

SkelMo: Universal Skeletal Motion Generation for 3D Rigged Shapes

Ye Tao^①, Yuxin Yao^{* ②}, Kendong Liu^③, Dapeng Wu^④, and Junhui Hou^{* ⑤}

Department of Computer Science,
City University of Hong Kong,
Hong Kong SAR, China



Fig. 1: SkelMo utilizes topology agnostic diffusion to generate skeletal animation according to the given skeleton and driving video.

Abstract. Motion generation for rigged shapes is vital for scalable 4D asset production. However, template-based methods are limited by specific topologies and fail to generalize across diverse morphologies. Conversely, per-case optimization is computationally expensive, susceptible to local optima, and highly sensitive to viewpoint-induced ambiguities. In this paper, we present SkelMo, a diffusion-based framework designed for category-agnostic skeletal animation generation from 2D video guidance. To overcome the scarcity of high-quality training data, we have curated a large-scale dynamic dataset comprising approximately 25,200 diverse 3D animations, each featuring complete textures, skeletal rigging, and a wide array of comprehensive animation sequences. To bridge the kinematic gap between 2D visual motion cues and heterogeneous 3D skeletal structures, we propose a structural-semantic injection mechanism. Our model integrates texture and semantic attributes directly into skeletal joint representations. This allows it to map perceived visual dynamics to specific joint hierarchies and their functional roles. This enables SkelMo to synthesize high-fidelity animations that maintain anatomical consistency across a vast range of unseen categories, from existing biological

* Corresponding authors. Email: {yuxinyao,jh.hou}@cityu.edu.hk

species to fantastical beings. Extensive experiments demonstrate that our approach significantly outperforms existing methods, setting a new state-of-the-art benchmark for robust and efficient 4D asset generation. Project Page: <https://research.davytao.me/skelmo/>.

Keywords: Motion Generation · Skeletal Animation · Diffusion Model

1 Introduction

Generating high-quality motion for rigged objects is a fundamental challenge in computer graphics and computer vision. In the digital realm, this technology addresses a critical bottleneck in 4D asset creation by reducing the reliance on manual animation. As a result, it has the potential to transform immersive entertainment, including video games, virtual reality (VR), and digital avatars. Beyond virtual environments, robust motion generation also serves as a crucial bridge for embodied AI, enabling physical robots to execute realistic and physically plausible movements in the real world. Traditionally, realistic motion is generated through two primary approaches. The first relies on expensive capture systems, such as calibrated multi-camera setups or IMU devices, which require specific environments and are often restricted to certain categories like human bodies. The second approach involves manual animation by professional artists, which is labor-intensive and difficult to scale for diverse virtual characters. Recently, some frameworks [6, 15, 26, 32, 41, 44] have been proposed to automatically synthesize motion from various input modalities, such as text or video. However, due to limited diversity in training data, these methods are primarily restricted to human bodies, quadrupeds, or simple articulated shapes like drawers. Any-Top [9] introduces an architecture for universal motion generation. However, in the absence of fine-grained conditional control and due to the limited scale of training data (fewer than 100 objects), its practicality remains inherently constrained. With the rapid advancement of video generative models [28, 37], synthesizing motion from monocular videos has emerged as a promising direction. A straightforward approach, as seen in [24, 33], involves optimizing skeletal poses to align the deformed shape with the video frames. Nevertheless, such methods are often sensitive to rapid movements and rely heavily on external priors like depth estimation or point tracking. Furthermore, the requirement for case-specific optimization leads to significant computational overhead.

In this paper, we propose SkelMo, a diffusion-based framework for category-agnostic skeletal animation generation. A major hurdle in training such a generative model is that high-quality 3D models with both skeletal rigging and diverse animations are scarce in existing open-source datasets. In response to this data scarcity, we curate a large-scale dynamic dataset consisting of approximately 25,200 carefully selected 3D animations. Each model is equipped with high-fidelity textures, complete skeletal rigging, and a wide array of skeletal animation sequences. The dataset is meticulously preprocessed with rest-pose normalization and orientation alignment to ensure cross-category consistency.

Building on this foundation, our video-conditioned model generates skeletal animation parameters for input 3D models, with the ability to generalize across diverse shape types. To achieve true category-agnostic generalization, a key challenge lies in aligning 2D visual motion cues with the diverse 3D skeletal topologies of heterogeneous objects. To address this, we encode the textural and semantic information of each 3D model into its skeletal joint representations, allowing the model to effectively correlate visual motion patterns from the video with object-specific features and achieve better generalization.

In summary, our contributions are as follows.

- We curate a high-quality dataset of approximately 25,200 3D animations featuring complete textures, skeletal rigging, and diverse animation sequences, all preprocessed for cross-category consistency.
- We propose SkelMo, a diffusion-based framework technically featured with a bidirectional video-skeleton fusion and a texture-semantic injection mechanism to enable precise alignment between visual motion cues and diverse 3D morphologies.
- We demonstrate that our approach pushes the performance of skeletal motion synthesis to a new height, setting a new state-of-the-art benchmark for generalization across unseen categories and complex motion sequences.

2 Related Work

Motion Capture. Motion capture is traditionally categorized into two primary domains: human-centric and non-human domains. For human subjects, state-of-the-art methods predominantly focus on recovering the pose and shape of parametric body models from monocular or multi-view inputs (e.g., HMR [11], Vibe [13], TCMR [5]). A dominant line of work regresses SMPL [18] or SMPL-X [20] parameters from images or video clips, leveraging strong template priors to ensure anatomical plausibility and stable tracking. These approaches span optimization-based fitting (e.g., SMPLify-X [20]) and modern feed-forward one-stage regressors with ViT-style backbones for robust whole-body recovery (e.g., OSX [16], Aios [27]). To obtain physically meaningful trajectories, recent monocular systems incorporate SLAM/visual odometry and motion priors to infer world-grounded global motion (e.g., Pace [14], Tram [29]), yet they still rely on a fixed human template and are not designed for arbitrary skeletons. Non-human motion capture and reconstruction is often studied through model-free (e.g., CMR [12], Lassie [38], Magicpony [30], and Lasr [35]) or model-based (e.g., SMAL [45]) animal pipelines. Although effective within their target domains, these methods are typically species- or template-specific, which limits their ability to generalize to the diverse and often non-animal skeletal structures found in arbitrary animatable assets. Concurrently, MoCapAnything [10] introduces a monocular video-based motion prediction method for category-agnostic skeletons. However, its performance is constrained by the reliance on estimated meshes and limited training data.




Fig. 2: Dataset analysis and statistics. (a) Category proportions. (b) Object counts follow a long-tailed distribution across categories. (c) Video-frame totals are substantial overall.

4D Content Generation. 4D generation, which focuses on dynamic 3D content synthesis, has recently attracted increasing attention. Current 4D generation methods primarily differentiate themselves through various combinations of 3D representations, temporal modeling techniques, and algorithmic paradigms. Early approaches (e.g., Sv4d [34], Sv4d2.0 [39], 4Diffusion [43] and DreamGaussian4D [21]) predominantly adopted an optimization-based framework to recover dynamic 3D models by lifting 2D video observations into the 3D space. While these methods are category-agnostic, they are often computationally expensive and struggle to maintain spatio-temporal consistency across long sequences. To address these efficiency bottlenecks, subsequent research (e.g., AnimateAnyMesh [31], DriveAnyMesh [23], and GVFD [42]) explored feed-forward networks to directly predict deformation fields or motion trajectories for temporal representation. Although these methods offer significantly faster inference, they frequently suffer from geometric artifacts where rigid regions undergo unnatural non-rigid distortions. Recently, as skeletal generative models (e.g., Rig-Anything [17], Magicarticulate [25], Puppeteer [24], and Shapegen4d [40]) have matured, skeleton-based animation generation has emerged as a more principled alternative for achieving structured and plausible motion. Despite its potential, research in this direction remains in its infancy, with AnyTop [9] being one of the few pioneering works to explore unconditional skeletal animation generation.

3 Dataset Construction

To empower SkelMo with the capability of animating arbitrary characters, we first establish a large-scale, category-diverse foundation. Unlike previous datasets restricted to specific species, our collection encompasses a wide range of biological taxonomies, including humans, stylized humanoids, quadrupeds, bipeds, and multi-legged creatures, as statistically summarized in Fig. 2. We use LLM-assisted labeling for high-level category statistics. Overall, the category distribution exhibits a mild long-tail pattern while still providing broad coverage across

diverse creature types. The dataset comprises approximately 25,200 unique 3D animations and paired rendered videos and corresponding skeletal motion sequences, totaling more than 3.5 million frames of high-fidelity skeletal data. Each entry features a meticulously rigged skeleton sourced from Articulation-XL [25], a curated subset of the Objaverse dataset [7, 8]. This refined skeletal structure provides the necessary foundation for accurate and anatomically sound motion transfer. To populate this dataset with diverse motions, we developed an automated pipeline that re-associates high-quality animation clips from the original Objaverse source and transfers them onto the refined skeletal structures using AutoRig-Pro [1]. This process effectively recovers missing motion data while ensuring anatomical consistency. Furthermore, we leverage the spatial reasoning capabilities of the Qwen large language model [2] to ensure that all models are oriented in a consistent and correct forward-facing direction. To support video-based conditioning, each animation sequence is rendered from a canonical front-view perspective, resulting in a large-scale collection of paired video-motion data. This data diversity allows SkelMo to learn a generalized mapping between visual appearance and underlying articulation, thereby overcoming the limitations of template-specific priors such as SMPL [18] or SMAL [45].

4 Proposed Method

Overview. Given a category-agnostic 3D mesh $\mathcal{M} = \{\mathcal{V}, \mathcal{F}\}$ with its associated skeleton $\mathcal{S}^0 = \{\mathcal{P}^0, \mathcal{G}\}$, where $\mathcal{V}$ and $\mathcal{F}$ denote the vertex positions and triangle faces, while $\mathcal{P}^0 \in \mathbb{R}^{J \times 3}$ and $\mathcal{G}$ represent the joint positions and topological hierarchy respectively. Along with a monocular driving video $\mathcal{X} = \{I^l\}_{l=1}^L$, which can be synthesized by a generative model, our objective is to produce a sequence of posed skeletons $\{\mathcal{S}^l = (\mathcal{P}^l, \mathcal{G})\}_{l=1}^L$. This sequence should faithfully animate the mesh $\mathcal{M}$ according to the motion dynamics observed in $\mathcal{X}$.

To achieve this, we develop a topology-agnostic diffusion model as our backbone architecture (Sec. 4.1). Building upon this foundation, we propose SkelMo, a diffusion-based framework illustrated in Fig. 3 that effectively bridges visual motion cues with structured skeletal hierarchies. SkelMo consists of two core technical modules. First, to bridge the gap between 2D visual semantics and 3D structural joints, we develop a skinning-aware texture-semantic injection module (Sec. 4.2). Second, we implement a bidirectional motion-topology alignment mechanism within a diffusion transformer (Sec. 4.3) to reconcile driving video cues with the target skeletal hierarchy. This systematic approach ensures that the generated motions are both high-fidelity to the source video and anatomically consistent with the target morphology.

4.1 Topology Agnostic Diffusion Model

To animate characters with arbitrary skeletal configurations, we require a generative backbone that is independent of specific joint templates or hierarchical




Fig. 3: Overview of SkelMo Pipeline. Given a mesh $\mathcal{M}$ with its skeleton $\mathcal{S}^0$ and driving video $\mathcal{X}$, our approach follows a structured pipeline: At each sampling step t , the driving video $\mathcal{X}$ is encoded using DINOv2 Encoder, and fed into the video branch. Simultaneously, the initial skeleton $\mathcal{S}^0$ is projected into the latent space as z^0 via a linear layer. This condition is then concatenated with the current noisy latents z_t , serving as a structural prior. Both noisy latent z_t and its corresponding video feature f are processed through the denoising network consisting of M repeated blocks, and a bidirectional video-skeleton fusion module then integrates these features. To incorporate geometric constraints, multi-view visual features are extracted from the static mesh $\mathcal{M}$, projected onto the vertices. We apply skinning weights-based aggregation to the mesh vertices to extract representative visual features. These features are then injected into the fused representation via cross attention, where skinning weights serve as a masking mechanism to ensure anatomically plausible feature alignment. For the last step, the resulting denoised latent $\hat{z}_0$ is passed through a linear layer to produce the final skeletal animation $\{\mathcal{P}^l \in \mathbb{R}^{J \times 3}\}_{l=1}^L$.

constraints. The primary challenge in such a topology-agnostic setting is maintaining spatial consistency between the driving 2D pixels and the 3D joints without a predefined structural prior.

Motion Representation. In contrast to traditional methods that rely on rotation-based relative representations, we adopt global Cartesian coordinates $\mathcal{P} \in \mathbb{R}^{L \times J \times 3}$ as our primary motion descriptor. While relative rotations excel in local joint smoothness, they introduce complex, non-linear mappings during 2D-to-3D lifting, often leading to spatial misalignment in video-conditioned tasks. By directly diffusing in the xyz coordinates space, we establish a more intuitive geometric correspondence with the driving video $\mathcal{X}$, which stabilizes the denoising process and ensures better global trajectory alignment.

Structural Conditioning. To ensure that the generated coordinates respect the target morphology, we incorporate the rest-pose skeleton $\mathcal{S}^0$ as a persistent structural prior. Specifically, the joint positions $\mathcal{P}^0$ are projected into a latent space $z^0 \in \mathbb{R}^{J \times C}$ and concatenated with the noisy latent z_t along the temporal

axis at each diffusion step t . This conditioning allows the model to perceive the character’s bone lengths and proportions as a constant geometric anchor, preventing anatomical distortions during the iterative sampling process.

Topology-Aware Denoising. Our denoising network consists of M transformer blocks optimized for skeletal data. To capture dependencies across arbitrary hierarchies, we employ the Skeletal Attention mechanism [9]. This module modulates the self-attention weights using a tree-based distance matrix derived from the topology $\mathcal{G}$, ensuring that joints with direct biological relationships (e.g., parent-child) maintain higher spatial correlation. The final denoised latent $\hat{z}_0$ is projected back to the coordinate space to produce the animation $\{\mathcal{P}^l\}_{l=1}^L$.

4.2 Skinning-Aware Texture-Semantic Injection

While the backbone captures geometric constraints, a purely structural representation lacks the semantic context of the shape’s appearance. To provide the model with a fine-grained understanding of the appearance of the object and its semantic structure, we design a specialized DINO-based feature injection mechanism. Specifically, for a given 3D mesh $\mathcal{M}$, we render a set of static images $\{I_i^0\}_{i=1}^N$ from N diverse viewpoints distributed around the object. We utilize a pretrained DINOv2 [19] model to extract dense visual descriptors $F_i = \Phi_{DINO}(I_i^0)$, where $F_i \in \mathbb{R}^{H \times W \times C}$ represents the spatial feature map for the i -th view. To establish a surface-aligned representation, these 2D features are back-projected onto the visible vertices of the 3D mesh. For each vertex $v \in \mathcal{V}$, its corresponding feature f_v is computed by aggregating the projected descriptors from all viewpoints where v is not occluded:

$$f_v = \text{Mean}(\{F_i(\pi_i(v)) \mid \text{Vis}(v, i) = 1\}), \quad (1)$$

where π_i denotes the projection function from the 3D space to the i -th image plane, and $\text{Vis}(v, i)$ is a binary visibility mask. To effectively integrate high-dimensional surface semantics into the skeletal space, we leverage skinning weights as a geometric prior to guide the feature injection process. Let $W = \{w_{j,v}\} \in \mathbb{R}^{J \times |\mathcal{V}|}$ denote the skinning weight matrix between J skeletal joints and mesh vertices. For the j -th joint, we normalize the weights such that the influence of all vertices sums up to unity:

$$\hat{w}_{j,v} = \frac{w_{j,v}}{\sum_{v \in \mathcal{V}} w_{j,v}}, \quad \text{s.t.} \sum_{v \in \mathcal{V}} \hat{w}_{j,v} = 1. \quad (2)$$

This normalization ensures that each joint acts as a localized aggregator for the visual semantics of its associated body parts. The aggregated joint feature $\tilde{F}_S \in \mathbb{R}^{J \times C}$ is then computed through a direct linear projection:

$$\tilde{F}_S = \hat{W} F_V, \quad (3)$$

where $F_V \in \mathbb{R}^{|\mathcal{V}| \times C}$ denotes the collected vertex features $\{f_v\}$. Each row of $\tilde{F}_S$ encodes the semantic profile of a corresponding joint, distilled from its physically

associated surface regions. Once the joint-specific features $\tilde{F}_S$ are computed, they are injected into the motion synthesis network. We employ a cross-attention mechanism where the queries Q are the hidden states of the skeletal joints within the network, while the keys K and values V are derived from the pre-computed $\tilde{F}_S$.

4.3 Bidirectional Video-Skeleton Fusion

Building upon the static geometry and semantic features derived above, we now introduce how SkelMo synchronizes these priors with dynamic motion cues from the driving video. This stage begins with a spatial-temporal video encoder to extract motion cues, followed by a bidirectional cross-modal interaction between motion cues and skeletal topology.

Video Encoder. Given a driving video $\mathcal{X}$ consisting of L frames, we first extract a sequence of raw visual descriptors $F_{\mathcal{X}} = \{f_{v,l}\}_{l=1}^L$ using a pre-trained DINOv2 backbone. To capture global motion context and inter-frame dependencies within the model, we employ an integrated video encoder consisting of successive spatial and temporal self-attention layers. In this stage, spatial attention captures the instantaneous pose semantics in each frame, while temporal attention encodes the dynamic trajectory across the sequence, yielding a motion-aware video representation $F_{\mathcal{X}}^{global}$.

Symmetrical Cross-modal Interaction. The core of our fusion module is a bidirectional cross-attention mechanism that treats the skeletal features $\tilde{F}_S$ and refined video features $F_{\mathcal{X}}^{global}$ as dual sequences. This symmetrical interaction is formalized through two parallel paths. In the first path, the skeletal joints query motion cues from the video frames to guide the synthesis:

$$F_{S \leftarrow \mathcal{X}} = \text{Softmax}\left(\frac{Q_S K_{\mathcal{X}}^T}{\sqrt{d}}\right) V_{\mathcal{X}}, \quad (4)$$

where Q_S is derived from the skeletal embeddings and $(K_{\mathcal{X}}, V_{\mathcal{X}})$ are projected from the video context. Simultaneously, the second path enables the video features to be conditioned on the target skeletal topology:

$$F_{\mathcal{X} \leftarrow S} = \text{Softmax}\left(\frac{Q_{\mathcal{X}} K_S^T}{\sqrt{d}}\right) V_S, \quad (5)$$

where d denotes the scaling dimension of the query and key features. This bidirectional alignment ensures that the motion cues are not merely broadcast to the skeleton but are filtered and refined based on the specific joint hierarchy of $\mathcal{S}$. The fused features from both paths are then integrated into a unified spatio-temporal representation, which is fed into the subsequent diffusion denoising blocks. This symmetrical interaction ensures that the generated animation maintains high fidelity to the source video while strictly adhering to the structural constraints of the target skeleton.

4.4 Training Objective

Our model is trained as a conditional diffusion model that generates 3D joint sequences. Following the standard diffusion formulation, we employ the simplified training objective. The training process aims to predict the noise ϵ added to the ground-truth joint coordinates. Given the noise-scaled joint positions z_t at diffusion step t , the loss function is defined as:

$$L_{\text{diff}} = \mathbb{E}_{z_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, C)\|^2], \quad (6)$$

where z_0 represents the ground-truth 3D joint coordinates, $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ is the Gaussian noise added at diffusion step t , and ϵ_θ is our network that predicts the noise given the noisy sample z_t , the timestep t , and the video conditioning C .

5 Experiments

5.1 Experiment Settings

Training Details. We trained our model from scratch. The model was trained for 80,000 steps with a batch size of 16. We used AdamW with a learning rate starting with 1×10^{-4} to optimize the model. We implemented a step decay schedule where the learning rate is multiplied by 0.99 every 10,000 steps. The experiments were conducted on 4 NVIDIA RTX 5880-48G GPUs.

Post-processing. To ensure kinematic consistency, we applied an IK-based refinement to the raw generated joints. Specifically, we projected the synthesized coordinates onto the predefined skeleton $\mathcal{S}^0$ through 150 iterations of Inverse Kinematics (IK), effectively eliminating bone stretching while maintaining the rigid geometric constraints of the target character.

Benchmark. To evaluate SkelMo, we selected a standardized subset of 100 randomly chosen models from our constructed dataset. We rendered a video of no more than 40 frames from the front view. The object was positioned at the origin, with the camera placed at a distance of 2 meters for rendering.

Baselines. We compare our model against several state-of-the-art (SOTA) methods to evaluate its performance in rigged shape motion generation. Ideally, the most relevant work to ours is the concurrent MocapAnything [10]; however, as their source code and models are not publicly available, a direct comparison is not feasible. We therefore select the following representative methods as our baselines. First, we consider Puppeteer [24], which is primarily a skeleton generation framework; however, one of its key applications involves generating animations via an optimization process guided by optical flow and depth estimation priors. Second, we evaluate against Motion 3-to-4 [4] and ActionMesh [22], which represent the SOTA in video-to-4D generation by focusing on the synthesis of dense vertex trajectories. To bridge the architectural gap between these paradigms and ensure a fair comparison, we implement a standardized post-processing protocol (details in the *Supplementary Material*). Specifically, we fit the vertex trajectories produced by Motion 3-to-4 and ActionMesh into skeletal

Table 1: Quantitative comparison with state-of-the-art methods. We evaluate our SkelMo against several baseline methods on our proposed test dataset. MPJPE (Mean Per Joint Position Error) and CD (Chamfer Distance) are reported to measure motion accuracy and surface fidelity, respectively. Bold values indicate the best performance.

Metrics	Motion 3-to-4 [4]	Puppeteer [24]	ActionMesh [22]	SkelMo (Ours)
MPJPE↓	0.230	0.161	0.197	0.054
CD↓	0.475	0.155	0.418	0.086

animations, allowing for a unified evaluation of motion quality and structural integrity across all methods.

Metrics. To quantify the accuracy of the generated skeletal motion, we employ two key metrics: Mean Per-Joint Position Error (MPJPE) for joint localization and Chamfer Distance (CD) for overall geometric fidelity. For baseline methods that directly output mesh sequences, we first extract the underlying skeletal animation via a fitting process to ensure a consistent and fair comparison. We further evaluate robustness under auto-rigged skeletons.

5.2 Comparison on Motion Generation

In the quantitative evaluation, as shown in Table 1, SkelMo outperforms vertex-trajectory-based generative methods, such as Motion 3-to-4 and ActionMesh, across all metrics. Specifically, our method achieves a reduction in MPJPE, consistently outperforming Motion 3-to-4 and ActionMesh as shown in Table 1. This lower error indicates a higher precision in capturing underlying motion dynamics. Furthermore, regarding the Chamfer Distance (CD), which measures geometric shape fidelity, we achieve a CD value that is lower than the competitive baselines, as detailed in Table 1. These elevated CD scores suggest that the baseline methods struggle to preserve the original structural integrity of the object during motion.

These numerical findings are consistent with the visual comparisons presented in Fig. 4. A primary limitation of pure vertex-trajectory generation is the susceptibility to non-rigid distortion. As observed in the case on the left, both Motion 3-to-4 and ActionMesh suffer from mesh adhesion and unnatural stretching artifacts. This phenomenon typically occurs because these frameworks independently predict vertex displacements, often lacking sufficient global constraints to maintain the spatial relationships between neighboring vertices. This leads to an elongated or "melted" appearance where the mesh topology is compromised. In contrast, our SkelMo maintains the rigid and articulated structures of the object more effectively, resulting in clean surfaces and stable geometry even during complex sequences. *The **Supplementary Material and Demo Video** provide more visual results.*




Fig. 4: Visual comparison between our SkelMo and other state-of-the-art methods. For each example, the left side displays the rendered 3D mesh, while the right side shows its corresponding skeletal structure.

5.3 Ablation Study

In the ablation studies, all models are trained following the training procedures described previously.

Effect of Texture-Semantic Injection. We investigate the impact of the texture-semantic injection module described in Sec. 4.2. As shown in Table 2, removing the DINO Feature Cross Attention module leads to a measurable performance decline. As shown in Fig. 5 (a), without the DINO-based semantic injection, the model struggles to identify the character’s boundary and specific body parts. This leads to collapsed skeletal structures (e.g., the Spiderman example in the 2nd row), where the joints cluster together because the model lacks the semantic "priors" of what a skeleton should look like under that specific texture.

Effect of Bidirectional Video-Skeleton Fusion. We evaluate the necessity of the bidirectional video-skeleton fusion module introduced in Sec. 4.3. By ablating

Table 2: Ablation study on key components. We evaluate the contribution of each module to the overall performance. "w/o Inject." denotes the removal of texture-semantic injection; "w/o Bidir." refers to the absence of bidirectional video-skeleton fusion; and "Rel. Repr." represents the use of relative skeletal coordinates. The Full Model achieves the best performance across all metrics.

Setting	w/o Inject.	w/o Bidir.	Rel. Repr.	Full Model
MPJPE↓	0.234	0.118	0.285	0.054
CD↓	0.235	0.147	0.340	0.086




Fig. 5: Visual comparison of different model variants. From left to right: (a) Results without texture-semantic injection. (b) Results without bidirectional video-skeleton fusion. (c) Results using relative skeletal representations instead of global coordinates. (d) Results of our full model. Each column presents the generated skeletal poses for different input characters and driving frames.

the $S \rightarrow \mathcal{X}$ and $\mathcal{X} \rightarrow S$ cross-attention mechanisms, we observe a consistent drop in performance, as shown in Table 2. As shown in Fig. 5 (b), while the model can still capture the general motion, it fails to achieve fine-grained alignment with the driving frame’s subtle nuances (e.g., the exact angle of the elbows or the extension of the limbs).

Effect of Relative Skeletal Representations. We examine the influence of skeletal representation methods as discussed in Sec. 4.1. Replacing global xyz coordinates with relative representations leads to a performance degradation, as evidenced in Table 2. In Fig. 5 (c), using relative coordinates instead of global ones leads to spatial disorientation. The model fails to maintain the correct global height and ground contact, often resulting in "floating" artifacts or inconsistent limb lengths. This suggests that global spatial context provides more robust guidance for the model to capture absolute motion trajectories accurately.

Robustness to Auto-rigged Skeletons. To evaluate robustness under imperfect rigging, we additionally test SkelMo with automatically generated skeletons (e.g., RigAnything [17]) instead of ground-truth rigs. Since auto-rigged skeletons introduce topology and joint layout discrepancies, skeletal-space metrics become inconsistent; thus, we report vertex-space metrics instead. As shown in Table 3, SkelMo maintains strong performance under auto-rigged skeletons with only minor degradation.

Table 3: Performance comparison under ground-truth (GT) and automatically generated (auto-rigged) skeletons. SkelMo shows consistent performance with only minor degradation under auto-rigged rigs.

Metrics ($\times 10^{-1}$)	Motion 3-to-4	Puppeteer	ActionMesh	SkelMo (w. GT)	SkelMo (w. Auto-rigged)
MSE↓	0.369	0.229	0.265	0.048	0.082
Vertex-CD↓	0.230	0.158	0.173	0.107	0.146

5.4 Application

Motion Transfer. Beyond standard motion generation, the topology-agnostic nature of SkelMo enables the challenging task of cross-entity motion transfer. In this setting, we replace the source video with a driving sequence featuring a completely different body shape or species (e.g., transferring motion from a quadruped to a humanoid, or vice versa). As illustrated in Fig. 6, our model successfully preserves the core motion dynamics of the driving video while faithfully adapting them to the geometric constraints of the target skeleton. The first two rows of Fig. 6 demonstrate intra-category transfer (human-to-human and quadruped-to-quadruped), where the model accurately captures fine-grained skeletal movements. More importantly, the third row highlights the capability of the model in cross-species transfer (e.g., human-to-quadruped skeleton).

This robust adaptation stems from our representation using global Cartesian coordinates, which provides a universal geometric language for motion, and the Skeletal Attention mechanism, which re-targets these coordinates to remain structurally valid for the new topology. Even with significant discrepancies in bone proportions and joint hierarchies, SkelMo produces natural and synchronized animations, highlighting its robustness in complex cross-category motion synthesis.

In-the-Wild Scenarios. To further evaluate the robustness and practical applicability of our method, we conduct experiments using driving videos from two distinct sources: AI-generated videos and real-world captures. Specifically, we utilize high-quality synthetic videos generated by AI (e.g., Doubao AI [3]) as driving signals to test the model’s performance on stylized motions. The first row of Fig. 7 demonstrates this by successfully transferring motion from




Fig. 6: Visualization of cross-category motion transfer results. Each row represents a specific transfer case across multiple frames. The leftmost column displays the target static meshes, while the subsequent columns show pairs of the driving video (source motion) and the resulting target skeleton. The first and second rows demonstrate intra-category transfer (human-to-human and quadruped-to-quadruped, respectively), while the third row illustrates cross-species transfer (e.g., human-to-quadruped). The extracted skeletons validate the model’s robustness in retargeting complex motions across significantly different skeletal morphologies and joint configurations.

an AI-generated character to a different target shape. Additionally, we apply our method to raw in-the-wild real-world videos to assess its stability under complex lighting and diverse backgrounds. The second row illustrates this performance on real-world captures where high-fidelity motion retargeting is consistently achieved. These results across both scenarios demonstrate strong generalization capabilities for various downstream applications.

6 Conclusion and Discussion

In this paper, we presented SkelMo, a diffusion-based framework designed to bridge the gap between 2D visual motion cues and 3D skeletal animation for category-agnostic rigged objects. By curating a large-scale, high-fidelity dataset of approximately 25,200 3D animations with complete rigging and diverse animations, we have addressed a critical data bottleneck in the field. By leveraging a bidirectional video-skeleton fusion and a texture-semantic injection mechanism, our architecture enables the model to transcend the limitations of species-specific




Fig. 7: Visualization of in-the-wild scenarios. Our method demonstrates strong generalization across diverse driving video sources. The first row showcases results driven by AI-generated videos (e.g., from Doubao AI), featuring stylized motions. The second row illustrates performance on real-world captures [36], highlighting the stability of the model under complex lighting and backgrounds. In both cases, high-fidelity motion retargeting is consistently achieved.

templates, achieving robust and physically plausible motion synthesis across diverse, unseen morphologies.

Despite these advancements, several promising directions remain for future exploration. A primary goal is to enrich our dataset with fine-grained textual captions that describe the stylistic nuances and intent of each motion sequence. Such semantic grounding will facilitate the transition from video-only conditioning to multimodal generative control, enabling users to synthesize complex animations through text instructions or audio signals. Furthermore, exploring object-environment interactions and physically-aware constraints will be crucial for extending the capabilities of SkelMo toward more complex embodied AI tasks. By expanding both the semantic depth of our data and the flexibility of our conditioning modalities, we aim to provide a truly universal engine for bringing diverse virtual and physical entities to life.

Acknowledgements

This work was supported in part by the Hong Kong Research Grants Council under Grants R1018-25F, N_CityU1114/25, and 11219324; and in part by the National Natural Science Foundation of China under Grant 62422118.

References

1. Auto-rig pro. <https://superhivemarket.com/products/auto-rig-pro>, blender add-on for character rigging and animation retargeting
2. Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., et al.: Qwen technical report. arXiv preprint arXiv:2309.16609 (2023)
3. ByteDance: Doubao. <https://www.doubao.com/> (2026), accessed: 2026-03-05
4. Chen, H., Chen, X., Zhang, Y., Xu, Z., Chen, A.: Motion 3-to-4: 3d motion reconstruction for 4d synthesis. In: CVPR (2026)
5. Choi, H., Moon, G., Chang, J.Y., Lee, K.M.: Beyond static features for temporally consistent 3d human pose and shape from a video. In: CVPR. pp. 1964–1973 (2021)
6. Chu, R., Liu, Z., Ye, X., Tan, X., Qi, X., Fu, C.W., Jia, J.: Command-driven articulated object understanding and manipulation. In: CVPR. pp. 8813–8823 (2023)
7. Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., et al.: Objaverse-xl: A universe of 10m+ 3d objects. NIPS **36**, 35799–35813 (2023)
8. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: CVPR. pp. 13142–13153 (2023)
9. Gat, I., Raab, S., Tevet, G., Reshef, Y., Bermano, A.H., Cohen-Or, D.: Anytop: Character animation diffusion with any topology. In: SIGGRAPH. pp. 1–10 (2025)
10. Gong, K., Wen, Z., He, W., Xu, M., Wang, Q., Zhang, N., Li, Z., Lian, D., Zhao, W., He, X., et al.: Mocapanything: Unified 3d motion capture for arbitrary skeletons from monocular videos. In: CVPR. pp. 7089–7099 (2026)
11. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7122–7131 (2018)
12. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: ECCV. pp. 371–386 (2018)
13. Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: CVPR. pp. 5253–5263 (2020)
14. Kocabas, M., Yuan, Y., Molchanov, P., Guo, Y., Black, M.J., Hilliges, O., Kautz, J., Iqbal, U.: Pace: Human and camera motion estimation from in-the-wild videos. In: 3DV. pp. 397–408. IEEE (2024)
15. Li, Z., Luo, M., Hou, R., Zhao, X., Liu, H., Chang, H., Liu, Z., Li, C.: Morph: A motion-free physics optimization framework for human motion generation. In: ICCV. pp. 14580–14589 (2025)
16. Lin, J., Zeng, A., Wang, H., Zhang, L., Li, Y.: One-stage 3d whole-body mesh recovery with component aware transformer. In: CVPR. pp. 21159–21168 (2023)
17. Liu, I., Xu, Z., Yifan, W., Tan, H., Xu, Z., Wang, X., Su, H., Shi, Z.: Riganything: Template-free autoregressive rigging for diverse 3d assets. TOG **44**(4), 1–12 (2025)
18. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. In: Seminal Graphics Papers: Pushing the Boundaries, Volume 2, pp. 851–866 (2023)
19. Oquab, M., Darcet, T., Moutakanni, T., Vo, H.V., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., Howes, R., Huang, P.Y., Xu, H., Sharma, V., Li, S.W., Galuba, W., Rabbat, M., Assran, M., Ballas, N., Synnaeve, G., Misra, I., Jegou, H., Mairal, J., Labatut, P., Joulin, A., Bojanowski, P.: Dinov2: Learning robust visual features without supervision (2023)

20. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: CVPR. pp. 10975–10985 (2019)
21. Ren, J., Pan, L., Tang, J., Zhang, C., Cao, A., Zeng, G., Liu, Z.: Dreamgaussian4d: Generative 4d gaussian splatting. arXiv preprint arXiv:2312.17142 (2023)
22. Sabathier, R., Novotny, D., Mitra, N.J., Monnier, T.: Actionmesh: Animated 3d mesh generation with temporal 3d diffusion. CVPR (2026)
23. Shi, Y., Liu, Y., Wu, Y., Liu, X., Zhao, C., Luo, J., Zhou, B.: Drive any mesh: 4d latent diffusion for mesh deformation from video. arXiv preprint arXiv:2506.07489 (2025)
24. Song, C., Li, X., Yang, F., Xu, Z., Wei, J., Liu, F., Feng, J., Lin, G., Zhang, J.: Puppeteer: Rig and animate your 3d models. NIPS **38**, 72152–72184 (2026)
25. Song, C., Zhang, J., Li, X., Yang, F., Chen, Y., Xu, Z., Liew, J.H., Guo, X., Liu, F., Feng, J., et al.: Magicarticulate: Make your 3d models articulation-ready. In: CVPR. pp. 15998–16007 (2025)
26. Sun, K., Litvak, D., Zhang, Y., Li, H., Wu, J., Wu, S.: Ponymation: Learning articulated 3d animal motions from unlabeled online videos. In: ECCV. pp. 100–119. Springer (2024)
27. Sun, Q., Wang, Y., Zeng, A., Yin, W., Wei, C., Wang, W., Mei, H., Leung, C.S., Liu, Z., Yang, L., et al.: Aios: All-in-one-stage expressive human pose and shape estimation. In: CVPR. pp. 1834–1843 (2024)
28. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models. arXiv preprint arXiv:2503.20314 (2025)
29. Wang, Y., Wang, Z., Liu, L., Daniilidis, K.: Tram: Global trajectory and motion of 3d humans from in-the-wild videos. In: ECCV. pp. 467–487. Springer (2024)
30. Wu, S., Li, R., Jakab, T., Rupperecht, C., Vedaldi, A.: Magicpony: Learning articulated 3d animals in the wild. In: CVPR. pp. 8792–8802 (2023)
31. Wu, Z., Yu, C., Wang, F., Bai, X.: Animateanymesh: A feed-forward 4d foundation model for text-driven universal mesh animation. ICCV (2025)
32. Xiao, L., Lu, S., Pi, H., Fan, K., Pan, L., Zhou, Y., Feng, Z., Zhou, X., Peng, S., Wang, J.: Motionstreamer: Streaming motion generation via diffusion-based autoregressive model in causal latent space. In: ICCV. pp. 10086–10096 (2025)
33. Xie, T., Chen, Y., Guo, Y., Yang, Y., Zhou, B., Terzopoulos, D., Jiang, Y., Jiang, C.: Animimic: Imitating 3d animation from video priors. In: CVPR. pp. 40266–40276 (2026)
34. Xie, Y., Yao, C.H., Voleti, V., Jiang, H., Jampani, V.: Sv4d: Dynamic 3d content generation with multi-frame and multi-view consistency. arXiv preprint arXiv:2407.17470 (2024)
35. Yang, G., Sun, D., Jampani, V., Vlasic, D., Cole, F., Chang, H., Ramanan, D., Freeman, W.T., Liu, C.: Lasr: Learning articulated shape reconstruction from a monocular video. In: CVPR. pp. 15980–15989 (2021)
36. Yang, G., Vo, M., Neverova, N., Ramanan, D., Vedaldi, A., Joo, H.: Banmo: Building animatable 3d neural models from many casual videos. In: CVPR (2022)

37. Yang, Z., Teng, J., Zheng, W., Ding, M., Huang, S., Xu, J., Yang, Y., Hong, W., Zhang, X., Feng, G., et al.: Cogvideox: Text-to-video diffusion models with an expert transformer. arXiv preprint arXiv:2408.06072 (2024)
38. Yao, C.H., Hung, W.C., Li, Y., Rubinstein, M., Yang, M.H., Jampani, V.: Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. NIPS **35**, 15296–15308 (2022)
39. Yao, C.H., Xie, Y., Voleti, V., Jiang, H., Jampani, V.: Sv4d 2.0: Enhancing spatio-temporal consistency in multi-view video diffusion for high-quality 4d generation. In: ICCV. pp. 13248–13258 (2025)
40. Yenphraphai, J., Mirzaei, A., Chen, J., Zou, J., Tulyakov, S., Yeh, R.A., Wonka, P., Wang, C.: Shapegen4d: Towards high quality 4d shape generation from videos. ICLR (2026)
41. Yi, H., Thies, J., Black, M.J., Peng, X.B., Rempe, D.: Generating human interaction motions in scenes with text control. In: ECCV. pp. 246–263. Springer (2024)
42. Zhang, B., Xu, S., Wang, C., Yang, J., Zhao, F., Chen, D., Guo, B.: Gaussian variation field diffusion for high-fidelity video-to-4d synthesis. In: ICCV. pp. 12502–12513 (2025)
43. Zhang, H., Chen, X., Wang, Y., Liu, X., Wang, Y., Qiao, Y.: 4diffusion: Multi-view video diffusion model for 4d generation. NIPS **37**, 15272–15295 (2024)
44. Zhang, Z., Wang, Y., Mao, W., Li, D., Zhao, R., Wu, B., Song, Z., Zhuang, B., Reid, I., Hartley, R.: Motion anything: Any to motion generation. arXiv preprint arXiv:2503.06955 (2025)
45. Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3d menagerie: Modeling the 3d shape and pose of animals. In: CVPR. pp. 6365–6373 (2017)