

# SkeletonGaussian: Editable 4D Generation through Gaussian Skeletonization

Lifan Wu   Ruijie Zhu   Yubo Ai   Tianzhu Zhang  
University of Science and Technology of China

{wusar, ruijiezhu, erebai}@mail.ustc.edu.cn, tzzhang@ustc.edu.cn

## Abstract

4D generation has made remarkable progress in synthesizing dynamic 3D objects from input text, images, or videos. However, existing methods often represent motion as an implicit deformation field, which limits direct control and editability. To address this, we propose SkeletonGaussian, a novel framework for generating editable, dynamic 3D Gaussians from monocular video input. Our approach introduces a hierarchical, articulated representation that decomposes motion into sparse, rigid motion explicitly driven by a skeleton and fine-grained, non-rigid motion. Concretely, we extract a robust skeleton and drive rigid motion via linear blend skinning, followed by a hexplane-based refinement for non-rigid deformations—enhancing interpretability and editability. Experimental results show that SkeletonGaussian surpasses existing methods in generation quality while enabling intuitive motion editing, establishing a new paradigm for editable 4D generation.

*Keywords: 4D Generation, Gaussian Splatting, Motion Editing, Skeleton Modeling, Dynamic 3D*

## 1. Introduction

Dynamic 3D generation, also referred to as 4D generation, aims to create dynamic 3D objects from input text, images, or videos. It has become a prominent research area, expanding creative possibilities in fields such as animation, game design, autonomous driving, and film production. In this paper, we focus on generating editable dynamic 3D Gaussian models [14, 53, 60] from monocular video input.

Recent advancements in text-to-3D generation [22, 38, 51, 63] and image-to-3D synthesis [23, 26, 39, 48, 49] have enhanced the creation of diverse 3D objects. Building upon these developments, novel techniques [14, 53, 60] have emerged in the field of 4D generation. These methods leverage Score Distillation Sampling (SDS) loss, derived from diffusion model priors [57, 43, 27, 19], to optimize 4D object representation models. Depending on the type of 4D model used, existing methods can be categorized into three main classes: dynamic Neural Radiance Field (NeRF)

generation [13, 32, 64, 37, 34, 35], dynamic 3D Gaussian generation [33], and dynamic mesh generation [19].

Despite these advances, current 4D object representation methods typically model motion as an implicit deformation field [3], which limits direct control and editability. Editing deformation fields [3] in 4D models often requires retraining the deformation field, making the process time-consuming and lacking real-time feedback. Moreover, the parameter requirements of the deformation method grow quadratically with time, making it challenging to apply this motion modeling approach to long-duration sequences. Additionally, implicit deformation representations are difficult to convert into standard skeleton or pose data, which obstructs seamless integration with widely used animation tools and pipelines (e.g., Blender [2]). Together, these limitations hinder the adoption of practical motion generation workflows.

To tackle these challenges, we aim to develop a high-quality 4D generation workflow that not only produces superior 4D results but also facilitates real-time motion editing. Inspired by recent advances in human reconstruction [36, 40, 11, 15], which integrate the SMPL model [28] into 4D Gaussian modeling, we introduce **SkeletonGaussian**. SkeletonGaussian is an innovative framework for editable 4D generation through Gaussian skeletonization. This framework introduces a lightweight, hierarchical articulated motion representation technique that captures motion details across multiple levels. Therefore, it enables efficient and high-quality 4D generation, while providing flexible editing capabilities.

SkeletonGaussian integrates linear blend skinning (LBS) and skeleton-driven articulated motion representations into 4D generation tasks. It decomposes object motion into two components: sparse rigid deformation, driven by the skeleton, and fine non-rigid deformation, which captures intricate motion details such as wrinkles in clothing and skin. Our 4D generation pipeline consists of three stages: static 3D Gaussian generation, rigid motion modeling, and non-rigid motion refinement. We adopt UniRig [66] as the default skeleton extractor for robust, category-agnostic rigging, while using Coverage Axis++ as an ablation baseline. By leveraging hierarchical motion structures, SkeletonGaussian effectively captures complex motion dynam-

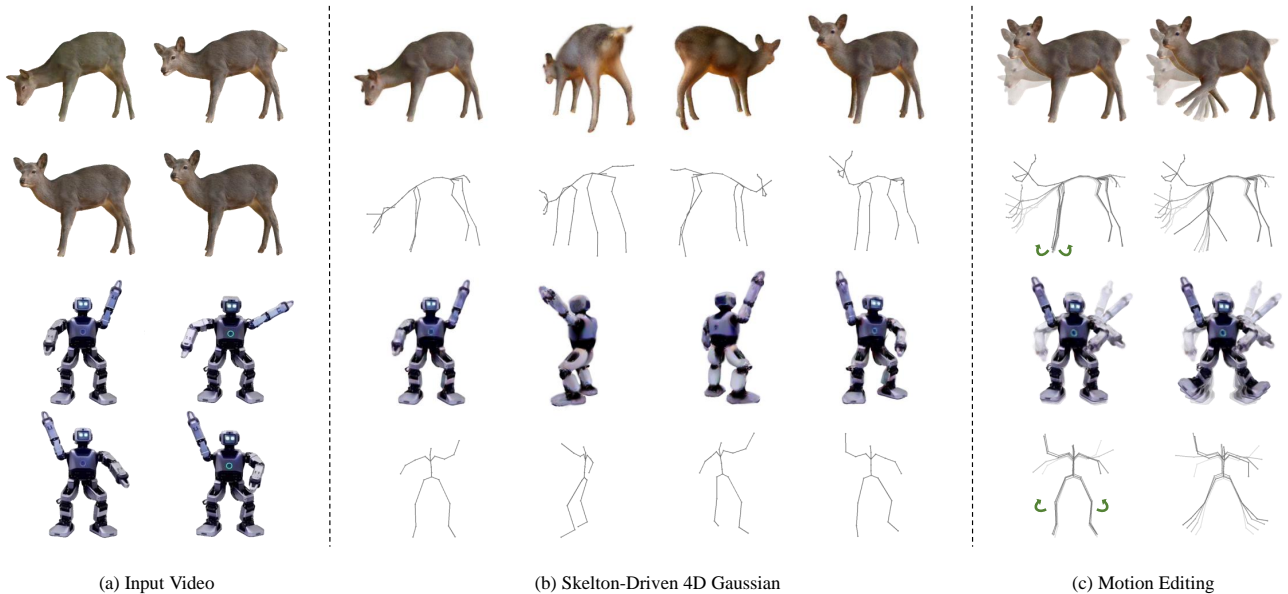


Figure 1: Given (a) an input monocular video, we propose a novel 4D generation method **SkeletonGaussian** which uses (b) a skeleton to drive the motion of 4D Gaussian model. SkeletonGaussian enables (c) direct motion editing through the skeleton’s explicit motion representation, allowing users to adjust skeleton poses to modify the motion of the objects directly.

ics, particularly in scenarios involving substantial transformations and intricate deformations. Moreover, SkeletonGaussian enables users to directly modify motion by editing the skeleton. It seamlessly integrates into existing 3D animation workflows, allowing for real-time motion adjustments without the need for computationally expensive optimization. The skeletal structure encodes the object’s physical topology, eliminating the need for auxiliary constraints, such as ARAP loss. Meanwhile, the explicit skeleton deformation method is highly parameter-efficient. The number of learnable pose parameters grows linearly with joints and time ( $\mathcal{O}(B \times T)$ ), reducing memory and training time compared to dense deformation fields. To validate the effectiveness of our method, we conduct quantitative and qualitative experiments using the Consistent4D [13] dataset. Our contributions are summarized as follows:

- We propose **SkeletonGaussian**, a skeleton-driven dynamic 3D Gaussian framework for motion modeling in generative tasks. By employing a hierarchical motion representation, SkeletonGaussian enhances motion fidelity while offering interpretable and editable pose controls. In contrast to dense deformation fields, the skeleton-based pose parameterization is more parameter-efficient, thereby reducing both storage demands and training times.
- Our explicit skeleton-based representation enables direct, real-time motion editing through the manipula-

tion of skeletal poses. The generated motions can be exported in standard skeleton and pose formats, ensuring seamless integration with animation pipelines such as Blender [2].

## 2. Related Work

**Skeleton-Based Motion Representations.** Skeleton-based motion representation is crucial in computer vision and graphics due to its manipulability and ability to model detailed object motion. It is extensively used in computer graphics, animation generation, and pose estimation. Linear Blend Skinning (LBS) is a commonly used technique for animating 3D models by applying transformations to a hierarchical skeleton, where the movement of joints influences the deformation of the model’s surface, enabling realistic motion and posing [16]. LBS is extensively used in contemporary 3D animation and motion modeling. In practical applications, models such as the Skinned Multi-Person Linear Model (SMPL) [28] employ LBS to combine joint articulation with skin meshes, resulting in realistic human motion and deformation. Similarly, FLAME [18] extends this method to facial animation, while SMAL [71] adapts it for animal modeling, demonstrating its versatility across different specialized fields.

**3D Skeleton Generation.** The extraction of skeletons from 3D representations, such as meshes or point clouds, is a well-established study area. Traditional methods relied on hand-crafted rules to extract geometric features. Techniques

such as Laplacian contraction [4] reduce point clouds to their topological structures, facilitating the extraction of key joints and skeletons. Offline approaches [31, 6, 52, 54, 17] have further refined this process. Recent methods [59, 21] employ deep neural networks to predict curve-based skeletons. In our system, we adopt **UniRig** [66] as the default skeleton extractor due to its generalizable rigging prior across categories, while also evaluating **Coverage Axis++** [52] as a baseline in our ablations (section 4.3). Both extractors are integrated into a unified pipeline with consistent axis correction and scale normalization.

**3D Deformation.** In 3D modeling, deformation techniques incorporate deformation fields into static 3D models. These techniques can be classified based on the model’s 3D representation: (1) **Mesh Deformation.** Classical mesh-based methods, such as Laplacian coordinates [25, 46] and cage-based techniques [61], focus on preserving geometric details during transformations, making them suitable for static objects. (2) **NeRF-Based Deformation.** Recent advancements in dynamic NeRF reconstruction utilize plane decomposition and 4D grids [3, 7] to achieve dynamic scene reconstruction by deforming canonical NeRF. (3) **3D Gaussian Deformation.** Recent advancements in 3D Gaussian splatting [14] significantly accelerate rendering. Dynamic 3D Gaussian methods [30, 53, 60] leverage deformation fields [3] to model 3D Gaussian motion. Some approaches, such as SC-GS [12] and BAGS [69], use sparse control points to represent 3D deformation. However, these methods rely on hexplane and MLP-based techniques to implicitly model control point movement and deformation, whereas our method explicitly models motion using a skeleton and joint pose parameterization. Recent techniques for 3D Gaussian-based dynamic human motion [36, 40, 11, 15] combine rigid skeletal skinning and non-rigid deformations for precise motion modeling. Our approach draws inspiration from these works, employing a skeleton-based deformation in 3D Gaussian rendering to provide an intuitive and efficient method for editable 4D generation.

**3D and 4D Generation.** In 3D generation, DreamFusion [38] first introduced score distillation sampling (SDS) [38, 51] loss, which optimizes NeRF [32] to produce high-quality 3D models. Building on advancements in 3D Gaussian techniques [14], DreamGaussian [49] leverages Gaussian splatting for 3D generation, significantly improving speed and performance. In 4D generation research, traditional motion generation methods [9, 47] are often limited to specific characters or datasets, reducing their applicability across diverse objects. Recently, diffusion model-based approaches [44, 1, 24, 42, 70, 13, 65, 62] address these limitations by integrating SDS loss into 4D framework, enabling more versatile and generalized motion generation. Approaches such as SC4D [55] introduce control points for motion transfer, enhancing editing flexibility in

4D generation. Additionally, Diffusion4D [20] and Stable Video 4D [56] achieve spatiotemporal consistency through specialized attention layers. STAG4D [65] initializes multi-view images anchored to input video frames, which are then used for multi-view SDS computation. Building on these advancements, our work introduces a novel skeleton-driven 3D Gaussian deformation framework for 4D generation tasks, offering a more intuitive and efficient approach to motion modeling and editing.

### 3. Method

As illustrated in figure 2, the 4D object generation pipeline of SkeletonGaussian consists of three stages: (1) **Static 3D Object Generation and Skeleton Extraction** (section 3.1): A static 3D object is initially generated using a 3D Gaussian generation method [14, 49]. Subsequently, an inherent skeletal structure is constructed for the 3D Gaussian model. (2) **Rigid Motion Modeling** (section 3.2): To capture the primary rigid motion of the object, a skeletal skinning network is employed, and the motion trajectories of each skeletal point are calculated using Forward Kinematics. Linear Blend Skinning (LBS) is then applied to deform the 3D Gaussian model according to these skeletal trajectories. During this stage, the **skeleton poses** are optimized to match the reference video sequences. (3) **Non-Rigid Motion Modeling** (section 3.3): Fine-grained non-rigid motions are represented through a hexplane [3] and a deformation MLP. At this stage, the skeletal skinning network remains frozen, while only the 3D Gaussian and the hexplane deformation field are trained to capture finer deformations.

Through these three stages, SkeletonGaussian generates a high-quality 4D object comprising a 3D Gaussian model, a skeletal pose representing the object’s rigid motion, and a fine-grained deformation field, achieving high-quality dynamic 3D object generation. The following sections detail each of these training stages.

#### 3.1. Static 3D Gaussian and Skeleton Generation

To generate a static 3D Gaussian and its corresponding skeletal structure, we select the middle frame of the video as the reference frame for constructing the initial static 3D Gaussian model  $\mathcal{G}_c$  in canonical space:

$$\mathcal{G}_c = \{\mathbf{p}_c, \mathbf{q}_c, \mathbf{s}, \sigma, \mathbf{c}\}, \quad (1)$$

where  $\mathbf{p}_c$ ,  $\mathbf{q}_c$ ,  $\mathbf{s}$ ,  $\sigma$ , and  $\mathbf{c}$  represent the position, quaternions, scale, opacity, and spherical harmonics coefficients of the 3D Gaussian in canonical space, respectively. The middle frame is chosen as the static reference because it minimizes the motion discrepancy with all other frames, thereby reducing task complexity. The static 3D Gaussian model is trained using both the multi-view SDS loss and the

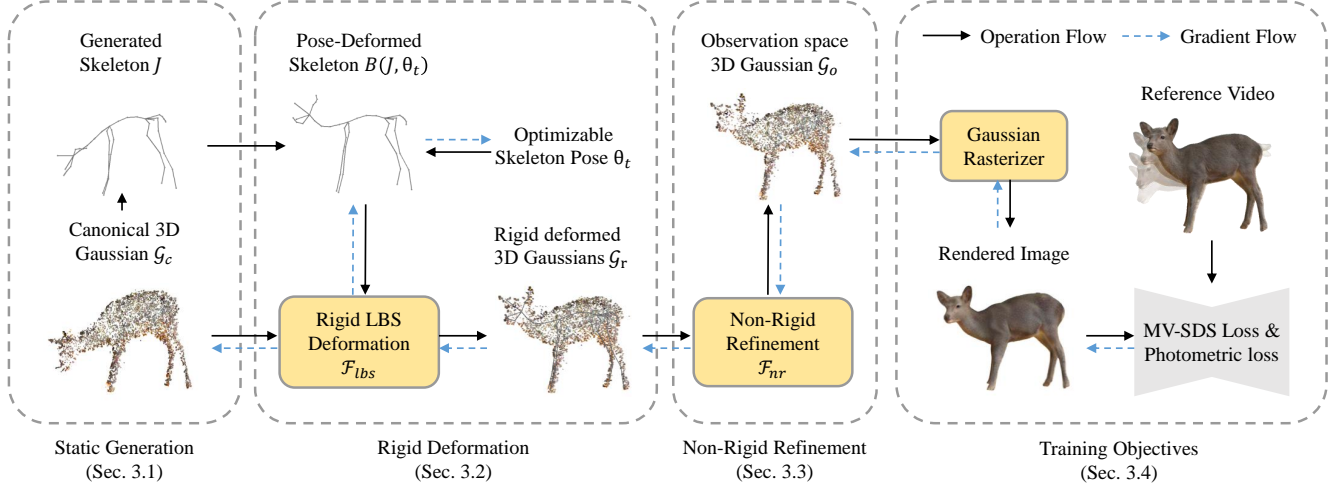


Figure 2: Pipeline of the SkeletonGaussian framework for 4D object generation, divided into three stages: (1) **Static 3D Object Generation and Skeleton Extraction**: Starting from a frame at the video’s midpoint, a static 3D Gaussian model  $\mathcal{G}_c$  (section 3.1) is generated in canonical space, from which an inherent skeletal structure is subsequently extracted. (2) **Rigid Motion Modeling**: Using LBS, rigid deformations  $\mathcal{F}_{lbs}$  (section 3.2) under various poses  $\theta_t$  are applied to rigidly deform  $\mathcal{G}_c$  into  $\mathcal{G}_r$ . During this stage, the skeleton poses  $\theta_t$  are optimized. (3) **Non-Rigid Motion Modeling**: To capture fine-grained deformations, a deformation field  $\mathcal{F}_{nr}$  (section 3.3) refines the motion of the rigidly deformed 3D Gaussian  $\mathcal{G}_r$ , transforming it into the observation space Gaussian  $\mathcal{G}_o$ .  $\mathcal{F}_{nr}$  comprises a hexplane [3] and an MLP. All three stages share the same **Training Objectives** (section 3.4). A differentiable Gaussian rasterizer renders images of the observation space 3D Gaussian  $\mathcal{G}_o$  from multiple viewpoints, comparing them to the reference video with photometric and MV-SDS losses for backpropagation.

photometric consistency loss. Further details are provided in section 3.4.

**Skeleton Generation.** To generate the kinematic tree structure of the skeleton joints  $\mathbf{J}$ , the mesh structure of the static object is first extracted from the static 3D Gaussian using occupation fields and the marching cubes algorithm [29]. We adopt a robust rigging pipeline built on UniRig [66] (default in our system), which predicts joint candidates and their connectivity to form an articulated skeleton for general objects. In practice, we support two invocation modes to enhance reproducibility and portability across environments: (1) an *internal* Python inference path, and (2) an *external* script path that caches results on disk and can be launched from sandboxed environments. Prior to building forward kinematics (FK), we apply standard preprocessing to ensure consistent coordinate conventions across extractors.

Based on the extracted skeletal points, we construct a kinematic tree by computing a Minimum Spanning Tree (MST) over candidate joints, and we preserve joint identifiers when available from the extractor to maintain consistency with rigging conventions. The kinematic tree provides a compact structural abstraction of the 3D Gaussian and serves as the control scaffold for subsequent motion generation.

### 3.2. 3D Gaussian Rigid Deformation

To model the primary motion of the 3D Gaussian, we use LBS to apply rigid deformation to the canonical 3D Gaussian  $\mathcal{G}_c$ , denoted as  $\mathcal{F}_{lbs}$ . Let  $\mathbf{J} = \{\mathbf{J}_b\}_{b=1}^B$  represent the set of static joint positions of the skeleton, and let  $\theta_t$  represent the skeleton’s pose at a specific time  $t$ . Under these conditions, the corresponding rigidly deformed 3D Gaussian  $\mathcal{G}_r$  is computed as follows:

$$\mathcal{G}_r = \mathcal{F}_{lbs}(\mathcal{G}_c; \mathbf{J}, \theta_t). \quad (2)$$

Note that while LBS results in non-rigid deformation, we term this “rigid deformation” to highlight it is driven by rigid skeletal joints, distinguishing it from the subsequent non-rigid refinement.

**Rigid Position and Rotation Transform.** The deformed 3D Gaussian point  $\mathcal{G}_r^i$  is computed by applying a transformation matrix  $\mathbf{T}_i$  to the 3D Gaussian point  $\mathcal{G}_c^i$ .  $\mathbf{T}_i$  is a weighted sum of the transformation matrices  $\mathbf{B}_k(\mathbf{J}, \theta_t)$  corresponding to the nearest  $K$  skeletal joints, with weights  $w_{k,i}$ :

$$\mathbf{T}_i = \sum_{k=1}^K w_{k,i} \mathbf{B}_k(\mathbf{J}, \theta_t). \quad (3)$$

The transformed position  $\mathbf{p}_r^i$  and the rotation  $\mathbf{q}_r^i$  of the de-

formed 3D Gaussian  $\mathcal{G}_c^i$  are calculated as follows:

$$\mathbf{p}_r^i = \mathbf{T}_i \mathbf{p}_c^i + \mathbf{t}_o, \quad \mathbf{q}_r^i = \mathbf{T}_{i(1:3,1:3)} \cdot \mathbf{q}_c^i, \quad (4)$$

where  $\mathbf{T}_{1:3,1:3}$  refers to the rotational component extracted from the transformation matrix  $\mathbf{T}_i$ . The term  $\mathbf{t}_o$  denotes the global translation of the root joint at time  $t$ .

**Forward Kinematics.** To compute skeletal animations and joint motions, we use the forward kinematics approach. Forward kinematics determines each joint’s transformation by recursively accumulating the transformations of its ancestor joints. It relies on a hierarchical skeleton tree structure, where the transformation matrix  $\mathbf{B}_k(\mathbf{J}, \theta_t)$  for each joint  $k$  is obtained by multiplying the local transformations  $\theta_t^j$  of all its ancestor joints  $j \in A(k)$ :

$$\mathbf{B}_k(\mathbf{J}, \theta_t) = \prod_{j \in A(k)} \theta_{t,j}. \quad (5)$$

**Skeletal Skinning Weights.** To compute the skinning weights  $w_{k,i}$  for each Gaussian point  $i$  relative to its surrounding skeleton joints, we apply the K-nearest neighbors (KNN) algorithm to identify the  $K$  nearest points. The weights are determined using inverse distance weighting, where the weight is inversely proportional to the distance  $d_{k,i}$  between point  $i$  and skeleton joint  $k$ :

$$w_{k,i} = \frac{\frac{1}{d_{k,i}}}{\sum_{k=1}^K \frac{1}{d_{k,i}}}. \quad (6)$$

We employ fixed inverse-distance KNN weights due to their simplicity and lack of training requirements. In future work, we plan to explore learning-based skinning weight fields to further improve deformation quality.

**Skeletal Pose Smoothness.** The skeletal pose is represented by a tensor  $\theta \in \mathbb{R}^{T \times B \times 4}$ , where  $T$  is the number of frames,  $B$  is the number of joints, and each entry  $\theta_{t,k}$  represents a 4D quaternion encoding the rotation of the  $k$ -th joint at the  $t$ -th frame. Additionally, we introduce a variable  $\mathbf{t}_o$  to record the global translation of the root joint. Directly optimizing skeleton poses can lead to overfitting, causing the model to capture noise from the training data and produce jitter in the generated motion. To mitigate this problem, we employ window smoothing to the skeleton poses, which uses a sliding window of size  $2w + 1$  during training to smooth the motion across  $w$  neighboring frames. For the local skeletal pose  $\theta_t$  at each time step  $t$ , we compute the average pose  $\bar{\theta}_t$  by averaging across frame  $t$  and its neighbors. This smoothed value  $\bar{\theta}_t$  is then used as input for LBS deformations. For simplicity, we use  $\theta_t$  instead of  $\bar{\theta}_t$  elsewhere. The formula is:

$$\bar{\theta}_t = \frac{1}{2w + 1} \sum_{i=-w}^w \theta_{t+i}. \quad (7)$$

### 3.3. Non-Rigid Refinement

LBS effectively captures global object motion but struggles to represent detailed motions, such as clothing wrinkles, due to the limited number of skeletal joints. To overcome this limitation, we propose a non-rigid deformation method to capture these detailed motions. Our approach employs a hexplane-based 4D representation to refine the motion of static 3D Gaussians. Specifically, we integrate a hexplane and an MLP to regress displacement, rotation, and scale changes of the 3D Gaussian. The non-rigid deformation function  $\mathcal{F}_{nr}$  which transforms the rigidly deformed 3D Gaussian  $\mathcal{G}_r$  into the observation space 3D Gaussian  $\mathcal{G}_o$  is given by:

$$\mathcal{G}_o = \mathcal{F}_{nr}(\mathcal{G}_r). \quad (8)$$

The observation space 3D Gaussian  $\mathcal{G}_o$  is then rendered through Gaussian rasterization. During this refinement field training stage, the skeletal skinning network is frozen, and only the 3D Gaussian and hexplane deformation field are trained to capture fine-grained deformations.

### 3.4. Training Objectives

Our objective is to generate a 4D Gaussian representation of the target object from an input video sequence by optimizing the static 3D Gaussian model, the skeleton poses  $\theta_t$ , and the refinement field parameters  $\mathcal{F}_{nr}$ . We begin by applying the multi-view diffusion model Zero123++ [43] to generate multi-view sequences  $I_{\text{anchor}}^t$  from the video inputs. These sequences act as spatiotemporal anchors, which are later used to compute the MV-SDS loss. The 3D Gaussian is then projected onto the screen to produce output images, which are compared with the anchor images  $I_{\text{anchor}}^t$  using the multi-view Score Distillation Sampling (SDS) loss  $\mathcal{L}_{MV-SDS}$  from Zero123 [27]. In addition, the loss function includes a reconstruction loss  $\mathcal{L}_{rec}$  and a foreground masking loss  $\mathcal{L}_{mask}$  between the reference image  $I_t^{ref}$  and the front-view rendered image. A regularization loss  $\mathcal{L}_{reg}$  is also applied to the deformation field  $\mathcal{F}_{nr}$  to enforce temporal smoothness in the motion. For a detailed explanation of the loss function, please refer to the Appendix section 7. The final optimization objective is given by:

$$\mathcal{L} = \mathcal{L}_{MV-SDS} + \lambda_1 \mathcal{L}_{rec} + \lambda_2 \mathcal{L}_{mask} + \lambda_3 \mathcal{L}_{reg}. \quad (9)$$

### 3.5. Generated Motion Editing

SkeletonGaussian provides an efficient approach for editing generated motion through its sparse, explicit skeleton-based representation. Users can intuitively adjust the motion of skeletal points by modifying poses at specific time steps, thereby altering the entire motion trajectory. We develop a GUI that simplifies the process of motion editing. This method also aligns with current motion




Figure 3: Visualizing 4D Object Motion with Skeleton Poses. We present generated 4D object motion and its corresponding skeleton poses, where the viewpoint rotates from left to right, and time progresses linearly from left to right.




Figure 4: Editing Generated Motion. We visualize the generated motion (top) and edited motion sequence (bottom). Users can directly adjust the skeleton poses of specific joints at different times to edit the object’s motion.

modeling techniques in computer graphics, enabling users to modify skeletal movements in popular 3D editors such as Blender [2]. Additionally, the hierarchical structure of the skeleton tree facilitates hierarchical motion editing, where adjustments to a parent node automatically propagate to its child nodes. Motion editing is illustrated in figure 4.

## 4. Experiments

### 4.1. Experiment Setup

**Implementation Details.** (1) **Static Stage:** We generate six anchor view videos  $I_t^i$  ( $i \in \{1...6\}$ ) using Zero123++ [43] from the input monocular video  $I_t^{ref}$ . The SDS loss is computed using Zero-1-to-3 [27]. 10000 3D Gaussian points are randomly initialized within a spherical canonical space. This stage is trained for 1500 steps to produce a static 3D Gaussian. Subsequently, a skeleton is generated using UniRig [66] (default). We support both an internal Python path and an external cached script path for cross-environment execution. (2) **Skeleton Training Stage:** Skeleton poses are trained for 2500 steps. A smoothing

window of size three is applied to the skeleton poses. (3) **Non-Rigid Motion Refinement:** A hexplane and a deformation MLP are trained for 7000 steps to capture the fine-grained motion. The detailed implementation and hyperparameters are provided in Appendix section 8. The entire training process takes approximately 1 hour on an RTX 3090 GPU, and the rendering process can be performed at 150 FPS in real time.

**Evaluation Dataset.** To fairly evaluate our method against the baselines, we use the Consistent4D dataset [13], which includes 4D animation assets from Sketchfab [45] for further animation assessment. The dataset comprises 12 synthetic and 12 real-world videos, each captured with a static vertically aligned camera focused on dynamic objects. Each video contains 32 frames over approximately 2 seconds.

**Evaluation Metrics.** We evaluate the quality of generated 4D videos based on their alignment with reference videos, spatio-temporal consistency, and motion fidelity. For each test object and method, we use a frontal view video as input to generate a corresponding dynamic 3D model, rendering four videos from azimuth angles of  $75^\circ$ ,  $15^\circ$ ,  $105^\circ$ , and  $195^\circ$  at a  $0^\circ$  elevation. These rendered videos are compared with the ground-truth videos in the dataset to evaluate the generation quality. Our evaluation metrics include CLIP [41], LPIPS [68], and FVD [50] for Video-to-4D evaluation. CLIP and LPIPS evaluate the semantic and perceptual similarities between generated and real images, while FVD computes frame quality and temporal consistency. Since DreamGaussian4D generates videos with only 16 frames, we use the FVD-16 score, which computes the FVD based on the first 16 frames.

**Baselines.** We compare SkeletonGaussian with several recent 4D generation methods capable of generating multi-view videos from a single-view video input, including Consistent4D [13], STAG4D [65], 4DGen [62], and DreamGaussian4D [42]. All baselines are evaluated on the Consis-




Figure 5: Qualitative Comparisons. We compare our method with STAG4D [65] and DreamGaussian4D [42]. For each instance, we render two viewpoints at two time steps. We also visualize the skeleton poses of SkeletonGaussian.

Table 1: **Quantitative Evaluation of 4D Generation on the Consistent4D Dataset.** SkeletonGaussian outperforms in both image quality and video frame consistency.

Method	CLIP $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$
Consistent4D [13]	0.877	0.161	1518.5
DreamGaussian4D [42]	0.913	0.143	994.11
4DGen(16 frames) [62]	0.909	0.137	913.10
STAG4D [65]	0.909	0.126	992.2
<b>SkeletonGaussian (Ours)</b>	<b>0.923</b>	<b>0.125</b>	<b>847.8</b>

tent4D dataset using their official code and configurations. Quantitative and qualitative comparisons are presented in figure 5 and table 1.

## 4.2. Comparisons

**Quantitative Comparisons.** As shown in table 1, our method outperforms STAG4D, DreamGaussian4D, and 4DGen on the Consistent4D [13] dataset in terms of reference view alignment (LPIPS, CLIP), indicating that our approach generates more realistic images. Furthermore, our

method achieves the lowest FVD score, demonstrating that our generated videos exhibit fewer temporal artifacts and better match real-world footage. These results highlight the effectiveness of SkeletonGaussian.

**Qualitative Comparison.** We compare the 4D outputs generated by our method with STAG4D and DreamGaussian4D in figure 5. For each video, we render the 4D results at two timestamps and from two perspectives: one from the front and the other from the back. Additionally, we visualize the skeletons of our method. Our approach achieves high-fidelity reconstruction with stable geometry and consistent texture. The results maintain fine details across frames, demonstrating robustness in both spatial and temporal aspects.

**User Study.** To validate our method, we conduct user studies to evaluate multi-view video synthesis and 4D outputs. We select 20 real-world and synthetic videos from the Objaverse [5] and Consistent4D [13] datasets. Participants compare results from three methods (SkeletonGaussian and three baselines) based on a novel camera view, choosing the most stable, realistic, and reference-like video. SkeletonGaussian is preferred by 32.5%, followed by STAG4D (27.5%), 4DGen (22.5%), and DreamGaussian4D (17.5%).



Figure 6: Qualitative evaluations of the ablation study. We visualize the skeleton poses and the objects at different time steps.

### 4.3. Ablation Studies

In this section, we evaluate the effectiveness of various motion modeling methods by analyzing the quality of the generated 4D Gaussians through a series of ablation studies. We assess the quality, memory requirements, and training time of different motion modeling approaches, providing quantitative comparisons in [table 2](#) and qualitative comparisons in [figure 6](#).

**Rigid-only (LBS).** Using only rigid LBS preserves articulated structure but underfits fine non-rigid motion (see [figure 6](#)). Thanks to the compact pose parameterization that scales as  $\mathcal{O}(B \times 3 \times T)$ , it achieves the smallest deformation-module VRAM and the shortest training time in [table 2](#). In our setting with  $B \approx 30$  and  $T = 32$ , this amounts to  $30 \times 32 \times 3$  scalars, i.e., only storing joint rotation angles; the rigid stage takes about 1000 steps ( 0.2 h).

**Non-rigid-only (HexPlane+MLP).** The non-rigid field captures detailed deformations but lacks an articulated prior, reducing temporal stability. Its parameter count scales with grid/plane resolutions and MLP widths, leading to the largest deformation-module VRAM and the longest training time; empirically the memory cost grows roughly  $\mathcal{O}(T^2)$  with sequence length, making long sequences hard to optimize. On Consistent4D with  $T = 32$ , we measure the deformation module VRAM at 136.40 MiB, and the deformation stage takes about 8000 steps ( 1.5 h). Quantitatively, it can slightly improve per-frame fidelity (LPIPS/CLIP) but still trails the Full model on overall temporal quality.

The **Full (Rigid+Non-rigid)** model combines both stages; its deformation-module VRAM is close to the non-rigid variant with a negligible skeleton overhead, and the total training time is roughly the sum of the two stages ( 1.7 h).

**Pose Smoothness.** The pose-smoothness regularizer improves the temporal continuity of articulated poses, yielding smoother motion and fewer jitters. Quantitatively, removing it degrades LPIPS/FVD, as summarized in [table 4](#).

**Skeleton Extractor (Coverage Axis++ vs UniRig).** We also ablate the skeleton extractor under the same pipeline. Coverage Axis++ [52] selects skeletal points via coverage heuristics and connects them via a Minimum Spanning Tree, whereas UniRig [66] provides a stronger, category-agnostic rigging prior with joint proposals and connectivity that we further regularize via FK. Replacing UniRig with Coverage Axis++ (row “Using Coverage Axis++” in [table 4](#)) weakens temporal stability and increases artifacts. Qualitatively, [figure 6](#) shows more stable and semantically aligned joints with UniRig, which improves control and reduces topological errors.

**Impact of Initial Frame Selection.** We further analyze how the choice of the initial reference frame affects the generation results. Our method typically uses the first frame as the canonical reference. Experiments indicate that selecting a frame with clear visibility and a neutral pose contributes to a more accurate canonical 3D Gaussian initialization. However, our skeleton-driven deformation mechanism provides strong geometric priors, making the system relatively robust to the initial frame selection. Even when initialized from frames with partial self-occlusions, the method can recover plausible motion dynamics through the subsequent rigid and non-rigid optimization stages. Quantitative results are shown in [table 3](#).

Table 2: Quantitative ablation on motion modeling. Metrics include CLIP/LPIPS/FVD and efficiency (VRAM in MiB and training time in minutes) measured for the deformation module only, excluding the static 3D Gaussian. We compare Rigid-only (LBS), Non-rigid-only (HexPlane+MLP), and Full (Rigid+Non-rigid).

Method	CLIP $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$	VRAM (MiB) $\downarrow$	Train Time (min) $\downarrow$
Rigid-only (LBS)	0.901	0.135	1012.6	0.01	12
Non-rigid-only (HexPlane+MLP)	0.909	0.126	992.2	136.40	90
<b>Full (Rigid+Non-rigid)</b>	<b>0.923</b>	<b>0.125</b>	<b>847.8</b>	136.41	102

Table 3: Ablation study on the impact of initial frame selection. We compare selecting the first frame (Frame 0), the middle frame (Frame 15), and a random frame as the initialization for the 3D Gaussian field. The results show consistent performance across different initial frames.

Initial Frame Selection	CLIP $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$
Frame 0 (First Frame)	0.921	0.126	851.2
Frame 15 (Middle Frame)	<b>0.923</b>	<b>0.125</b>	<b>847.8</b>
Random Frame	0.919	0.128	858.5

Table 4: Compact ablations on pose smoothness and skeleton extractor. ‘‘Coverage Axis++’’ swaps UniRig in the Full (Rigid+Non-rigid) model; ‘‘w/o Pose Smoothness’’ drops the pose-smoothness regularizer; ‘‘Full’’ uses UniRig with pose smoothness.

Method	CLIP $\uparrow$	LPIPS $\downarrow$	FVD $\downarrow$
Using Coverage Axis++	0.918	0.128	890.4
w/o Pose Smoothness	0.906	0.131	1034.3
Full (Rigid+Non-rigid)	<b>0.923</b>	<b>0.125</b>	<b>847.8</b>

## 5. Discussion

**Limitations and Future Directions.** We observe that incorrect skeleton retrieval can degrade the quality of the generated results. Specifically, in some cases, severe topological errors in skeleton extraction can diminish the quality of the generated results. Additionally, there may be cases where objects do not have a clear skeleton structure, and our method performs poorly in these situations. The hexplane deformation field demonstrates error compensation capabilities for mild skeletal inaccuracies, helping to address this issue. Furthermore, we are developing an adaptive skeleton error-correcting mechanism that dynamically adjusts the skeleton structure during training. Please refer to Appendix section 9 for a detailed analysis of failure cases.

Currently, our method does not support multi-object motion, thus limiting its applicability in scenarios involving

multiple objects. Future work could address this limitation by incorporating independent skeletons for each object. Additionally, we are developing the integration of predefined skeleton templates, such as SMPL [28], to initialize the 3D Gaussian and skeleton structure using vertex and joint positions. We also successfully integrate the human pose estimation method ViTPose [58] into SkeletonGaussian to initialize the skeleton poses. This integration is expected to improve the accuracy and quality of 4D motion generation significantly. Furthermore, our approach can seamlessly integrate with skeleton-controlled video generation techniques, such as ControlNet [67, 8, 10]. Using 3D skeletons as conditional inputs for diffusion models in 2D images opens new possibilities for 4D generation. Additionally, enhanced skeletal control provides a novel representation of motion, which could be applied to motion-tracking tasks.

## 6. Conclusion

This paper introduces SkeletonGaussian, a framework for generating editable 4D Gaussian-based models from monocular video. By explicitly decomposing motion into rigid skeletal movements and fine-grained non-rigid details, this framework improves control and interpretability in 4D Gaussian modeling. SkeletonGaussian operates in three phases: constructing a static 3D Gaussian model, modeling rigid motion through skeletal LBS, and refining non-rigid motion using a hexplane-based deformation field. This hierarchical structure enables intuitive motion editing by adjusting skeleton poses and aligning seamlessly with standard animation workflows. Experimental results demonstrate that SkeletonGaussian delivers superior quality over existing methods, offering a new paradigm for editable 4D motion generation.

## References

- [1] S. Bahmani, I. Skorokhodov, V. Rong, G. Wetzstein, L. Guibas, P. Wonka, S. Tulyakov, J. J. Park, A. Tagliasacchi, and D. B. Lindell. 4d-fy: Text-to-4d generation using hybrid score distillation sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7996–8006, 2024. 3

- [2] Blender. Blender. <https://www.blender.org/>, 2024. Accessed: 2024-10-22. 1, 2, 6
- [3] A. Cao and J. 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. 1, 3, 4
- [4] J. Cao, A. Tagliasacchi, M. Olson, H. Zhang, and Z. Su. Point cloud skeletons via laplacian based contraction. In *2010 Shape Modeling International Conference*, pages 187–197. IEEE, 2010. 3
- [5] M. Deitke, D. Schwenk, J. Salvador, L. Weihs, O. Michel, E. VanderBilt, L. Schmidt, K. Ehsani, A. Kembhavi, and A. Farhadi. Objaverse: A universe of annotated 3d objects. *arXiv preprint arXiv:2212.08051*, 2022. 7
- [6] Z. Dou, C. Lin, R. Xu, L. Yang, S. Xin, T. Komura, and W. Wang. Coverage axis: Inner point selection for 3d shape skeletonization. In *Computer Graphics Forum*, volume 41, pages 419–432. Wiley Online Library, 2022. 3
- [7] S. Fridovich-Keil, G. Meanti, F. R. Warburg, B. Recht, and A. Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12479–12488, 2023. 3
- [8] 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. 9
- [9] F. G. Harvey, M. Yurick, D. Nowrouzezahrai, and C. Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020. 3
- [10] L. Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8153–8163, 2024. 9
- [11] L. Hu, H. Zhang, Y. Zhang, B. Zhou, B. Liu, S. Zhang, and L. Nie. Gaussianavatar: Towards realistic human avatar modeling from a single video via animatable 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 634–644, 2024. 1, 3
- [12] Y.-H. Huang, Y.-T. Sun, Z. Yang, X. Lyu, Y.-P. Cao, and X. Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4220–4230, 2024. 3
- [13] Y. Jiang, L. Zhang, J. Gao, W. Hu, and Y. Yao. Consistent4d: Consistent 360  $\{\backslash\text{deg}\}$  dynamic object generation from monocular video. *arXiv preprint arXiv:2311.02848*, 2023. 1, 2, 3, 6, 7
- [14] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 1, 3
- [15] M. Kocabas, J.-H. R. Chang, J. Gabriel, O. Tuzel, and A. Ranjan. Hugs: Human gaussian splats. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 505–515, 2024. 1, 3
- [16] J. P. Lewis, M. Corder, and N. Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 811–818, 2023. 2
- [17] P. Li, B. Wang, F. Sun, X. Guo, C. Zhang, and W. Wang. Q-mat: Computing medial axis transform by quadratic error minimization. *ACM Transactions on Graphics (TOG)*, 35(1):1–16, 2015. 3
- [18] T. Li, T. Bolkart, M. J. Black, H. Li, and J. Romero. Learning a model of facial shape and expression from 4d scans. *ACM Trans. Graph.*, 36(6):194–1, 2017. 2
- [19] Z. Li, Y. Chen, and P. Liu. Dreammesh4d: Video-to-4d generation with sparse-controlled gaussian-mesh hybrid representation. *arXiv preprint arXiv:2410.06756*, 2024. 1
- [20] H. Liang, Y. Yin, D. Xu, H. Liang, Z. Wang, K. N. Plataniotis, Y. Zhao, and Y. Wei. Diffusion4d: Fast spatial-temporal consistent 4d generation via video diffusion models. *arXiv preprint arXiv:2405.16645*, 2024. 3
- [21] C. Lin, C. Li, Y. Liu, N. Chen, Y.-K. Choi, and W. Wang. Point2skeleton: Learning skeletal representations from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4277–4286, 2021. 3
- [22] C.-H. Lin, J. Gao, L. Tang, T. Takikawa, X. Zeng, X. Huang, K. Kreis, S. Fidler, M.-Y. Liu, and T.-Y. Lin. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 300–309, 2023. 1
- [23] Y. Lin, H. Han, C. Gong, Z. Xu, Y. Zhang, and X. Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. *arXiv preprint arXiv:2309.17261*, 2023. 1
- [24] H. Ling, S. W. Kim, A. Torralba, S. Fidler, and K. Kreis. Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8576–8588, 2024. 3
- [25] Y. Lipman, O. Sorkine, M. Alexa, D. Cohen-Or, D. Levin, C. Rössl, and H.-P. Seidel. Laplacian framework for interactive mesh editing. *International Journal of Shape Modeling*, 11(01):43–61, 2005. 3
- [26] M. Liu, C. Xu, H. Jin, L. Chen, M. Varma T, Z. Xu, and H. Su. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *Advances in Neural Information Processing Systems*, 36, 2024. 1
- [27] R. Liu, R. Wu, B. Van Hoorick, P. Tokmakov, S. Zakharov, and C. Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 1, 5, 6
- [28] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. Smpl: A skinned multi-person linear model. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 851–866, 2023. 1, 2, 9
- [29] W. E. Lorensen and H. E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Seminal graphics: pioneering efforts that shaped the field*, pages 347–353, 1998. 4

- [30] J. Luiten, G. Kopanas, B. Leibe, and D. Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 3
- [31] L. Meyer, A. Gilson, O. Scholz, and M. Stamminger. Cherypicker: Semantic skeletonization and topological reconstruction of cherry trees, 2023. 3
- [32] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 3
- [33] Z. Pan, Z. Yang, X. Zhu, and L. Zhang. Fast dynamic 3d object generation from a single-view video. *arXiv preprint arXiv:2401.08742*, 2024. 1
- [34] K. Park, U. Sinha, J. T. Barron, S. Bouaziz, D. B. Goldman, S. M. Seitz, and R. Martin-Brualla. Nerfies: Deformable neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5865–5874, 2021. 1
- [35] K. Park, U. Sinha, P. Hedman, J. T. Barron, S. Bouaziz, D. B. Goldman, R. Martin-Brualla, and S. M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *arXiv preprint arXiv:2106.13228*, 2021. 1
- [36] C. Pokhariya, I. N. Shah, A. Xing, Z. Li, K. Chen, A. Sharma, and S. Sridhar. Manus: Markerless grasp capture using articulated 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2197–2208, 2024. 1, 3
- [37] G. Pons-Moll, F. Moreno-Noguer, E. Corona, and A. Pumarola. D-nerf: Neural radiance fields for dynamic scenes. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2021. 1
- [38] B. Poole, A. Jain, J. T. Barron, and B. Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv preprint arXiv:2209.14988*, 2022. 1, 3
- [39] G. Qian, J. Mai, A. Hamdi, J. Ren, A. Siarohin, B. Li, H.-Y. Lee, I. Skorokhodov, P. Wonka, S. Tulyakov, et al. Magic123: One image to high-quality 3d object generation using both 2d and 3d diffusion priors. *arXiv preprint arXiv:2306.17843*, 2023. 1
- [40] Z. Qian, S. Wang, M. Mihajlovic, A. Geiger, and S. Tang. 3dgs-avatar: Animatable avatars via deformable 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5020–5030, 2024. 1, 3
- [41] 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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 6
- [42] J. Ren, L. Pan, J. Tang, C. Zhang, A. Cao, G. Zeng, and Z. Liu. Dreamgaussian4d: Generative 4d gaussian splatting. *arXiv preprint arXiv:2312.17142*, 2023. 3, 6, 7
- [43] 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. 1, 5, 6
- [44] U. Singer, S. Sheynin, A. Polyak, O. Ashual, I. Makarov, F. Kokkinos, N. Goyal, A. Vedaldi, D. Parikh, J. Johnson, et al. Text-to-4d dynamic scene generation. *arXiv preprint arXiv:2301.11280*, 2023. 3
- [45] Sketchfab. Sketchfab 3d models. <https://sketchfab.com/3d-models>, 2024. Accessed: 2024-10-22. 6
- [46] O. Sorkine and M. Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, volume 4, pages 109–116. Citeseer, 2007. 3
- [47] P. Starke, S. Starke, T. Komura, and F. Steinicke. Motion in-betweening with phase manifolds. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(3):1–17, 2023. 3
- [48] J. Sun, B. Zhang, R. Shao, L. Wang, W. Liu, Z. Xie, and Y. Liu. Dreamcraft3d: Hierarchical 3d generation with bootstrapped diffusion prior. *arXiv preprint arXiv:2310.16818*, 2023. 1
- [49] J. Tang, J. Ren, H. Zhou, Z. Liu, and G. Zeng. Dreamgaussian: Generative gaussian splatting for efficient 3d content creation. *arXiv preprint arXiv:2309.16653*, 2023. 1, 3
- [50] 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. 6
- [51] H. Wang, X. Du, J. Li, R. A. Yeh, and G. Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12619–12629, 2023. 1, 3
- [52] Z. Wang, Z. Dou, R. Xu, C. Lin, Y. Liu, X. Long, S. Xin, L. Liu, T. Komura, X. Yuan, et al. Coverage axis++: Efficient inner point selection for 3d shape skeletonization. *arXiv preprint arXiv:2401.12946*, 2024. 3, 8
- [53] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, and X. 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, 3
- [54] S. Wu, H. Huang, M. Gong, M. Zwicker, and D. Cohen-Or. Deep points consolidation. *ACM Transactions on Graphics (ToG)*, 34(6):1–13, 2015. 3
- [55] Z. Wu, C. Yu, Y. Jiang, C. Cao, F. Wang, and X. Bai. Sc4d: Sparse-controlled video-to-4d generation and motion transfer. *arXiv preprint arXiv:2404.03736*, 2024. 3
- [56] 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. 3
- [57] 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. 1
- [58] Y. Xu, J. Zhang, Q. Zhang, and D. Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in Neural Information Processing Systems*, 35:38571–38584, 2022. 9

- [59] Z. Xu, Y. Zhou, E. Kalogerakis, and K. Singh. Predicting animation skeletons for 3d articulated models via volumetric nets. In *2019 international conference on 3D vision (3DV)*, pages 298–307. IEEE, 2019. 3
- [60] Z. Yang, X. Gao, W. Zhou, S. Jiao, Y. Zhang, and X. Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20331–20341, 2024. 1, 3
- [61] W. Yifan, N. Aigerman, V. G. Kim, S. Chaudhuri, and O. Sorkine-Hornung. Neural cages for detail-preserving 3d deformations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 75–83, 2020. 3
- [62] Y. Yin, D. Xu, Z. Wang, Y. Zhao, and Y. Wei. 4dgen: Grounded 4d content generation with spatial-temporal consistency. *arXiv preprint arXiv:2312.17225*, 2023. 3, 6, 7
- [63] C. Yu, Q. Zhou, J. Li, Z. Zhang, Z. Wang, and F. Wang. Points-to-3d: Bridging the gap between sparse points and shape-controllable text-to-3d generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 6841–6850, 2023. 1
- [64] Y.-J. Yuan, Y.-T. Sun, Y.-K. Lai, Y. Ma, R. Jia, and L. Gao. Nerf-editing: geometry editing of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18353–18364, 2022. 1
- [65] Y. Zeng, Y. Jiang, S. Zhu, Y. Lu, Y. Lin, H. Zhu, W. Hu, X. Cao, and Y. Yao. Stag4d: Spatial-temporal anchored generative 4d gaussians. *arXiv preprint arXiv:2403.14939*, 2024. 3, 6, 7
- [66] J.-P. Zhang, C.-F. Pu, M.-H. Guo, Y.-P. Cao, and S.-M. Hu. One model to rig them all: Diverse skeleton rigging with unirig. *arXiv preprint arXiv:2504.12451*, 2025. 1, 3, 4, 6, 8
- [67] L. Zhang, A. Rao, and M. Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023. 9
- [68] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6
- [69] T. Zhang, Q. Gao, W. Li, L. Liu, and B. Chen. Bags: Building animatable gaussian splatting from a monocular video with diffusion priors. *arXiv preprint arXiv:2403.11427*, 2024. 3
- [70] Y. Zhao, Z. Yan, E. Xie, L. Hong, Z. Li, and G. H. Lee. Animate124: Animating one image to 4d dynamic scene. *arXiv preprint arXiv:2311.14603*, 2023. 3
- [71] S. Zuffi, A. Kanazawa, D. W. Jacobs, and M. J. Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6365–6373, 2017. 2