

Not All Frame Features Are Equal: Video-to-4D Generation via Decoupling Dynamic-Static Features

Liyang Yang¹, Chen Liu², Zhenwei Zhu¹, Ajian Liu³, Hui Ma¹, Jian Nong¹, Yanyan Liang^{1†}

¹ Macau University of Science and Technology ² The University of Queensland

³ Institute of Automation, Chinese Academy of Sciences (CASIA)

† Corresponding Author

Abstract

Recently, the generation of dynamic 3D objects from a video has shown impressive results. Existing methods directly optimize Gaussians using whole information in frames. However, when dynamic regions are interwoven with static regions within frames, particularly if the static regions account for a large proportion, existing methods often overlook information in dynamic regions and are prone to overfitting on static regions. This leads to producing results with blurry textures. We consider that decoupling dynamic-static features to enhance dynamic representations can alleviate this issue. Thus, we propose a dynamic-static feature decoupling module (DSFD). Along temporal axes, it regards the regions of current frame features that possess significant differences relative to reference frame features as dynamic features. Conversely, the remaining parts are the static features. Then, we acquire decoupled features driven by dynamic features and current frame features. Moreover, to further enhance the dynamic representation of decoupled features from different viewpoints and ensure accurate motion prediction, we design a temporal-spatial similarity fusion module (TSSF). Along spatial axes, it adaptively selects similar information of dynamic regions. Hinging on the above, we construct a novel approach, DS4D. Experimental results verify our method achieves state-of-the-art (SOTA) results in video-to-4D. In addition, the experiments on a real-world scenario dataset demonstrate its effectiveness on the 4D scene. Project page: <https://github.com/LiyangCV/DS4D>.

1. Introduction

Generating dynamic 3D (4D) content [7, 11–13, 26, 28, 32, 39] from video is an essential research topic involving the field of computer vision and computer graphics. However, it is rather formidable to predict accurate motion from a few viewpoints while ensuring high-quality generation.

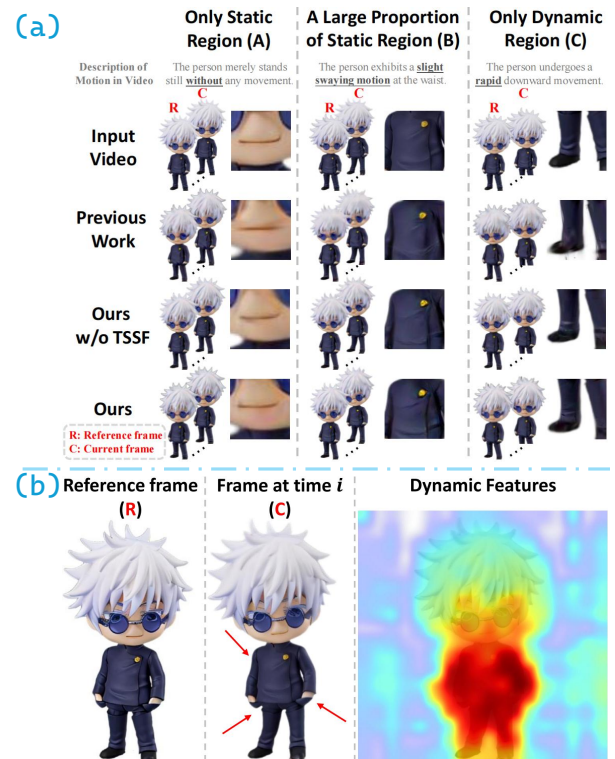


Figure 1. (a) Illustration of the issues caused by different proportions of dynamic and static regions. Previous work [51] generates the 4D content with obviously blurry textures in the dynamic regions with B-type video input. In contrast, our methods with decoupling dynamic-static features generates high-quality results with clear textures. (b) The visualization of dynamic features in our method. The red region highlights the primary region of interest in the dynamic region in B-type videos. This part is also the dynamic region between frames R and C. It demonstrates our method can successfully decouple dynamic-static features.

There are two main streams in current approaches to improving generation quality. Inference-based methods [15, 30, 43, 52] can generate high-quality 4D content

by capturing temporal-spatial correlations in its 4D diffusion model. Another stream is optimization-based methods [34, 41, 49, 51]. These methods generate 4D content through distilling spatial prior knowledge from pre-trained multi-view [17, 33] diffusion models. However, *previous methods only model temporal-spatial correlations using whole information in frames, but fail to explicitly differentiate between the dynamic and static regions within a frame*. If static regions account for a significant portion, previous methods overlook dynamic information. Thus they tend to overfit static regions, resulting in a diminished capacity to perceive texture variations in dynamic regions. As shown in Fig.1 (a), previous work (STAG4D [51]) produces results with blurry texture (e.g., wrinkles in clothes), particularly when confronted with B-type input video[†].

To tackle above problems, we present a novel approach DS4D, which decouples dynamic and static features along temporal and spatial axes to enhance dynamic representations for high-quality 4D generation. In particular, we propose dynamic-static feature decoupling module (DSFD) to obtain decoupled features. We assume some regions of current frame features exhibit significant differences relative to reference frame features. Such regions always contain important knowledge of texture, shape variations, and motion trends under current timestamps. Therefore, it can be regarded as the dynamic feature of the current frame. In contrast, the remaining parts are the static features. Based on this assumption, we decompose the dynamic components between each frame feature and reference frame features along the temporal axes. The dynamic component (also termed dynamic features), as difference features between two frame features, is able to represent the dynamic information (as shown in Fig.1 (b)). At last, we acquire decoupled features driven by dynamic components and corresponding frame features.

Note that spatial occlusion leads to the failure of dynamic components captured from a specific viewpoint to inadequately represent the inherent dynamic information in 4D space. To mitigate this issue, we design temporal-spatial similarity fusion module (TSSF) which engages in enhancing features’ dynamic representations. Firstly, Gaussian points are initialized by utilizing a large reconstruction model [45]. Then, point features are obtained by retrieving decoupled features for Gaussian points via view projection in TSSF. Subsequently, TSSF produces fused Gaussian features by adaptively selecting information on dynamic regions containing similar texture, shape, and motion representations from point feature space at same timestamps along spatial axes. Finally, fused Gaussian features with strong dynamic representation, are used for 4D generation.

The contributions can be summarized as follows:

- To our best knowledge, we are the first to propose a

[†]There is a large proportion of static regions within the video.

framework DS4D that decouples dynamic-static information in frames along temporal-spatial axes to enhance dynamic representations for high-quality 4D generation.

- Leveraging significant differences between frame features, we propose dynamic-static feature decoupling module (DSFD) to decouple dynamic and static features.
- We tackle the issue of inadequate dynamic information in 4D space resulting from spatial occlusion by designing temporal-spatial similarity fusion module (TSSF). TSSF enhances dynamic representations in features.
- Experimental results on Consistent4D dataset [9] and Objaverse dataset [4] demonstrate our DS4D outperforms other SOTA methods in terms of video quality, motion fidelity, and temporal-spatial consistency. Furthermore, the experiments on a complex real-world scenario dataset [11] verify its effectiveness on 4D scenes.

2. Related Works

2.1. 3D Generation

3D generation aims to produce 3D assets using images or text descriptions. The early works generate 3D objects in voxels [3, 42, 46, 53], meshes [23, 37, 38] or point clouds [1, 6, 19] forms. Recently, thanks to the popular application of NeRF [21, 22] and 3D Gaussian Splatting [10, 50], and the successful pre-trained text-to-image diffusion model [31] in 2D generation tasks, most works including DreamFusion [27] Dreamgaussian [35] that focus on 3D generation try to use Score Distillation Sampling (SDS) loss [27] to explore the possibility of inspiring potential 3D-aware from diffusion. However, these works always suffer from over-saturation, over-smoothing, and multi-face problems. Because it is a challenge to distill the invisible views while ensuring multi-view consistency based on 2D diffusion. In contrast, Zero-1-to-3 [17] and Zero123++ [33] directly train a 3D-aware diffusion model using multi-view images and corresponding camera poses. Zero-1-to-3 and Zero123++ usually generate more stable and high-quality novel views thanks to the spatial perception built into them. Therefore, in our method, we use Zero123++ to predict the novel views at all timestamps to provide additional multi-view sequences for optimization.

2.2. 4D Generation

Compared to 3D generation, 4D generation is a more challenging task. Current 4D representations are two main streams, including NeRF-based [5, 21, 25, 28, 36, 40] and 3D Gaussians-based [10, 13, 14, 16, 18, 39, 48]. Consistent4D [9], as the NeRF-based method, leverages the prior knowledge from pre-trained 2D diffusion models to optimize dynamic NeRF by Score Distillation Sampling (SDS) optimization. Dreamgaussian4D [29], STAG4D [51] and SC4D [41] both attempt to introduce dynamic 3D Gaussians

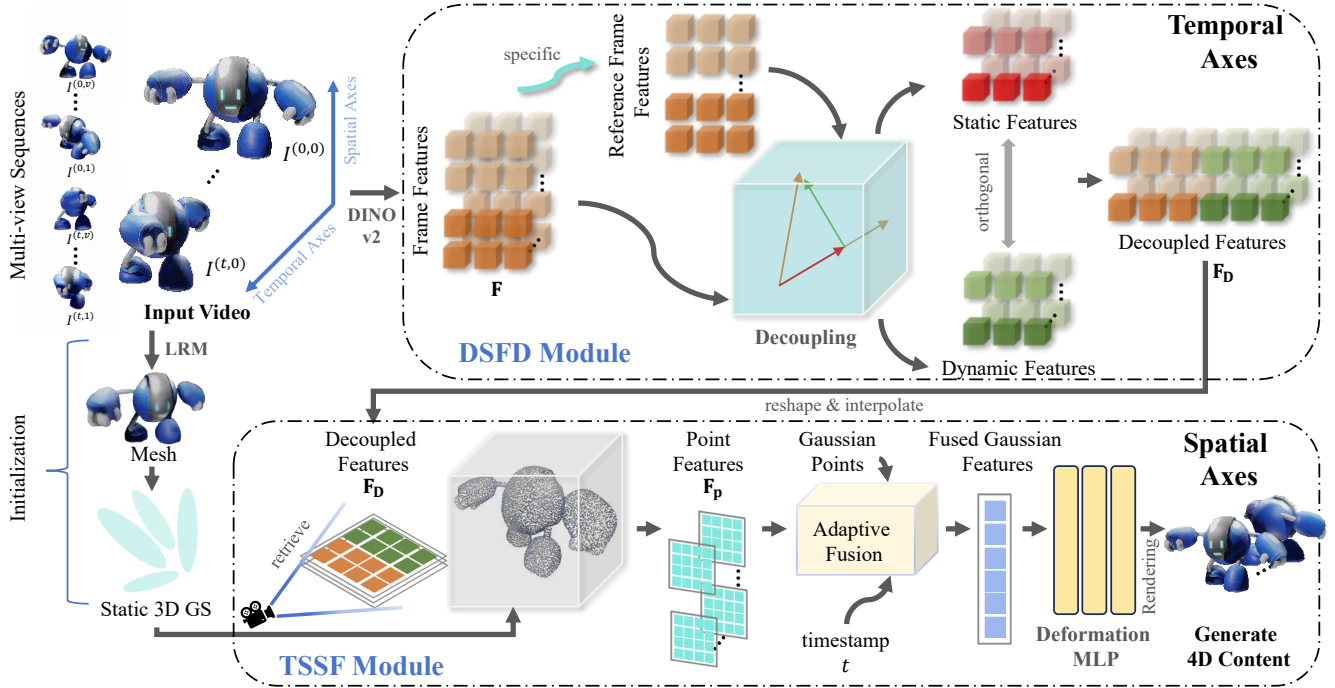


Figure 2. **Overview of our proposed DS4D.** Given an input video and corresponding multi-view sequences, our DS4D decouples features of the frame at time t based on the reference frame in DSFD module. Next, we acquire point features by retrieving each decoupled feature for Gaussian points via view projection, and we obtain fused Gaussian features by adaptively selecting similar dynamic information from point feature space in TSSF module. Finally, through Deformation MLP, our method generates 4D content.

in 4D generation. Dreamgaussian4D notably reduces the cost of optimization time by improving the training strategies. STAG4D can generate anchor multi-view videos for optimization via a training-free strategy. SC4D optimizes sparse-controlled dynamic 3D Gaussians through SDS loss. However, these methods model temporal-spatial correlations using whole video but ignore the dynamic information in frames when static regions account for a large portion. It leads to them easily overfitting on static parts.

3. Methodology

3.1. Overview

In this section, we provide an overview of DS4D, which includes initialization in Sec.3.2, dynamic-static feature decoupling (DSFD) in Sec.3.3, and temporal-spatial similarity fusion (TSSF) in Sec.3.4. The overall framework of our proposed method is illustrated in Fig.2.

3.2. Initialization

Formally, we begin with a single-view video. For each frame in the video, we adopt Zero123++ [33] to infer the \dagger pseudo multi-view images of frames. Then we obtain frame sequences, $\mathbf{I} = \{I^{(0,0)}, I^{(0,1)}, \dots, I^{(i,j)}, \dots, I^{(t,v)}\}$, where t represents the number of frames, and v represents the number of views of each frame.

Frame Feature Extraction. Owing to the superior feature extraction capability of the visual foundation model, we employ it as the feature extractor to obtain high-quality features. Thanks to the strong feature extraction ability of DINOv2 [24], we leverage DINOv2 to extract frame features $\mathbf{F} = \{f^{(0,0)}, f^{(0,1)}, \dots, f^{(i,j)}, \dots, f^{(t,v)}\}$ from sequences $\mathbf{I} = \{I^{(0,0)}, I^{(0,1)}, \dots, I^{(i,j)}, \dots, I^{(t,v)}\}$, where $f^{(i,j)} \in \mathbb{R}^{P \times D}$, P and D are the number of tokens and token dimension, respectively. Each feature $f^{(i,j)}$ encodes the geometry and texture of the corresponding frame.

Gaussian Points Initialization. Previous works [41, 51] initialize Gaussian points randomly, leading to unstable topology during optimization and ultimately compromising the quality of results. To solve the above problem, we initialize static 3D Gaussian points through pre-defined point clouds. Specifically, we employ a large reconstruction model [45] to generate the point clouds from the middle frame. This initialization approach provides a geometric prior and ensures the stability of subsequent optimization.

3.3. Dynamic-Static Feature Decoupling (DSFD)

We employ a set of frame features \mathbf{F} , which consists of all frame features and their corresponding multi-view features,

[†]The discussion about the impact of multi-view sequence quality and the robustness of our methods can be seen in supplementary material (Sec.6.1).

to decouple dynamic-static features. Specifically, our decoupling process involves two aspects: 1) Finding out reference frame features; 2) Decoupling dynamic and static features along temporal axes.

Finding out reference frame features. Decoupling frame-by-frame features usually consumes considerable computation time. Hence, we specify special frames as references for decoupling. It is critical to determine whether reference frame features can accurately represent the semantic content of the entire video. Considering that we initialize static Gaussian points at the middle frame, the whole model needs to capture the texture, shape, and motion transformation from the middle frame to any frame during training and inference. Therefore, we choose the frame features $f^{(\frac{k}{2},j)}$ of the middle timestamp as one of the reference frame features, which represent the semantic. Furthermore, to represent the motion trends in video, we need a kind of feature that can describe the motion variations of the entire video. Hence, we also specify the \dagger average frame features $\bar{f}^{(\bar{t},j)}$ as reference frame features, which represent the average motion variations.

Decoupling dynamic and static features along temporal axes. To reliably decouple dynamic and static information in frame features, we need to obtain the regions of current frame features that exhibit significant differences relative to reference frame features. These regions contain important knowledge of texture, shape variations, and motion trends under current timestamps. The regions with significant differences mean dynamic parts between the two frame features, while regions with similar information mean static parts between the two frame features. Therefore, we propose the dynamic-static feature decoupling module (DSFD), as shown in Fig.3.

Concretely, given specific j -th view in timestamp i , we project frame features $f^{(i,j)}$ onto reference frame features r^j . **The projection represents the semantic overlapping part of the two frames. In other words, regions in the two frames with no texture and shape transformations are identified as the static part.** The process is mathematically formulated as follows:

$$f_{\text{static}}^{(i,j)} = \left(\frac{f^{(i,j)} \cdot r^j}{\|r^j\|_2} \right) \frac{r^j}{\|r^j\|_2}, \quad (1)$$

where \cdot denotes dot product.

The difference between current frame features and static features means the orthogonal vector. Then, we identify the orthogonal vector as the semantic difference between the current frame and the reference frame. Hence, the orthogonal vector can be regarded as dynamic features:

$$f_{\text{dynamic}}^{(i,j)} = f^{(i,j)} - f_{\text{static}}^{(i,j)}. \quad (2)$$

At each timestamp, we obtain the dynamic features in a particular view. Dynamic features contain significant differ-




Figure 3. An overview of the decoupling architecture in DSFD. The decoupling architecture is shown on the left. The demonstration of projection is shown on the right.

ence information, which represents the motion knowledge of the current frame relative to the reference frame. Meanwhile, frame features contain texture and shape attributes of the current frame, which are useful for subsequently texture and shape learning. Hence, we leverage dynamic features to supplement dynamic information in the frame features while retaining the inherent attributes of the frame features.

Here, we append each token of the current frame features $f^{(i,j)}$ with dynamic features $f_{\text{dynamic}}^{(i,j)}$. Then, the appended features $f_d^{(i,j)}$ are integrated based on both the temporal i and spatial j , which are named decoupled features $\mathbf{F}_D = \{f_d^{(0,0)}, f_d^{(0,1)}, \dots, f_d^{(i,j)}, \dots, f_d^{(t,v)}\}$, where t and v denote the number of frames and number of views, respectively.

3.4. Temporal-Spatial Similarity Fusion (TSSF)

Since the features in \mathbf{F}_D obtained from a particular viewpoint do not adequately represent the complete dynamic information in the 4D space, we design temporal-spatial similarity fusion module (TSSF) to adaptively select similar dynamic information from decoupled features under different views at the same timestamp.

Inspired by [20], we retrieve the decoupled features from Gaussian points by view projection. Then we integrate decoupled features into point features based on the given camera pose. The Gaussian points are well-aligned with decoupled features thanks to the projection operations. Therefore, point features $\mathbf{F}_P = \{f_p^{(0,0)}, f_p^{(0,1)}, \dots, f_p^{(i,j)}, \dots, f_p^{(t,v)}\}$ encapsulate extensive dynamic information from decoupled features. In fact, among different viewpoints, the texture, shape, and motion representations from the same spatial area of the object are similar in the same timestamps. Based on this fact, utilizing the semantic similarities in point features across viewpoints and aggregating them contributes to the enhancement of dynamic representations in point features. Hence, we propose an adaptive fusion in Fig.4.

At the same timestamp, we design global awareness fusion (GA) in adaptive fusion to merge similar dynamic information from point features along spatial axes. However,

\dagger Average frame features $\bar{f}^{(\bar{t},j)}$ are derived from the average of all frame features from a particular view.




Figure 4. An overview of adaptive fusion in TSSF.

under conditions of significant differences among multi-views, the effect of merging in GA is affected by that dynamic regions are noticeably obscured in specific viewpoints. Therefore, we need to reduce the impact of point features under these views for merging while retaining the useful dynamic information in these features. Based on this condition, we design distance awareness fusion (DA).

Global awareness (GA). At time i , the fully-connection layer (FC) generates score maps $\mathbf{W} = \{w^{(0,0)}, \dots, w^{(t,v)}\}$ for the point features of each view,

$$\mathbf{W} = \text{Softmax}(\text{FC}(\mathbf{F}_p)). \quad (3)$$

Then, we fuse them into fused point features by the weighted summation of all point features according to their score map. The fused features f_a^i can be calculated as:

$$f_a^i = \sum_{j=0}^v w^{(i,j)} f_p^{(i,j)}. \quad (4)$$

Distance awareness (DA). Frames from single-view video are regarded as front view. The front frame reveals real and visible motion areas of the object at various timestamps compared to the other views. This information within the front frame requires preservation as much as possible while supplementing the useful dynamic information from other views for accuracy motion prediction. Thus, at time i , we calculate the L_1 distance between point features from the front view and point features from the other views. The distance quantifies the difference information between point features extracted from the front view and those derived from other viewpoints. The difference information contains dynamic information that cannot be represented by the point features of the front view. Then, we append point features from other views with distance features and merge them by Eq. (3) and Eq. (4). The fused point features from other views are appended to point features from the front view. We merge these features by Eq. (3) and Eq. (4) to acquire the fused point features f_a^i . Through the above pre-processing, we reduce the impact of point features under other views according to the front view while retaining useful dynamic information in point features under other views.

In addition to fused features through adaptive fusion, we also present dynamic Gaussian features as 3D Gaussian features with HexPlane [2] for inherently regularizing features of fields and guaranteeing their smoothness. Formally, the dynamic Gaussian features at time i can be defined as:

$$f_{hg}^i = \text{HexPlane}(X_i, s_i, \gamma_i, \sigma, \zeta), \quad (5)$$

where $X_i = (x_i, y_i, z_i)$ denotes the position of Gaussian points, s_i and γ_i denote scale and rotation at time i . σ and ζ represent opacity and spherical harmonic coefficients of the radiance. Then, we combine dynamic Gaussian features f_{hg}^i and fused point features f_a^i at time i , and map them to fused Gaussian features via learnable linear transformation. The fused Gaussian features represent Gaussian intrinsic properties while capturing rich dynamic information. Therefore, fused Gaussian features are essential for accurately predicting deformation using Deformation MLP.

3.5. Training Objectives

Following [51], we calculate SDS loss and the photometric loss based on rendered views and ground truth images. Inspired by [8, 45, 54], we introduce the LPIPS loss to minimize the similarity between pseudo multi-view images and rendered views. More details can be found in the supplementary material.

4. Experiments

4.1. Implementation Details

The implementation details are provided in the section. More details can be found in supplementary material.

Datasets In our article, we use **four** challenging datasets.

- 1) We utilize the dataset provided by Consistent4D.
- 2) We introduce a subset of Objaverse [4, 15]. The objects in Objaverse have more complexity motion than Consistent4D dataset.
- 3) We introduce some challenging videos from online sources.
- 4) In Sec.5, to evaluate the effectiveness of our method in real-world scenarios, we utilize three real-world scenarios provided by Neu3D’s [11] dataset.

Metrics Following Consistent4D [9] and STAG4D [51], we measure the 4D generation quality of our methods using CLIP, LPIPS, FVD and FID-VID. In Sec.5, following [39], we evaluate our results using PSNR, LPIPS, SSIM, and structural dissimilarity index measure (D-SSIM).

Training Our two models, including DS4D-GA (using GA in TSSF) and DS4D-DA (using DA in TSSF) use the same training setting.

4.2. Comparisons with Existing Methods

In this section, we compare our methods DS4D-GA and DS4D-DA with the SOTA methods on Consistent4D dataset and Objaverse dataset, including Consistent4D [9], Dreamgaussian4D [29], STAG4D [51], SC4D [41], 4Diffusion

Methods	Optimization	Consistent4D dataset				Objaverse dataset			
		CLIP \uparrow	LPIPS \downarrow	FVD \downarrow	FID-VID \downarrow	CLIP \uparrow	LPIPS \downarrow	FVD \downarrow	FID-VID \downarrow
Consistent4D [9]	✓	0.9085	0.1316	1041.2242	28.6471	0.8491	0.2222	1814.5652	48.7921
Dreamgaussian4D [29]	✓	0.9145	0.1517	844.9087	37.9977	0.8127	0.2017	1545.3009	58.3686
STAG4D [51]	✓	0.9078	0.1354	986.8271	26.3705	0.8790	0.1811	1061.3582	30.1359
SC4D [41]	✓	0.9117	0.1370	852.9816	26.4779	0.8490	0.1852	1067.7582	40.5130
4Diffusion [52]	✗	0.8734	0.2284	1551.6363	149.6170	-	-	-	-
L4GM [30]	✗	0.9158	0.1497	898.0604	31.4996	-	-	-	-
DS4D-GA (Ours)	✓	0.9206	0.1311	799.9367	26.1794	0.8868	0.1761	890.2646	26.6717
DS4D-DA (Ours)	✓	0.9225	0.1309	784.0235	24.0492	0.8881	0.1759	870.9489	25.3836

Table 1. Evaluation and comparison of the performance on Consistent4D dataset and Objaverse dataset. The best score is highlighted in bold. All the experiments of the methods are carried out using the code from their official GitHub repository. For a fair comparison, the experiment of and L4GM on Objaverse dataset are disregarded since they are inference-based methods trained on this dataset.

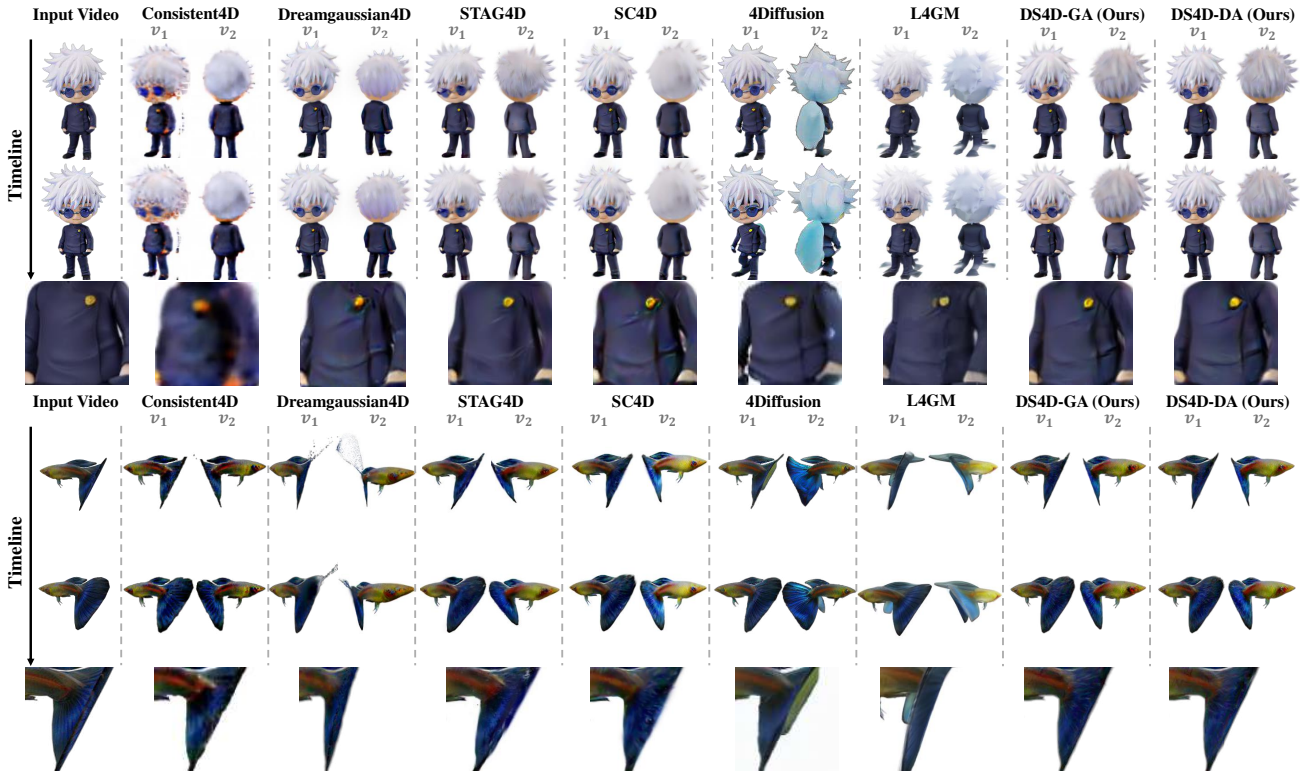


Figure 5. Qualitative comparison on video-to-4D generation. For each method, we render results under two novel views at two timestamps. More comparison examples, and some of the 360° generation results (video file) can be found in supplementary materials.

[52] and L4GM [30]. We conduct qualitative and quantitative comparisons, respectively. The superior performance demonstrates the effectiveness of our methods.

Quantitative comparisons. The quantitative results are shown in Tab. 1. Our methods DS4D-GA and DS4D-DA, consistently outperform other methods in all metrics. Specifically, our methods notably exceed the SOTAs in FVD, indicating that our generation results have fewer temporal artifacts than others. Furthermore, the experiments on Objaverse dataset present that our methods significantly outperform other SOTAs by a large margin. It demonstrates

our methods are superior in terms of temporal-spatial consistency, fidelity, and quality of generation results and underscores the robustness of our methods. In summary, such significant improvements are attributed to our method of decoupling dynamic-static features, which can explicitly distinguish dynamic and static regions within frame features.

Qualitative comparisons. The qualitative results on Consistent4D dataset are presented in Fig. 5. Besides, the results on Objaverse dataset are shown in Fig. 6. Obviously, whether in the example with a larger proportion of static regions (the person in Fig. 5 and warrior in Fig. 6) or the exam-

ple with a balanced proportion of dynamic and static regions (guppie in Fig. 5), the results, generated by Consistent4D, Dreamgaussian4D, STAG4D and SC4D, have different degrees of blurriness in the details, especially in the areas with motion trends in the current and subsequent frames. Although 4Diffusion and L4GM generate videos with clearer textures than other previous methods with the help of a large amount of 4D data prior, it is easy to generate the abnormal shape or textures with inconsistent details in some areas (e.g, the back of the person). These methods do not differentiate between dynamic and static information, leading them to easily overlook information in dynamic regions when faced with large proportions of static. In contrast, our methods address the challenges arising from varying proportions of dynamic and static regions in frames, resulting in the achievement of high-quality 4D generation.




Figure 6. Qualitative comparison on video-to-4D generation based on Objaverse dataset.

4.3. Ablation Experiments

To evaluate the effectiveness of different components in our methods, we conduct ablation experiments on Consistent4D and Objaverse datasets as shown in Tab.2. More analysis can be found in supplementary material.

Effect of DSFD and Decoupling. Validating the effect of DSFD also means Validating the effect of decoupling since DSFD consists of decoupling. Therefore, we conduct experiments in two aspects: a) Adding frame features in the model without decoupling (labeled as D). b) Performing decoupling on frame features in the model (labeled as E, also termed DSFD). In particular, the performance of D and E both have significant improvement compared to A and B, especially FVD score. For D, the frame features provide enough structure prior knowledge and texture information for Deformation MLP. It assists Deformation MLP in capturing shape deformation. Nevertheless, D fails to differentiate between dynamic and static regions of frames clearly. Consequently, D is prone to overfitting in static areas, resulting in blurry details. For example, the blurred details of legs from D in Fig.7. In contrast, DSFD, which decouples dynamic-static features, mitigates the issue.

Effect of TSSF. To validate the effect of TSSF, we add




Figure 7. Ablation on different experiment settings from Tab.2.




Figure 8. (a) Visualization on the heatmap of dynamic features in DSFD. The red region highlights the primary region of interest in dynamic features. (b) The score map of point features in TSSF-GA and TSSF-DA. The red area indicates high attention to dynamic information in point features of a specific view.

the TSSF based on model E, including model F, G (TSSF-GA) and H (TSSF-DA). Among them, we replace adaptive fusion with average pooling in model F. Specifically, we observe that E, G, H can generate more texture details than other models. Because they mitigate the problem brought by spatial occlusion by aggregating dynamic components from different viewpoints. However, F only averages the point features from DSFD along spatial axes, which ignores merging similar dynamic information. The model G and F adaptively select dynamic information, which plays a crucial role in enhancing dynamic representations. Compared with E, the texture of 4D content generated by F and H is more refined in Fig.7. This underscores the efficacy of merging similar dynamic information in TSSF, including TSSF-GA and TSSF-DA.

4.4. Visualization

In this section, we perform several visualizations of features in DSFD and TSSF, respectively.

DSFD. Fig.8 (a) presents the heatmap of dynamic features obtained by DSFD decoupling features from the current and reference frame features. The red area indicates the primary region of interest in the features. We observe the motion trends come from the trunk and body of elephant in the front view (left of the figure). Thus, dynamic features show strong concern in this area. Limited to the viewpoint range, the motion trajectory for novel views (right of the figure) is more interested in the elephant’s head. No doubt, dynamic features also present a high response in similar areas. In conclusion, our method acquires accurate dynamic features using DSFD, as supported by the visualizations.

Methods	Point Initialization	Consistent4D dataset				Objaverse dataset			
		CLIP \uparrow	LPIPS \downarrow	FVD \downarrow	FID-VID \downarrow	CLIP \uparrow	LPIPS \downarrow	FVD \downarrow	FID-VID \downarrow
A	\times	0.9133	0.1341	953.6300	27.3747	0.8736	0.1816	1072.1292	28.2410
B	\checkmark	0.9151	0.1313	913.371	27.1357	0.8763	0.1801	1062.9398	28.0977
C. w/ LPIPS Loss	\checkmark	0.9163	0.1311	899.5714	27.0836	0.8773	0.1804	1016.2576	27.8062
D. w/ Frame Features	\checkmark	0.9174	0.1350	888.6579	26.8486	0.8805	0.1778	1005.7503	27.6334
E. w/ DSFD	\checkmark	0.9186	0.1333	861.6075	26.5403	0.8827	0.1765	989.0834	26.9199
F. w/ TSSF-Average Pooling	\checkmark	0.9194	0.1313	839.6600	26.5071	0.8848	0.1761	951.9127	26.8412
G. w/ TSSF-GA (Ours)	\checkmark	0.9206	0.1311	799.9367	26.1794	0.8868	0.1761	890.2646	26.6717
H. w/ TSSF-DA (Ours)	\checkmark	0.9225	0.1309	784.0235	24.0492	0.8881	0.1759	870.9489	25.3836

Table 2. The ablation experiments on Consistent4D dataset and Objaverse dataset. Each setup is based on a modification of the immediately preceding setups. The best score is highlighted in bold.

Method	PSNR \uparrow	SSIM \uparrow	D-SSIM \downarrow	LPIPS \downarrow
4D-GS [39]	32.1598	0.9483	0.0132	0.1422
Ours	32.3964	0.9494	0.0123	0.1434

Table 3. Evaluation of the performance on Neu3D’s dataset.

TSSF. Fig.8 (b) shows the score map of adaptively selecting similarity dynamic information from point features of different views by TSSF-GA and TSSF-DA. Specifically, the leg movements of the triceratops indicate the primary trend in motion between the middle and current time. The red area indicates the high attention to dynamic information in point features of a specific view. Meanwhile, this area is also similar to the dynamic area of other views in Fig.8 (b). Hence, it reveals that the two approaches can capture a certain degree of similar dynamic information from different viewpoints. However, TSSF-GA is interested in the back of triceratops rather than legs since the front legs are nearly invisible in the novel view v_2 . The unseen dynamic region in the novel view influences the accuracy of selecting similar dynamic information. In contrast, TSSF-DA predicts highest scores on the legs under view v_1 and v_2 , thanks to reducing the impact of novel views in TSSF-DA. It indicates TSSF-DA can alleviate issues caused by the fact that dynamic regions are noticeably obscured in novel views.

5. Discussion on Real-World Scenario

Generating a dynamic scene from input real-world videos is an essential task. In this section, we conduct experiments on three scenes from Neu3D’s dataset to demonstrate the effectiveness of DSFD and TSSF on complexity real-world dynamic scene. We set 4D-GS [39] as the baseline, which has shown promising results in this task. To fair comparison, we directly insert our DSFD and TSSF into 4D-GS. More details can be found in supplementary material.

The results are shown in Tab.3 and Fig.9. Our method achieves competitive results, demonstrating significant potential even for intricate real-world scenarios. Due to static nature of a significant portion of the real-world scene, the previous method easily encounters challenges in exploring




Figure 9. Visualization of flame steak compared with 4D-GS.

dynamic regions, consequently impacting the quality of outputs. Hence, it is necessary to decompose dynamic and static regions to mitigate this issue. In essence, significant differences exhibit texture variations and motion trends between frames. Thus it is beneficial to utilize this nature to decouple dynamic-static information at feature-level, which enables us to enhance fine-grained dynamic representations. Indeed, there are additional methods available to decouple effectively. For instance, we can introduce optical flow and utilize 3D-aware foundation model [44, 47] to extract depth features used to decouple.

6. Conclusion

In this paper, we propose a novel framework DS4D, which decouples dynamic-static information along temporal-spatial axes to enhance dynamic representations for high-quality 4D generation. Its DSFD decouples features by regarding the regions of current features that possess significant differences relative to reference frame features as dynamic features. Moreover, TSSF is designed to enhance dynamic representations by selecting similar dynamic information. Overall, our method can produce high-quality 4D content and shows promise in 4D scene generation.

Acknowledgements This work is supported by the National Key Research and Development Plan under Grant 2021YFE0205700, Science and Technology Development Fund of Macau project 00123/2022/A3, 0096/2023/RIA2.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International conference on machine learning*, pages 40–49. PMLR, 2018. 2
- [2] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 130–141, 2023. 5
- [3] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016. 2
- [4] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. 2, 5
- [5] Jiemin Fang, Taoran Yi, Xinggang Wang, Lingxi Xie, Xiaoopeng Zhang, Wenyu Liu, Matthias Nießner, and Qi Tian. Fast dynamic radiance fields with time-aware neural voxels. In *SIGGRAPH Asia 2022 Conference Papers*, pages 1–9, 2022. 2
- [6] Matheus Gadelha, Rui Wang, and Subhransu Maji. Multiresolution tree networks for 3d point cloud processing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 103–118, 2018. 2
- [7] Hang Gao, Ruilong Li, Shubham Tulsiani, Bryan Russell, and Angjoo Kanazawa. Monocular dynamic view synthesis: A reality check. *Advances in Neural Information Processing Systems*, 35:33768–33780, 2022. 1
- [8] Yicong Hong, Kai Zhang, Jiuxiang Gu, Sai Bi, Yang Zhou, Difan Liu, Feng Liu, Kalyan Sunkavalli, Trung Bui, and Hao Tan. Lrm: Large reconstruction model for single image to 3d. *arXiv preprint arXiv:2311.04400*, 2023. 5
- [9] Yanqin Jiang, Li Zhang, Jin Gao, Weiming Hu, and Yao Yao. Consistent4d: Consistent 360° dynamic object generation from monocular video. In *The Twelfth International Conference on Learning Representations*, 2024. 2, 5, 6
- [10] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 2
- [11] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5521–5531, 2022. 1, 2, 5
- [12] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6498–6508, 2021.
- [13] Zhiqi Li, Yiming Chen, and Peidong Liu. Dreammesh4d: Video-to-4d generation with sparse-controlled gaussian-mesh hybrid representation. *Advances in Neural Information Processing Systems*, 37:21377–21400, 2024. 1, 2
- [14] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8508–8520, 2024. 2
- [15] Hanwen Liang, Yuyang Yin, Dejia Xu, Hanxue Liang, Zhangyang Wang, Konstantinos N Plataniotis, Yao Zhao, and Yunchao Wei. Diffusion4d: Fast spatial-temporal consistent 4d generation via video diffusion models. *arXiv preprint arXiv:2405.16645*, 2024. 1, 5
- [16] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. Gaufré: Gaussian deformation fields for real-time dynamic novel view synthesis. *arXiv preprint arXiv:2312.11458*, 2023. 2
- [17] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 2
- [18] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *2024 International Conference on 3D Vision (3DV)*, pages 800–809. IEEE, 2024. 2
- [19] Shitong Luo and Wei Hu. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2837–2845, 2021. 2
- [20] Luke Melas-Kyriazi, Christian Ruppert, and Andrea Vedaldi. Pc2: Projection-conditioned point cloud diffusion for single-image 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12923–12932, 2023. 4
- [21] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2
- [22] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 2
- [23] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3504–3515, 2020. 2

- [24] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3
- [25] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 2
- [26] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. 1
- [27] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [28] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 2
- [29] Jiawei Ren, Liang Pan, Jiaxiang Tang, Chi Zhang, Ang Cao, Gang Zeng, and Ziwei Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 2, 5, 6
- [30] Jiawei Ren, Kevin Xie, Ashkan Mirzaei, Hanxue Liang, Xiaohui Zeng, Karsten Kreis, Ziwei Liu, Antonio Torralba, Sanja Fidler, Seung Wook Kim, et al. L4gm: Large 4d gaussian reconstruction model. *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 1, 6
- [31] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022. 2
- [32] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16632–16642, 2023. 1
- [33] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 2, 3
- [34] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation. *arXiv preprint arXiv:2301.11280*, 2023. 2
- [35] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [36] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12959–12970, 2021. 2
- [37] Nanyang Wang, Yinda Zhang, Zhuwen Li, Yanwei Fu, Wei Liu, and Yu-Gang Jiang. Pixel2mesh: Generating 3d mesh models from single rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pages 52–67, 2018. 2
- [38] Chao Wen, Yinda Zhang, Chenjie Cao, Zhuwen Li, Xiangyang Xue, and Yanwei Fu. Pixel2mesh++: 3d mesh generation and refinement from multi-view images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2166–2180, 2022. 2
- [39] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20310–20320, 2024. 1, 2, 5, 8
- [40] Tianhao Wu, Fangcheng Zhong, Andrea Tagliasacchi, Forrester Cole, and Cengiz Oztireli. D²nerf: Self-supervised decoupling of dynamic and static objects from a monocular video. *Advances in neural information processing systems*, 35:32653–32666, 2022. 2
- [41] Zijie Wu, Chaohui Yu, Yanqin Jiang, Chenjie Cao, Fan Wang, and Xiang Bai. Sc4d: Sparse-controlled video-to-4d generation and motion transfer. *arXiv preprint arXiv:2404.03736*, 2024. 2, 3, 5, 6
- [42] Haozhe Xie, Hongxun Yao, Shengping Zhang, Shangchen Zhou, and Wenxiu Sun. Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images. *International Journal of Computer Vision*, 128(12): 2919–2935, 2020. 2
- [43] Yiming Xie, Chun-Han Yao, Vikram Voleti, Huaizu Jiang, and Varun Jampani. Sv4d: Dynamic 3d content generation with multi-frame and multi-view consistency. *arXiv preprint arXiv:2407.17470*, 2024. 1
- [44] Haofei Xu, Songyou Peng, Fangjinhua Wang, Hermann Blum, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Depthsplat: Connecting gaussian splatting and depth. *arXiv preprint arXiv:2410.13862*, 2024. 8
- [45] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 2, 3, 5
- [46] Liying Yang, Zhenwei Zhu, Xuxin Lin, Jian Nong, and Yanyan Liang. Long-range grouping transformer for multi-view 3d reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18257–18267, 2023. 2

- [47] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024. [8](#)
- [48] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642*, 2023. [2](#)
- [49] Yuyang Yin, Dejia Xu, Zhangyang Wang, Yao Zhao, and Yunchao Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. [2](#)
- [50] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. [2](#)
- [51] Yifei Zeng, Yanqin Jiang, Siyu Zhu, Yuanxun Lu, Youtian Lin, Hao Zhu, Weiming Hu, Xun Cao, and Yao Yao. Stag4d: Spatial-temporal anchored generative 4d gaussians. In *European Conference on Computer Vision*, pages 163–179. Springer, 2024. [1](#), [2](#), [3](#), [5](#), [6](#)
- [52] Haiyu Zhang, Xinyuan Chen, Yaohui Wang, Xihui Liu, Yunhong Wang, and Yu Qiao. 4diffusion: Multi-view video diffusion model for 4d generation. *arXiv preprint arXiv:2405.20674*, 2024. [1](#), [6](#)
- [53] Zhenwei Zhu, Liying Yang, Ning Li, Chaohao Jiang, and Yanyan Liang. Umiformer: Mining the correlations between similar tokens for multi-view 3d reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18226–18235, 2023. [2](#)
- [54] Zi-Xin Zou, Zhipeng Yu, Yuan-Chen Guo, Yangguang Li, Ding Liang, Yan-Pei Cao, and Song-Hai Zhang. Triplane meets gaussian splatting: Fast and generalizable single-view 3d reconstruction with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10324–10335, 2024. [5](#)