, and evolving appearances using textual prompts, significantly simplifying the workflow while maintaining high-quality results.

We demonstrate the limitations of traditional textured mesh animation in handling complex temporal changes in Fig. 16. Specifically, we employ the Text2Tex [10] to generate the texture for the input mesh in T-pose and render it from multiple viewpoints. To ensure a fair comparison, we composite the rendered results with the background generated by our method. Notably, the 'ghost' and 'snowman' examples exhibit visible seams during animation due to self-occlusions that commonly appear in dynamic poses but cannot be accurately predicted during T-pose texture generation. Although the texture is still globally consistent, Text2Tex not just fails to model dynamic effects like appearance transformation but also results in empty texels and disrupts the visual continuity of the animation. For the rendered video results, please kindly refer to our supplementary videos.



"a mossy stone monster dances in a mysterious forest."



"a spirit in neon tilts its head, cyberpunk style."

Figure 17. **Multi-view results from conditioned video diffusion models.** The conditioned video diffusion models struggle to maintain consistent multi-view representations of a single object due to the depth condition being view-dependent and lacking global geometry information.

#### D.3 Comparison with Text-to-4D Methods

Although the setting of the Text-to-4D task is different from Tex4D as the mesh sequence is not given, we also provide a comprehensive comparison for Text-to-4D methods like SV4D [65] and L4GM [44] as in Fig. 5. In our experiments, we found that these Text-to-4D methods usually fail to generate plausible results and have these limitations:

- Multiview video-based methods (e.g., SV4D) struggle with consistency under significant motion. In our early attempts, we also found that the diversity of multiview attention-based inference heavily depends on dataset quality. SV4D handles only simple character animations.
- Animatable Gaussian-based methods (DreamGaussian4D, L4GM) built on LGM suffer from blurry and static textures, as decoupling geometry and appearance simultaneously is more challenging.

#### Please find the method that has best spatial and temporal consistency. The prompt is "a Stormtrooper swimming" Videos are capturing the same object from different views.













Figure 18. User Study. We provide more visual examples and include quantitative results from our user study. We evaluate the videos from three metrics: Appearance Quality, Spatial and Temporal Consistency, and Fidelity to the Prompt.



Figure 19. Complex UV with seams created by Blender Smart UV (angle limit=0.5). Tex4D is robust with different UV mappings. First column: UV without texels. Last column: rendered video frames. Other columns: texture generated by Tex4D. Zoom in for texture details.

## E User Study

Our study included 67 participants who provided 1,104 valid responses across six scenes from previous works, each rendered from two different viewpoints. We show each participant 30 pairs of videos synthesized by different methods, capturing the same object from different views. For each pair, each participant is asked three questions in sequence:

- Which method has better appearance quality?
- Which method has better spatial and temporal consistency?
- Which method has better fidelity to the prompts?

Our study involved 67 participants who provided a total of 1104 valid responses based on six different scenes drawn from six previous works, with each scene producing videos from two different views.

In addition, we further invite 24 participants who provided 455 valid trials to statistically evaluate user preferences (rating from 1 to 5) across Appearance quality, Spatio-temporal consistency, and Prompt Consistency. Users consistently favored Tex4D over all baselines.

## **F** Broader Discussion

As Tex4D is the first work to generate dynamic textures, we provide a broader analysis of our method in this section and hope to provide more insights for future works.

## F.1 Robustness of Different UV Mappings

To validate the robustness of our method in different UV mappings, we test Tex4D using Blender's smart UV strategy. Specifically, we set the angle limit for adjacent face normals to 30 degrees for splitting with more seams. The rendered results are also visually robust as shown in Fig 19.

## F.2 Highly Structured Texture Generation

To further demonstrate the robustness of our method, we test Tex4D with the task of highly structured texture generation that is commonly used in relief mapping and displacement mapping. We create a low-poly house as the base mesh and test our method with different prompts as shown in Fig. 20. Tex4D could generate multi-view consistent and highly structured texture, demonstrating our method is robust across different tasks.



Figure 20. **Highly structured texture generation.** A low-poly house is created with different prompts to demonstrate the capability of Tex4D to generate diverse, highly structured textures. Zoom in for texture details.



100 frames

200 frames

300 frames

Figure 21. Long Texture Sequence Generation. Our method could maintain plausible textures with the increase in animation sequences. However, the quality may degrade due to the training sets of video diffusion models usually do not contain such long sequences.

### F.3 Runtime Breakdown

In Tab. 2, we provide a runtime breakdown including different keyframe numbers and latent resolutions. We have tested the latent resolutions with  $96 \times 96$  and  $64 \times 64$ , and the keyframe numbers 3, 10, 17 and 25 respectively.

Keyframes	3	10	17	25
Time (min) w/ res. (96×96)	4.1	13.6	26.6	34.4
Time (min) w/ res. $(64 \times 64)$	2.3	6.8	10.7	14.4

Table 2. Quantitative results of runtime breakdown.

tively in the video generation task with a total frame number of 72. In our default setting, we use the keyframe interval of 3 and latent size  $96 \times 96$  for better generation quality.

#### F.4 Long Texture Sequence Generation

In this section, we extend the original mesh sequence to longer lengths to study the robustness of Tex4D at long texture sequence generation. Specifically, we use the "snowman" case, which originally has 100 frames. We extend the mesh sequence to 200 and 300 frames by repeating the animation. Tex4D is robust in generating long sequence textures while preserving the dynamics. We provide the result of intermediate frames in Fig. 21. However, the texture details may degrade due to the video diffusion models, which tend to distort high-frequency components in long videos, as indicated by [36]. Moreover, the training sets of video diffusion models usually do not contain such long sequences, which hinders the model to learn long sequence appearance as indicated by [12].

#### F.5 Analysis of Appearance Details

During our early experiments and literature review, we found video diffusion models tend to generate relatively smooth results compared with image diffusion models, as indicated by DiffusionRenderer [30] and I2VGen-XL [77]. As a result, in some cases, the clown's face lacks fine details within small pixel patches as shown in Fig. 9.

We analyze the underlying causes from three perspectives. Intuitively, maintaining high-frequency details temporally is harder than a smooth one, and video diffusion models tend to comprise the spatial details for better temporal consistency. Technically, this problem should be attributed to the data limitations, including dataset scales and low-quality patches from websites (e.g., WebVid10M [2]) used for video diffusion models, as analyzed by VideoCrafter2 [12]. Specifically for Tex4D, the latents

aggregation module (Sec. 4.2) may potentially wipe out details because of the average operation on latents, although with the weighted cosine maps. A possible step is to train a tiny adapter like TexFusion [7] to optimize the latents aggregation instead of RGB mapping.

Some concurrent works attempt to address this problem, either from the data distribution view (FreqPrior [72]) or the U-Net architecture view (CogVideoX [69]). As our methods can be easily integrated into any video diffusion model, we anticipate improving the quality of Tex4D with the advancement of video diffusion models.

## F.6 Societal Impacts

**Potential Social Impacts** Although our method offers broad potential for controllable generation tasks, like other frameworks for image and video synthesis, it could be misused for harmful purposes (such as generating deceptive content or fake media). As such, responsible and cautious use is essential when applying it in real-world scenarios.

**Safeguards** During inference, we enable the NSFW filter provided by the underlying models to help prevent the generation of explicit or inappropriate content, thereby safeguarding users from unwanted exposure. The training datasets used in our base model, CTRL-Adapter, have already filter out the image/video samples with harmful contents. Specifically, Panda70M filters harmful content using an automated pipeline and replaces names with "person" via NLTK. Similarly, LAION-POP applies a custom NSFW classifier to exclude unsafe samples.

## **G** Limitations

## G.1 Panoramic Background Modeling

One limitation of our method is the lack of seamless integration between the generated textures and the background, resulting in a disjointed appearance where the foreground and background elements may seem artificially stitched together. However, the dynamic textures remain globally consistent and can be directly applied to the downstream tasks, as shown in Fig. 10. To the best of our knowledge, no existing work tackles the foreground and background texture generation together because the task is computationally expensive, and the scene-level dataset is limited. Addressing the scene-level 4D texturing remains an open challenge that we aim to explore in future work.

## G.2 Computation Time

We notice that our method is relatively computationally intensive compared with other texture synthesis methods. The running time of our method primarily depends on the foundation model CTRL-Adapter, which takes approximately 5 minutes to generate a 24-frame video. We anticipate efficiency improvements with advancements in conditioned video diffusion models to further enhance the practicality of Tex4D.

## **H** Licenses

We use standard licenses from the community and provide the following links to the licenses for the datasets, codes, and models that we used in this paper. For further information, please refer to the specific link.

PyTorch [1]: BSD-style HuggingFace Transformers [62]: Apache License 2.0 HuggingFace Diffusers [59]: Apache License 2.0 CTRL-Adapter [32]: Apache License 2.0 ControlNet [74]: Apache License 2.0 I2VGen-XL [76]: MIT License