

DreaMo: Articulated 3D Reconstruction From A Single Casual Video

Tao Tu^{1*} Ming-Feng Li^{2*} Chieh Hubert Lin³ Yen-Chi Cheng⁴ Min Sun^{1,5} Ming-Hsuan Yang³

¹National Tsing Hua University ²Carnegie Mellon University

³University of California, Merced ⁴University of Illinois Urbana-Champaign ⁵Amazon

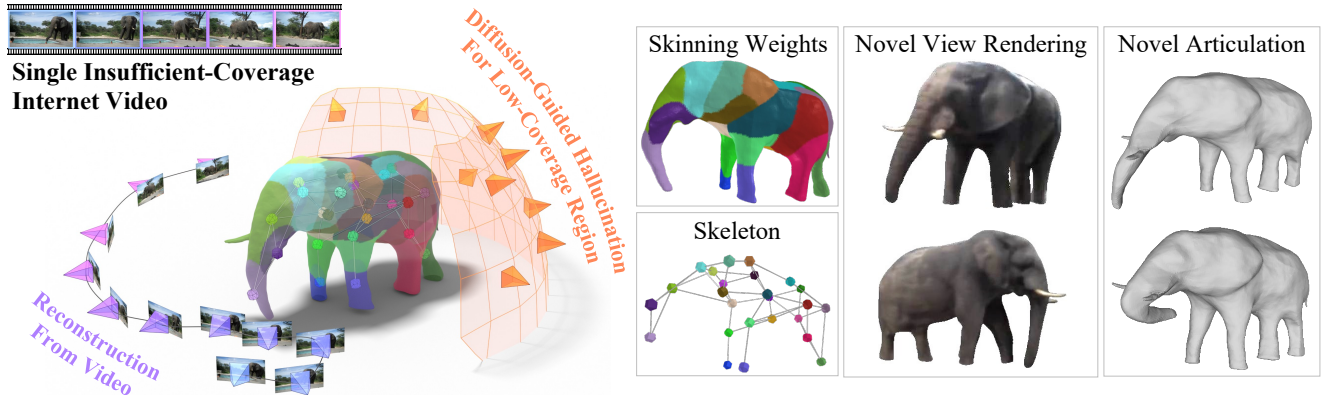


Figure 1. Casual everyday videos often lack sufficient view coverage of the subject, posing a challenge for existing articulated shape reconstruction methods. In contrast, DreaMo not only learns deformable 3D reconstruction from visible viewpoints but also hallucinates the invisible regions with a view-conditioned diffusion model. As a result, DreaMo achieves high-quality rendering of novel views, detailed geometry reconstruction, and provides interpretable and controllable skeletons with skinning weights.

Abstract

Articulated 3D reconstruction has valuable applications in various domains, yet it remains costly and demands intensive work from domain experts. Recent advancements in template-free learning methods show promising results with monocular videos. Nevertheless, these approaches necessitate a comprehensive coverage of all viewpoints of the subject in the input video, thus limiting their applicability to casually captured videos from online sources. In this work, we study articulated 3D shape reconstruction from a single and casually captured Internet video, where the subject’s view coverage is incomplete. We propose DreaMo that jointly performs shape reconstruction while solving the challenging low-coverage regions with view-conditioned diffusion prior and several tailored regularizations. In addition, we introduce a skeleton generation strategy to create human-interpretable skeletons from the learned neural bones and skinning weights without any predefined skeleton structures. We conduct our study on a self-collected internet video collection characterized by incomplete view coverage. DreaMo shows promising quality in novel-view rendering, detailed articulated shape reconstruction, and skeleton generation. Extensive qualitative and quantitative studies validate the efficacy of each pro-

posed component, and show existing methods are unable to solve correct geometry due to the incomplete view coverage.

1. Introduction

Articulated 3D models have extensive applications in movie production, gaming, and virtual reality. These models offer users flexible motion controls, making them highly suitable for content creation across various scenarios and specifications. However, manual creation of such models is expensive and time-consuming, and the quality heavily depends on artists’ skill, often leading to shapes and appearances that deviate from realism. Therefore, there are ongoing explorations into extracting articulated 3D models directly from video data due to its high accessibility on the Internet and the low hardware requirements.

Retrieving a 3D model from casual videos without constraints is a challenging and ill-conditioned problem. The dynamic movement of the subject hampers triangulation, subsequently introducing complexity to the geometry extraction. In addition, self-occlusion often impedes the retrieval of crucial geometry and motion cues necessary for full-shape reconstruction. Therefore, recent methods [14, 65] often utilize parametric template models from existing 3D scans of humans and animals to regulate and guide the

geometry of invisible surfaces.

While effective, collecting 3D scans for arbitrary object categories and wildlife is challenging. BANMo [59] represents a recent endeavor to reconstruct the non-rigid 3D model from videos without template shape priors. Despite BANMo demonstrating promising reconstruction quality when providing multiple videos of the same subject, we observed that it requires dense camera coverage of the same subject from multiple video sequences. This requirement limits the application of retrieving 3D shapes from Internet videos or casual video captures. The real-world videos are not dedicated to 3D shape reconstruction and, thus often have insufficient view coverage of the subject from diverse angles. Through experiments, we show that existing 3D reconstruction methods cannot handle such types of videos.

To address the aforementioned issues, we propose DreaMo, a template-free 3D articulated shape reconstruction framework tailored with joint reconstruction and hallucination. DreaMo simultaneously reconstructs 3D shape using neural radiance field [33], and hallucinates plausible geometry of invisible regions using the diffusion prior [29, 38, 44]. We analyze several design choices and show that careful parameter selection during distilling information from the diffusion model is critical for preserving high-quality surface texture (Section 3.2). In addition, we introduce three regularization techniques to stabilize the learned neural bones, mitigating the generation of eccentric bumps and fragmented structures (Section 3.3). To further enhance the interpretability and controllability of the learned 3D model, we propose a strategy for generating a skeleton based on the neural bones and the learned radiance field without any predefined skeleton structures (Section 3.4).

To validate the performance of DreaMo and the limitations of existing methods on diverse animal species under different capture settings, we collect a set of short video clips with insufficient view coverage from the Internet. Through extensive quantitative and qualitative comparisons, we show DreaMo produces more plausible geometry, texture, and skeleton, compared to existing state-of-the-art approaches in articulated 3D reconstruction. Our ablation study on each proposed design further supports the significance of these design choices.

2. Related Work

Model-based reconstruction. Model-based methods [3, 4, 14, 17, 28, 68] build articulated 3D models from images or videos by utilizing prior parametric models [31, 36, 55, 69, 70] derived from an extensive collection of 3D scans of humans or toy animals. Despite achieving remarkable reconstruction results, collecting 3D scan data is practically challenging, particularly for wildlife animals.

Image-based reconstruction. Due to the ease of obtaining monocular image data from the Internet, prior studies

focus on learning 3D shapes from images with weak 2D supervision, such as keypoints or silhouettes. Some approaches [9, 13, 15, 18, 20, 25, 54, 64] learn a category-specific model from an image collection and perform test-time 3D shape reconstruction using a single image. Additionally, part-based methods [60, 61] assemble 3D parts to construct articulated 3D models from a limited image collection, a scenario akin to ours. However, unlike video, images lack temporal information and smooth transitions across frames, which are valuable cues for 3D reconstruction.

Video-based reconstruction. Video-based methods [6, 10, 12, 19, 21, 22, 24, 53, 57, 58] for articulated 3D reconstruction can leverage the temporal information inherent in given video sequences. Inspired by the promising outcomes in novel-view synthesis research [8, 27, 30, 33, 35, 39], recent 3D articulation reconstruction methods employ differential rendering to minimize reconstruction loss. While some of these methods can generate plausible 3D models, they require specific inputs such as multi-view videos [37, 65], predefined 3D skeletons [37, 49, 51, 52], or 3D rest-pose point clouds [48]. Among them, BANMo [59] shows promising 3D reconstruction results for articulated objects solely using video data. However, it still demands several videos featuring dense camera coverage of the same subject, limiting its feasibility for a single casual video. REACTO [47] is a concurrent work that also tackles template-free articulated reconstruction, but rather focusing on common objects with single articulation. BAGS [67] is a concurrent work that provides a solution for Gaussian-Splatting-based [16] rendering, however, the method still suffers from noisy artifacts in the hallucinated region.

Distillation from diffusion models. 2D diffusion models [42, 44, 45, 66] have demonstrated promising results in generating realistic 2D images. Building upon this, recent works [23, 26, 32, 38, 41, 43, 46] achieve 3D model reconstruction by employing a pretrained 2D diffusion model as a prior for view synthesis [38, 50]. To gain control of the camera viewpoint over the 2D diffusion model, Zero-1-to-3 [29] finetunes the diffusion model on a synthetic dataset [7] and demonstrates the zero-shot ability to generate novel views for a specified subject. Such view control capability allows us to harness the generative power of the 2D diffusion model by imagining parts of the subject that are not observed in the input video, enabling us to improve the reconstruction of the target subject.

3. Approach

We tackle the problem of articulated 3D shape reconstruction from a single video. The recovered 3D shape encapsulates the time-dependent motions observed in the video, and the learned neural bones with skinning weights support novel articulations controlled by the users. Similar to BANMo [59], our DreaMo learns a rest-pose neural 3D

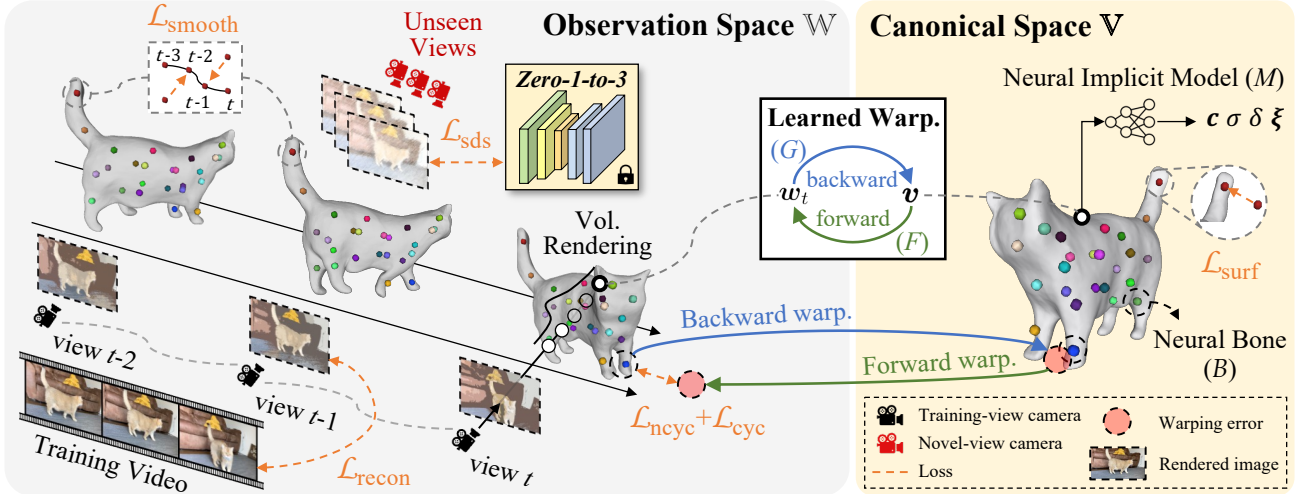


Figure 2. **DreaMo** is a template-free articulated 3D shape reconstruction framework, which jointly performs training-view reconstruction and unseen-region hallucination. It learns a rest-pose neural 3D model (M) using a neural implicit function in the canonical space (\mathbb{V}). A forward warping model (F) transforms this 3D model into observation space (\mathbb{W}), where the video frames supervise the model to capture time-dependent motions. Conversely, a backward warping model (G) performs the inverse operation of F . DreaMo uses Zero-1-to-3 to hallucinate and complete the unseen regions. Meanwhile, we introduce several regularization terms (\mathcal{L}_{nyc} , $\mathcal{L}_{\text{surf}}$, $\mathcal{L}_{\text{smooth}}$) to improve the placement of neural bones and reduce geometric artifacts.

model represented by a neural implicit function in the time-invariant canonical space \mathbb{V} and a separate time-dependent warping implicit function deforms the canonical space features to the observation space \mathbb{W} . All the image supervision involving volume rendering [33] occurs in the observation space. Figure 2 shows the overview of DreaMo.

3.1. Articulated 3D Reconstruction Model

Neural implicit model in canonical space (M). The neural implicit function M in the canonical space represents a rest pose of the 3D reconstruction target. Following previous works in neural rendering [33, 59, 63], we model color c , signed distance function (SDF) δ , density σ , and semantic feature ξ by neural implicit functions with

$$c, \sigma, \delta, \xi = M(v), \quad (1)$$

where $v \in \mathbb{V}$ is a coordinate in canonical space. We implement M with Multi-Layer Perceptron (MLP) and follow VolSDF [63] to convert the SDF values δ to density σ by the cumulative distribution function of Laplacian distribution.

Linear blend skinning. Similar to BANMo [59], in order to provide explicit representation for users to articulate the reconstructed 3D model, we randomly initialize Gaussian ellipsoids and use them to represent a set of neural bones $\mathcal{B} = \{\mu_b, \Sigma_b\}_{b=1}^B$, where B is the number of neural bones, μ_b is the bone position, and Σ_b represents the bone orientation and scale. The transformation of a skin point at coordinate v is determined by weighting over the B bone transformations based on the skinning weights, following the linear blend skinning method [11]. A higher skinning weight implies that the transformation of the associated bone will in-

troduce larger deformation to the skin point v . Specifically, we determine the per-bone skinning weights for v by the Mahalanobis distance between v and the b -th neural bone, which are subsequently normalized using a softmax layer.

Time-dependent warping models (F and G). To represent the time-dependent motion of the reconstructed subject, we use a forward warping model F to transform points in canonical space to the observation space of time t with $w = F(v, t) \in \mathbb{W}$. We also jointly learn a backward warping function $v = G(w, t)$ to approximate the inverse of F . More specifically, we implement the warping functions as a combination of linear blend skinning that blends a set of bone transformations $\{\mathcal{J}_{t,b}\}_{b=1}^B$ for modeling the subject’s deformation, and a global transformation $P_t(\cdot) \in \text{SE}(3)$ that represents the camera transformation.

Volume rendering. We use volume rendering [33, 59] to render images at different time steps t . To render a pixel u at time t , we cast a ray from a posed pinhole camera at time t , and sample a sequence of 3D points w_i^u along the ray in the observation space, where i is the index of the point on the ray. We first obtain the point-wise canonical space color $c_i^{u,t}$, density $\sigma_i^{u,t}$, SDF value $\delta_i^{u,t}$, and semantic features $\xi_i^{u,t}$ in Equation (1) with

$$c_i^{u,t}, \sigma_i^{u,t}, \delta_i^{u,t}, \xi_i^{u,t} = M(G(w_i^u, t)). \quad (2)$$

Subsequently, the pixel color $c^{u,t}$ and silhouette $o^{u,t}$ in the camera space are obtained using volume rendering:

$$c^{u,t} = \sum_i \tau_i c_i^{u,t}, \quad o^{u,t} = \sum_i \tau_i, \quad (3)$$

where $\tau_i = \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$ are the weights along the ray,

$\alpha_i = 1 - \exp(-\sigma_i^{u,t} \Delta_i)$ are the alpha compositing values, and Δ_i is the distance between the i -th sample and the subsequent sample along the ray. Likewise, the camera space semantic features $\xi^{u,t}$ are rendered by replacing the pixel color $c_i^{u,t}$ with semantic feature $\xi_i^{u,t}$ in Equation (3).

3D model manipulation. To articulate the rest-pose 3D model reconstructed by DreaMo into a user-defined pose, the user only needs to specify the bone transformations from the rest to the target pose. With these bone transformations, for any points in the space, we can derive the skinning weights contributed by each bone and determine the deformation of all the points on the object’s surface. In practice, we convert the implicit model in canonical space into an explicit mesh representation $\mathcal{V} = \{\nu_i \in \mathbb{R}^3\}$ using the marching cube algorithm. Subsequently, we morph the mesh by deforming each vertex ν_i to the target pose. Finally, the color of each vertex is determined by querying the implicit model in canonical space.

3.2. View-Conditioned Diffusion Model as Prior

One of our main challenges is recovering appropriate geometry for unseen surfaces of the deformable subject. In order to support the low-coverage or unseen view angles, we use Zero-1-to-3 [29] to synthesize the novel view conditioning on a source image of the same time step and the camera pose of the novel view. These synthetic supervisions are distilled into the articulated 3D reconstruction model using Score Distillation Sampling (SDS) [38].

However, we find naively updating all the trainable parameters with SDS will destroy the high-frequency details of the reconstructed surface texture (Figure 8). This leads to a blurry and over-saturated appearance similar to the relevant works in 3D asset generation [29, 38, 50]. We hypothesize this is due to the randomness of the diffusion model, which causes the novel views inconsistent with the training views and makes the radiance field unable to converge.

As a remedy, recall that we only intend to enhance the geometry of the low-coverage surface, we discover making the SDS gradients only update the geometry-relevant parameters can achieve the goal without compensating the surface texture. In practice, we only update the parameters of the neural bones \mathcal{B} . The ablation study in Table 2 and Figure 8 shows the effectiveness of this parameter selection.

3.3. Learning

Our end-to-end full training objective is a weighted sum of reconstruction loss $\mathcal{L}_{\text{recon}}$, SDS loss \mathcal{L}_{SDS} , and regularization terms \mathcal{L}_{cyc} , $\mathcal{L}_{\text{ncyc}}$, $\mathcal{L}_{\text{surf}}$, $\mathcal{L}_{\text{smooth}}$.

Reconstruction Loss. The reconstruction loss minimizes the discrepancies between DreaMo’s rendered images and the ground-truth images in training views. We also train the model to predict optical flow as it aids in learning correct deformation by providing approximated pixel correspon-

dence across time. For a pixel u at time t , the optical flow $\mathbf{f}_{t \rightarrow t'}^u$ represents the warping vector between u and its new pixel location u' at t' after object deformation and view-point transformation. Such u' is computed by backward warping the observation space points at t to the canonical space, which is time-invariant, forward warping to the observation space at t' , and then applying perspective camera projection $\Gamma_{t'}(\cdot)$ derived from the learned global transformation $\mathbf{P}_{t'}$ and the camera intrinsic at time t' . The resulting u' can be expressed as $u' = \sum_i \tau_i \Gamma_{t'}(F(G(\mathbf{w}_i^u, t)), t')$. Finally, the reconstruction loss is

$$\mathcal{L}_{\text{recon}} = \mathbb{E}_{(u,t,t')} \|\hat{\mathbf{c}}^{u,t} - \mathbf{c}^{u,t}\|_2 + \|\hat{\xi}^{u,t} - \xi^{u,t}\|_2 + \|\hat{o}^{u,t} - o^{u,t}\|_2 + \|\hat{\mathbf{f}}_{t \rightarrow t'}^u - \mathbf{f}_{t \rightarrow t'}^u\|_2, \quad (4)$$

where we randomly sample N combinations of (u, t, t') from all unique combinations for each training iteration.

Regularization. To encourage the inverse relationship between the forward and backward warping functions, we use a training-view cyclic consistency loss \mathcal{L}_{cyc} similar to BANMo. In addition, since our problem involves large regions unseen by the training views, we introduce the novel-view cyclic consistency loss $\mathcal{L}_{\text{ncyc}}$:

$$\mathcal{L}_{\text{cyc}} = \sum_m \tau_m \|\mathbf{w}_m - F(G(\mathbf{w}_m, t), t)\|_2, \quad (5)$$

$$\mathcal{L}_{\text{ncyc}} = \sum_n \tau_n \|\mathbf{w}_n - F(G(\mathbf{w}_n, t), t)\|_2, \quad (6)$$

where \mathbf{w}_m and \mathbf{w}_n are coordinates of points on the training-view and novel-view rays, respectively.

In addition, learning the neural bones without additional constraints allows them to scatter all over the space. This creates unusual geometry and floaters in the regions with insufficient view coverage. Therefore, we introduce a surface constraint loss $\mathcal{L}_{\text{surf}}$ to keep the bones beneath the SDF surface:

$$\mathcal{L}_{\text{surf}} = \|\max\{\delta, 0\}\|_2 \quad (7)$$

where δ is the SDF value of the bones. Meanwhile, we observe that the learned transitions of bones in the low-coverage or self-occluded regions often exhibit unnatural jiggles. These unnatural motions also frequently produce broken geometry and floaters. For this, we design a smooth transition loss $\mathcal{L}_{\text{smooth}}$ to encourage both the translation and rotation of each bone to have steady change over time:

$$\mathcal{L}_{\text{smooth}} = \sum_{b=1, t=1}^{B, T-1} \frac{\text{ang}(\mathbf{R}_b^t, \mathbf{R}_b^{t+1}) + \|\mathbf{s}_b^t - \mathbf{s}_b^{t+1}\|_2}{B(T-1)}, \quad (8)$$

where $(\mathbf{R}_b^t | \mathbf{s}_b^t) = \mathbf{J}_b^t$, and we compute the relative angle of rotations $\text{ang}(\mathbf{R}_1, \mathbf{R}_2)$ by $\arccos((\text{tr}(\mathbf{R}_1 \mathbf{R}_2^T) - 1)/2)$.

3.4. Skeleton Generation

In this section, we describe how DreaMo generates a skeleton from the learned neural bones, and illustrate the




Figure 3. **Skeleton generation strategy.** Initially, we extract the rest-pose surface of the neural implicit 3D model using marching cubes. Subsequently, each vertex is assigned to a neural bone based on the maximum skinning weights. An edge between bones is established if there exists a sufficient vertex connection.

strategy in Figure 3. The skeleton defines the pairwise relationship between bones, indicating whether two bones have close interactions to jointly control a shared set of points on the SDF surface or the vertices of the mesh converted from the SDF. We assess such a relationship using the skinning weights, which define the contribution weights of the neural bones to each particular point in the space.

With the canonical space model M learned with the objectives described in Section 3.3, we can extract the rest-pose mesh by running marching cubes. We denote the mesh as a collection of vertices $\mathcal{V} = \{\nu_i \in \mathbb{R}^3\}$ and faces $\Lambda = \{\lambda_j\}$ where each element λ_j denotes a triangle face formed by three vertices in \mathcal{V} . Additionally, we extract the skinning weight of each vertex based on its coordinates and then identify the bone $b \in \mathcal{B}$ with the highest skinning weight to control that vertex. After iterating through all the vertices, we obtain clusters of vertices, each assigned with the corresponding bone ID. Our next target is to find the bone pairs that control a significant number of shared faces, which implies the two bones will deform a considerable number of common faces during articulation. We traverse all the face edges of the mesh. For each pair of bones (b_1, b_2) , we count the number of edges that connect two vertices assigned to these two bones. Finally, we decide whether each pair of bones is connected using such surface connectivity count and simple threshold to remove extreme outliers (e.g., two bones sharing only a single face).

4. Experimental Results

Implementation details of DreaMo. Our canonical space implicit model, warping models, neural bones, and global transformation network are all MLPs with residual connections similar to NeRF [33]. To model the fine-grain deformations, we learn another MLP to predict delta skinning weights and add it to the linear skinning weights before the softmax layer as in BANMo [59]. To obtain the SDS loss, we randomly sample a relative azimuth between $[-90^\circ, 90^\circ]$ and a relative elevation in $[-10^\circ, 45^\circ]$ as the

viewpoint condition of Zero-1-to-3 [29]. During volume rendering, we sample 64 points along the radiance for most of the losses, except the SDS loss only uses 16 samples to reduce memory consumption. Throughout the experiments, the weights of our loss terms $\mathcal{L}_{\text{recon}}$, \mathcal{L}_{sds} , \mathcal{L}_{cyc} , $\mathcal{L}_{\text{ncyc}}$, $\mathcal{L}_{\text{surf}}$, $\mathcal{L}_{\text{smooth}}$ are 10^{-1} , 10^{-4} , 10^{-2} , 1, 10^{-1} , and 10^{-2} , respectively. We reduce the frequency of computationally expensive losses to accelerate the training, specifically, we update $\mathcal{L}_{\text{ncyc}}$ every 3 iterations and \mathcal{L}_{sds} every 10 iterations. For the skeleton generation, we empirically set the threshold for surface connectivity count to 100.

Datasets. We collected 42 animal video clips with diverse species and insufficient view coverage from the Internet. To ensure the videos satisfy our challenging insufficient-coverage scenario, we compute and ensure the videos have low azimuth viewpoint coverage during the data collection. We compute the azimuth coverage by evenly dividing the azimuth of canonical space into 36 equally angle bins, finding the bins covered by any learned global transformation P_t (see Section 3.1), then determining the occupancy ratio. The average azimuth viewpoint coverage is 31%. Our video set contains 28 different animal species, and each video only contains one animal subject. The average duration of the videos is 15.7 seconds. We follow the data preprocessing protocol from our baseline BANMo. In particular, an intermediate step utilizes optical flow [56] to filter out frames with small motions. The final average number of video frames utilized by DreaMo and BANMo is 124.4.

Metrics. We consider two aspects: (a) consistency between the input video and the reconstructed 3D model, and (b) visual quality of the re-rendered images.

For the consistency evaluation, we only consider semantic consistency. In our problem setting, it is infeasible to hold out testing video frames from training for pixel-level reconstruction metrics, as the subject’s articulations in test frames are unknown to DreaMo and non-trivial to discover. It is difficult to find the correct object pose alignment between the canonical space model and the target in unseen video frames. To measure the semantic consistency, we use the image encoder of CLIP [40] to extract image features and compare the cosine similarity between the training and the rendered video frames. Specifically, we compute the cosine similarity in two different strategies: *exhaustive* that compares each rendered novel view with all the training video frames, and *per-time* that only compares the video frames at the same time step.

For quality assessment, we use Kernel Inception Distance (KID) [5, 34], which is a popular quality measurement metric in generative modeling research specialized for small sample sizes. For each time step, we render a novel view at the same time step, resulting in an image set at the same size as the number of training video frames. We calculate the KID between these two sets of images, and then report

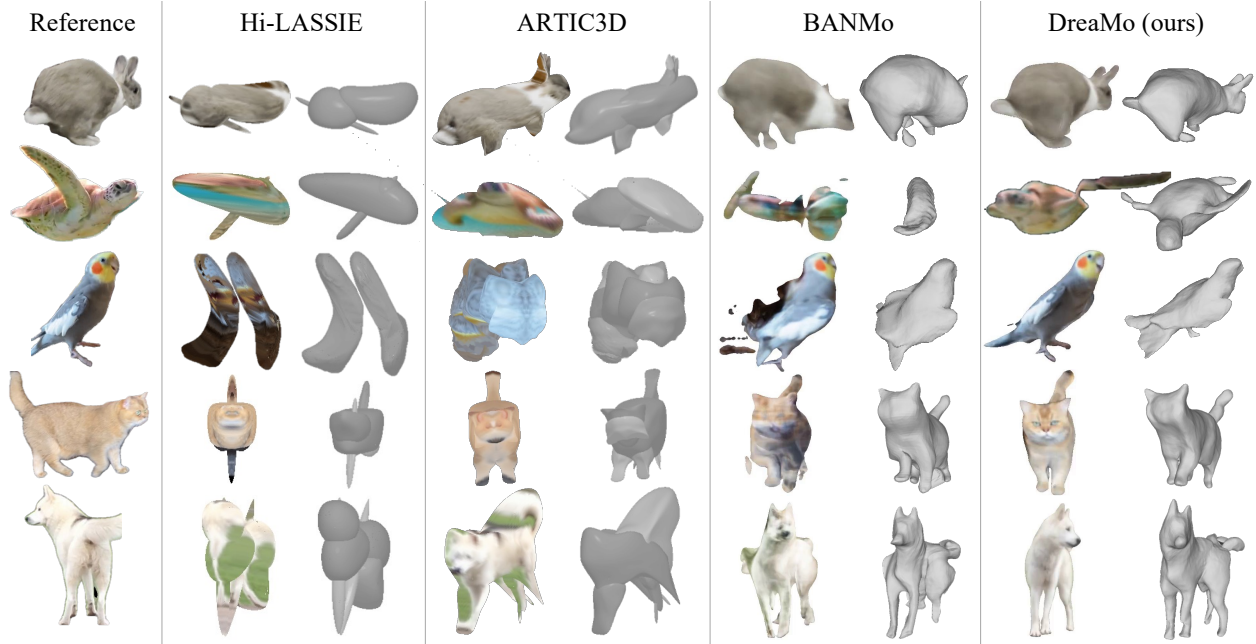


Figure 4. **3D reconstruction comparison among state-of-the-arts.** DreaMo generates more complete 3D shapes and maintains richer high-frequency details in the novel view rendering. For each method, we show the rendered novel-view image and the reconstructed shape.

both the mean and the standard deviation.

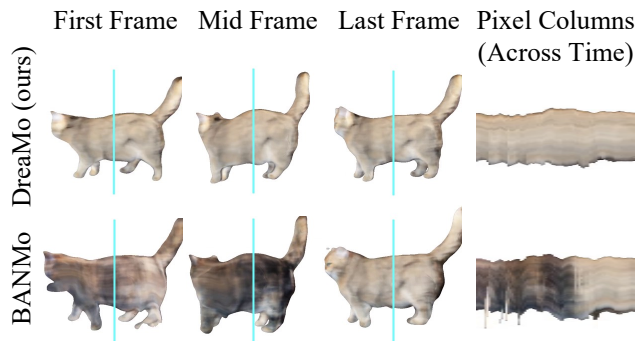


Figure 5. **DreaMo has better visual consistency.** We render the video from the same novel-view angle, retrieve a pixel column from each time step (cyan line), and aggregate them on the right. BANMo often shows a drastic color shift over time.

4.1. 3D Reconstruction

In Figure 4, we compare with state-of-the-art methods in articulated 3D shape reconstruction.

Video-based methods. Our most relevant baseline is BANMo [59], the state-of-the-art video-based 3D reconstruction method, which has shown outstanding quality when input videos have a very dense camera coverage. We use the latest implementation [2] from the author and ensure the fairness of comparisons by adding back the missing sinkhorn loss mentioned in the original paper of

Table 1. **DreaMo maintains better semantical consistency and better visual quality.** We measure the semantical consistency with CLIP similarity and quantify visual quality using KID. † denotes evaluated on a subset of 11 videos.

Method	CLIP (↑)		KID (↓)
	Exhaustive	Per-time	Mean ± Stddev
Hi-LASSIE [61]	0.742	0.744	0.0876 ± 0.0021
BANMo [59]	0.792	0.797	0.0576 ± 0.0021
DreaMo (ours)	0.813	0.817	0.0488 ± 0.0020
ARTIC3D† [62]	0.774	0.778	0.0822 ± 0.0036
DreaMo† (ours)	0.866	0.870	0.0350 ± 0.0024

BANMo [59]. In Figure 4, BANMo shows blurry reconstruction with fragmented geometry due to insufficient view coverage in certain regions. With the same setting, DreaMo successfully maintains clean and plausible geometry in the novel-view angles. In Figure 5, BANMo struggles to maintain appearance consistency across time in these low-coverage regions. In Figure 6, we show our regularizations result in better neural bone placements, which prevents the model from creating irregular shape artifacts. The quantitative comparisons in Table 1 also indicate our DreaMo maintains better semantic consistency from all angles while presenting better visual quality in KID measurement. Please refer to the supplementary material for more reconstruction results.

Image-based methods. Besides BANMo, we further com-

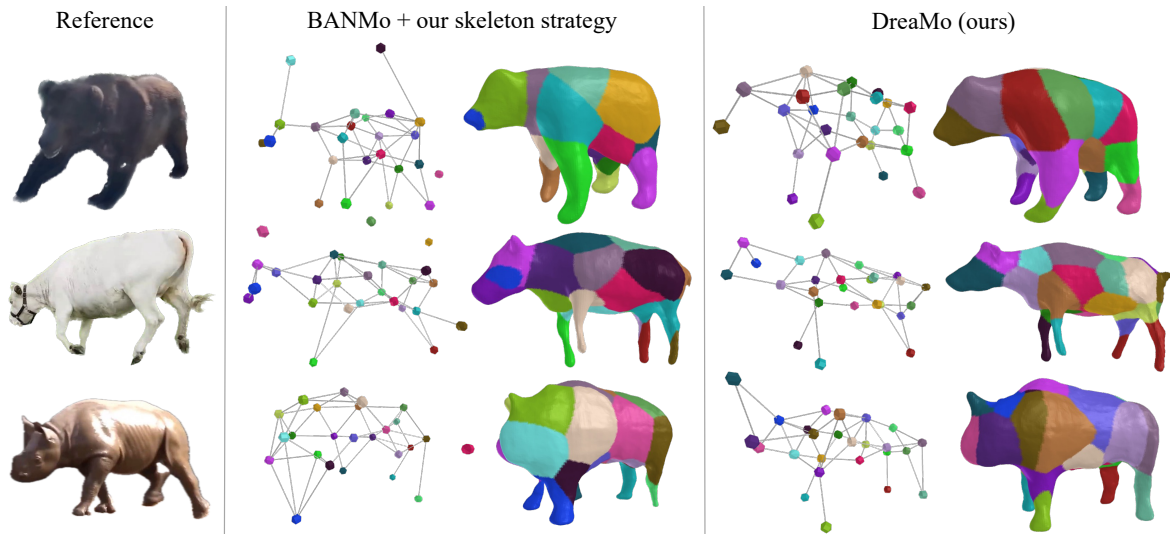


Figure 6. **DreaMo generates intuitive skeletons aligned with mesh.** We apply our skeleton generation strategy to both BANMo and DreaMo, then show the skeleton and the mesh converted from the canonical model. The face colors of the mesh denote the bone ID with the highest skinning weight. Our DreaMo better aligns with the 3D shape without leftover bones drifting far away from the object surface.

pare our DreaMo with LASSIE [61] and ARTIC3D [62], which are weakly relevant works in template-based articulated 3D-shape reconstruction from image collection of the same animal species. These methods consider fitting the appearance of an articulated template shape with images of the same animal species with diverse viewing angles. Despite not having the same level of visual diversity as an image collection, we put our best effort into examining these methods using video frames in replace of the image collection. We run Hi-LASSIE using the official codebase [1]. To satisfy the additional requirements of the skeleton initialization from Hi-LASSIE, for each video, we manually select a frame where most of the subject’s parts are visible and suitable for skeleton initialization. As of ARTIC3D, we sought the authors to run their algorithm on an 11-video subset from our dataset, considering the computational costs.

In Figure 4, both Hi-LASSIE and ARTIC3D often fail to solve the correct correspondence between the neural body parts and the subject in the video. Such a problem is reinforced in their part-based optimization when each part struggles to match the appearance of the reference, thus resulting in an unrealistic appearance as a whole in the end. In contrast, our DreaMo models the entire subject using a single neural implicit field, facilitates the model to solve all the body parts at once, and results in a more consistent shape. Moreover, DreaMo leverages optical flow across video frames to better understand the pixel correspondence from the limited data, while image-based methods do not consider such information, and integrating optical flow information into their framework is non-trivial. We also quantitatively compare with Hi-LASSIE and ARTIC3D in Table 1, our method consistently shows better semantic con-

sistency in CLIP score and visual quality in KID.

4.2. Skeleton Generation

We evaluate the skeletons generated by DreaMo against those produced by BANMo [59] in Figure 6. While BANMo [59] can produce reasonable object surfaces, its bone placement does not well represent the object shape, with some bones distributed far outside the object surfaces. In contrast, since DreaMo generates more accurate shapes and encourages neural bones to stay within the reconstructed surface, it results in detailed and interpretable skeletons capturing features such as limbs and heads. Please refer to the supplementary material for more practical usages of DreaMo’s output skeletons.

4.3. Articulating 3D Model

In Figure 7, we demonstrate the controllability of DreaMo by articulating the learned neural bones. Following the 3D model manipulation described in Section 3.1, we manually adjust the positions of the neural bones and transform the skin points (*i.e.* vertices of the mesh) to produce novel poses. With the accurately learned bone placement, skinning weights, and 3D shapes from DreaMo, the skin points can be reasonably transitioned corresponding to the movement of the neural bones, leading to plausible results in novel poses.

4.4. Ablation Study

We quantitatively analyze the contribution from different components of DreaMo in Table 2. The results show that each component contributes to the performance improvements in semantic consistency in the CLIP score and vi-




Figure 7. **Controlling DreaMo by manipulating generated skeletons.** We generate plausible 3D shapes of target objects in novel poses by controlling the generated skeletons. We manually adjust the bone positions and warp the corresponding skin, i.e., the mesh vertices in the figure, to demonstrate new poses of the articulated 3D objects.

Table 2. **Ablation study.** We show that every proposed component contributes to the final improvements in semantic consistency (CLIP) and visual quality (KID).

Method	CLIP (\uparrow)		KID (\downarrow)
	Exhaust.	Per-time	Mean \pm Stddev
No \mathcal{L}_{nyc} , $\mathcal{L}_{\text{smooth}}$, $\mathcal{L}_{\text{surf}}$	0.799	0.801	0.0542 \pm 0.0016
No $\mathcal{L}_{\text{smooth}}$, $\mathcal{L}_{\text{surf}}$	0.805	0.808	0.0507 \pm 0.0019
No \mathcal{L}_{nyc}	0.808	0.812	0.0535 \pm 0.0021
No $\mathcal{L}_{\text{smooth}}$	0.811	0.815	0.0512 \pm 0.0019
No $\mathcal{L}_{\text{surf}}$	0.807	0.810	0.0501 \pm 0.0020
No \mathcal{L}_{sds}	0.797	0.800	0.0539 \pm 0.0020
\mathcal{L}_{sds} for all params	0.794	0.796	0.0585 \pm 0.0022
DreaMo (ours)	0.813	0.817	0.0488 \pm 0.0020

sual quality in the KID metric. Among all components, the diffusion prior makes the most significant improvement by guiding DreaMo in updating articulation parameters to minimize \mathcal{L}_{sds} . Additionally, in Figure 8, we show that updating all parameters with \mathcal{L}_{sds} results in overly-smoothed textures. This effect may potentially stem from the randomness of the denoising diffusion process, which generates view-inconsistent results that hinder the model from optimizing detailed textures. Please refer to the supplementary material for an analysis of various SDS strategies.

The proposed regularization schemes also consistently improve the performance of DreaMo. These schemes help avoid placing the neural bones in empty space far from the object’s surface, or sudden discontinuous transitions between frames. These improvements also contribute to the interpretable skeletons, as demonstrated in the Section 4.2.




Figure 8. Naively updating all the trainable parameters with SDS hinders the high-frequency texture details.

5. Conclusion

We present DreaMo, a template-free framework to reconstruct plausible articulated 3D models from a single casual video with incomplete view coverage. To overcome the insufficient supervision in the unseen or low-coverage regions, DreaMo leverages the view-conditioned diffusion model to hallucinate and complete the 3D shape with a plausible and coherent appearance. Besides, we show the effectiveness of our proposed regularization schemes that improve the placement of the neural bones and reduce the irregular reconstruction artifacts. We further present a skeleton generation strategy to transform the learned neural bones and skinning weights into interpretable skeletons. Through extensive qualitative and quantitative experiments, we show DreaMo achieves state-of-the-art quality in articulated 3D shape reconstruction in our single video setting.

References

- [1] Hi-lassie. <https://github.com/google/hi-lassie>, 2023. 7

- [2] Lab4d. <https://github.com/lab4d-org/lab4d>, 2023. 6
- [3] Marc Badger, Yufu Wang, Adarsh Modh, Ammon Perkes, Nikos Kolotouros, Bernd G Pfrommer, Marc F Schmidt, and Kostas Daniilidis. 3d bird reconstruction: a dataset, model, and shape recovery from a single view. In *ECCV*, 2020. 2
- [4] Benjamin Biggs, Oliver Boyne, James Charles, Andrew Fitzgibbon, and Roberto Cipolla. Who left the dogs out? 3d animal reconstruction with expectation maximization in the loop. In *ECCV*, 2020. 2
- [5] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. In *ICLR*, 2018. 5
- [6] Christoph Bregler, Aaron Hertzmann, and Henning Biermann. Recovering non-rigid 3d shape from image streams. In *CVPR*, 2000. 2
- [7] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023. 2
- [8] Chen Gao, Ayush Saraf, Johannes Kopf, and Jia-Bin Huang. Dynamic view synthesis from dynamic monocular video. In *ICCV*, 2021. 2
- [9] Shubham Goel, Angjoo Kanazawa, and Jitendra Malik. Shape and viewpoint without keypoints. In *ECCV*, 2020. 2
- [10] Paulo FU Gotardo and Aleix M Martinez. Non-rigid structure from motion with complementary rank-3 spaces. In *CVPR*, 2011. 2
- [11] Alec Jacobson, Zhigang Deng, Ladislav Kavan, and JP Lewis. Skinning: Real-time shape deformation. In *ACM SIGGRAPH 2014 Courses*, 2014. 3
- [12] Yasamin Jafarian and Hyun Soo Park. Learning high fidelity depths of dressed humans by watching social media dance videos. In *CVPR*, 2021. 2
- [13] Tomas Jakab, Ruining Li, Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Farm3d: Learning articulated 3d animals by distilling 2d diffusion. In *3DV*, 2024. 2
- [14] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 1, 2
- [15] Angjoo Kanazawa, Shubham Tulsiani, Alexei A Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *ECCV*, 2018. 2
- [16] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 2
- [17] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. 2
- [18] Filippos Kokkinos and Iasonas Kokkinos. To the point: Correspondence-driven monocular 3d category reconstruction. *NeurIPS*, 2021. 2
- [19] Chen Kong and Simon Lucey. Deep non-rigid structure from motion. In *ICCV*, 2019. 2
- [20] Nilesh Kulkarni, Abhinav Gupta, David F Fouhey, and Shubham Tulsiani. Articulation-aware canonical surface mapping. In *CVPR*, 2020. 2
- [21] Suryansh Kumar. Non-rigid structure from motion: Prior-free factorization method revisited. In *WACV*, 2020. 2
- [22] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *ICCV*, 2017. 2
- [23] Verica Lazova, Eldar Insafutdinov, and Gerard Pons-Moll. 360-degree textures of people in clothing from a single image. In *3DV*, 2019. 2
- [24] Xueting Li, Sifei Liu, Shalini De Mello, Kihwan Kim, Xiaolong Wang, Ming-Hsuan Yang, and Jan Kautz. Online adaptation for consistent mesh reconstruction in the wild. *NeurIPS*, 2020. 2
- [25] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3d reconstruction via semantic consistency. In *ECCV*, 2020. 2
- [26] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023. 2
- [27] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *ICCV*, 2021. 2
- [28] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM TOG*, 2021. 2
- [29] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *ICCV*, 2023. 2, 4, 5
- [30] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *CVPR*, 2023. 2
- [31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM TOG*, 2015. 2
- [32] Gal Metzer, Elad Richardson, Or Patashnik, Raja Giryes, and Daniel Cohen-Or. Latent-nerf for shape-guided generation of 3d shapes and textures. In *CVPR*, 2023. 2
- [33] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021. 2, 3, 5
- [34] Anton Obukhov, Maximilian Seitzer, Po-Wei Wu, Semen Zhydenko, Jonathan Kyl, and Elvis Yu-Jing Lin. High-fidelity performance metrics for generative models in pytorch, 2020. 5
- [35] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021. 2
- [36] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 2

- [37] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 2
- [38] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023. 2, 4
- [39] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 2
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 5
- [41] Amit Raj, Srinivas Kaza, Ben Poole, Michael Niemeyer, Nataniel Ruiz, Ben Mildenhall, Shiran Zada, Kfir Aberman, Michael Rubinstein, Jonathan Barron, Yuanzhen Li, and Varun Jampani. Dreambooth3d: Subject-driven text-to-3d generation. In *ICCV*, 2023. 2
- [42] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 2
- [43] Elad Richardson, Gal Metzger, Yuval Alaluf, Raja Giryes, and Daniel Cohen-Or. Texture: Text-guided texturing of 3d shapes. *arXiv preprint arXiv:2302.01721*, 2023. 2
- [44] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [45] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 2022. 2
- [46] Uriel Singer, Shelly Sheynin, Adam Polyak, Oron Ashual, Iurii Makarov, Filippos Kokkinos, Naman Goyal, Andrea Vedaldi, Devi Parikh, Justin Johnson, et al. Text-to-4d dynamic scene generation. In *ICML*, 2023. 2
- [47] Chaoyue Song, Jiacheng Wei, Chuan Sheng Foo, Guosheng Lin, and Fayao Liu. Reacto: Reconstructing articulated objects from a single video. In *CVPR*, 2024. 2
- [48] Shih-Yang Su, Timur Bagautdinov, and Helge Rhodin. Npc: Neural point characters from video. In *ICCV*, 2023. 2
- [49] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. *NeurIPS*, 2021. 2
- [50] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023. 2, 4
- [51] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Vid2actor: Free-viewpoint animatable person synthesis from video in the wild. *arXiv preprint arXiv:2012.12884*, 2020. 2
- [52] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humanerf: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, 2022. 2
- [53] Shangzhe Wu, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. Dove: Learning deformable 3d objects by watching videos. *IJCV*, 2023. 2
- [54] Shangzhe Wu, Ruining Li, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. Magicpony: Learning articulated 3d animals in the wild. In *CVPR*, 2023. 2
- [55] Donglai Xiang, Hanbyul Joo, and Yaser Sheikh. Monocular total capture: Posing face, body, and hands in the wild. In *CVPR*, 2019. 2
- [56] Gengshan Yang and Deva Ramanan. Learning to segment rigid motions from two frames. In *CVPR*, 2021. 5
- [57] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Huiwen Chang, Deva Ramanan, William T Freeman, and Ce Liu. Lasr: Learning articulated shape reconstruction from a monocular video. In *CVPR*, 2021. 2
- [58] Gengshan Yang, Deqing Sun, Varun Jampani, Daniel Vlasic, Forrester Cole, Ce Liu, and Deva Ramanan. Viser: Video-specific surface embeddings for articulated 3d shape reconstruction. *NeurIPS*, 2021. 2
- [59] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *CVPR*, 2022. 2, 3, 5, 6, 7
- [60] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Lassie: Learning articulated shapes from sparse image ensemble via 3d part discovery. *NeurIPS*, 2022. 2
- [61] Chun-Han Yao, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Hi-lassie: High-fidelity articulated shape and skeleton discovery from sparse image ensemble. In *CVPR*, 2023. 2, 6, 7
- [62] Chun-Han Yao, Amit Raj, Wei-Chih Hung, Yuanzhen Li, Michael Rubinstein, Ming-Hsuan Yang, and Varun Jampani. Artic3d: Learning robust articulated 3d shapes from noisy web image collections. *arXiv preprint arXiv:2306.04619*, 2023. 6, 7
- [63] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 2021. 3
- [64] Yufei Ye, Shubham Tulsiani, and Abhinav Gupta. Shelf-supervised mesh prediction in the wild. In *CVPR*, 2021. 2
- [65] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yan-shun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM TOG*, 2021. 1, 2
- [66] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 2
- [67] Tingyang Zhang, Qingzhe Gao, Weiyu Li, Libin Liu, and Baoquan Chen. Bags: Building animatable gaussian splatting from a monocular video with diffusion priors. *arXiv preprint arXiv:2403.11427*, 2024. 2
- [68] Silvia Zuffi, Angjoo Kanazawa, Tanya Berger-Wolf, and Michael J Black. Three-d safari: Learning to estimate zebra pose, shape, and texture from images in the wild. In *ICCV*, 2019. 2

- [69] Silvia Zuffi, Angjoo Kanazawa, and Michael J Black. Lions and tigers and bears: Capturing non-rigid, 3d, articulated shape from images. In *CVPR*, 2018. 2
- [70] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *CVPR*, 2017. 2