

Ponymation: Learning Articulated 3D Animal Motions from Unlabeled Online Videos

Keqiang Sun^{1*}, Dor Litvak^{2,3*}, Yunzhi Zhang², Hongsheng Li¹,
Jiajun Wu^{2†}, and Shangzhe Wu^{2†}

¹CUHK MMLab ²Stanford University ³UT Austin
<https://keqiangsun.github.io/projects/ponymation>

Abstract. We introduce a new method for learning a generative model of articulated 3D animal motions from raw, unlabeled online videos. Unlike existing approaches for 3D motion synthesis, our model requires *no* pose annotations or parametric shape models for training; it learns purely from a collection of unlabeled web video clips, leveraging semantic correspondences distilled from self-supervised image features. At the core of our method is a video *Photo-Geometric Auto-Encoding* framework that decomposes each training video clip into a set of explicit geometric and photometric representations, including a rest-pose 3D shape, an articulated pose sequence, and texture, with the objective of re-rendering the input video via a differentiable renderer. This decomposition allows us to learn a generative model over the underlying articulated pose sequences akin to a Variational Auto-Encoding (VAE) formulation, but *without* requiring any external pose annotations. At inference time, we can generate new motion sequences by sampling from the learned motion VAE, and create plausible 4D animations of an animal automatically within seconds given a single input image.

Keywords: 3D animal motion · 4D generation · Unsupervised learning

1 Introduction

We share the planet with a wide variety of lively animals. Similarly to humans, they navigate and interact with the physical world, demonstrating various sophisticated motion patterns. In fact, the first film in history, “The Horse in Motion,” was a sequence of photographs that captured a galloping horse, created by Eadweard Muybridge in 1887 [50]. Films capture only sequences of 2D projections of 3D animal movements. Further modeling dynamic animals in 3D is not only useful for numerous mixed reality and content creation applications, but also provides computational tools for biologists to study animal behaviors.

While a lot of efforts have been invested in capturing and modeling 3D human motions using computer vision techniques, significantly less attention has been paid to animals. Existing learning-based approaches require an extensive amount

* Equal contribution. † Equal advising.

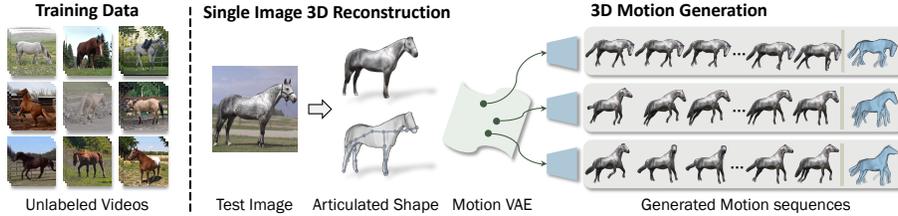


Fig. 1: Learning 3D Animal Motions from Unlabeled Online Videos. Given a collection of monocular videos of an animal category sourced from the Internet as training data, our method learns a *generative* model of the articulated 3D motions together with a monocular 3D reconstruction model, without relying on any shape templates or pose annotations. At inference time, the model generates new 3D motion sequences and turns a single test image in 4D animations fully automatically.

of 3D scans [48, 61, 62], parametric shape models [9, 32, 34, 59, 92], multi-view videos [20, 25, 44], or geometric annotations, such as keypoints [22, 23, 26, 57–59, 67], as supervision for training. Collecting large-scale 3D training data involves specialized capture devices and intensive labor, which can only be justified for specific objects, like humans, that are of utmost value in applications.

In this work, we would like to learn a *generative* model of the 3D motions of an animal category, which will allow us to sample new 3D motion sequences and generate 4D animations fully automatically within seconds in a feedforward fashion. Crucially, unlike existing 3D motion synthesis approaches on human bodies [22, 26, 31, 57, 58, 83, 95], we do *not* rely on explicit manual supervision for training, such as keypoints or template shapes. Instead, we propose to learn this 3D motion generative model purely from raw, unlabeled videos sourced from the Internet. This task is also different from video synthesis methods [11, 65, 76] that operate purely on 2D images. We would like to obtain an *explicit* 3D motion representation, in the form of a 3D mesh and a sequence of articulated 3D poses, which can easily facilitate downstream applications, including fine-grained controllable 3D animation and motion pattern analysis.

Learning 3D motions from unstructured online video collections is an extremely ill-posed task, as each video clip depicts only a short sequence of 2D projections of a *unique* 4D instance, with unique shape, appearance, motion, and viewpoint that are *not* assumed to reappear in another clip. This task, therefore, requires registering these unique video clips in a single canonical 3D model to learn a distribution of the underlying 3D motions of the animals. To address this challenge, we take advantage of recent advancements in self-supervised image representation learning [12], and distill semantic correspondences across different instances from self-supervised image features produced by a pre-trained DINO-ViT [12]. Furthermore, we assume a coarse description of the motion skeleton of the animal, *e.g.*, “quadruped,” which effectively constrains the space of deformation akin to Non-Rigid Structure-from-Motion [10] and provides a succinct representation for modeling the 3D motion.

Building on top of these insights, we design a video *Photo-Geometric Auto-Encoding* framework for learning 3D motion generative models from unlabeled

videos. At its core is a spatio-temporal transformer that automatically decomposes a video clip into a set of geometric and photometric factors, including a rest-pose 3D mesh, appearance, viewpoint, and a motion latent code that encapsulates the 3D motion of the instance. This motion latent code is then decoded into a sequence of articulated 3D poses, which are used to animate the rest-pose mesh and re-render a 2D video clip using a differentiable renderer. This allows us to train the entire model end-to-end like a “Variational Auto-Encoder” (VAE) over the space of articulated 3D motions, using only 2D image reconstruction losses on the RGB frames, DINO features, and object masks, with pseudo-ground-truth masks obtained from off-the-shelf detectors [37].

At inference time, we can generate new 3D motion sequences by sampling from the motion VAE latent space. If further given a single image of an animal, our model can reconstruct its articulated 3D shape and appearance in a feed-forward fashion, and generate 4D animations fully automatically within seconds.

To summarize, this paper makes several contributions:

- We propose a new method for learning a *generative* model of articulated 3D animal motions from *unlabeled* Internet videos, without any shape templates or pose annotations;
- We design a spatio-temporal transformer architecture that effectively extracts motion information from input video clips into a latent VAE;
- At inference time, the model generates diverse 3D motion sequences and turns a single image into 4D animations automatically in seconds;

2 Related Work

Learning 3D Animals from Image Collections. While modeling dynamic 3D objects traditionally requires motion capture markers or simultaneous multi-view captures [1, 17, 24], recent learning-based approaches have demonstrated the possibility of learning 3D deformable models simply from raw single-view image collections [33, 43, 47, 68, 78, 80, 90, 91]. Most of these methods require additional geometric supervision besides object masks for training, such as keypoint [33, 42] and viewpoint annotations [18, 54, 66], template shapes [21, 38, 39], semantic correspondences [30, 43, 78, 89, 90], and strong geometric assumptions like symmetries [79–81] and viewpoint distributions [14, 15, 52, 53, 63, 71, 72]. Among these, MagicPony [78] demonstrates impressive results in learning articulated 3D animals, such as horses, using only single-view images with object masks and self-supervised image features as training supervision. However, it reconstructs static images individually, ignoring the dynamic motions of the underlying 3D animals underneath those images. In this work, we focus on learning a generative model of 3D animal motions from videos instead of independent images.

Deformable Shapes from Monocular Videos. Reconstructing deformable shapes from monocular videos is a long-standing problem in computer vision. Early approaches with Non-Rigid Structure from Motion (NRSfM) reconstruct deformable shapes from 2D correspondences, by incorporating heavy constraints on the motion patterns [3, 10, 13, 16, 82]. DynamicFusion [51] further integrates additional depth information from depth sensors. NRSfM pipelines have recently

been revived with neural representations. In particular, LASR [84] and its follow-ups [81, 85–87] optimize deformable 3D shapes over a small set of monocular videos, leveraging 2D optical flows in a heavily engineered optimization procedure. DOVE [77] proposes a learning-based framework that learns a category-specific single-image 3D reconstruction model from a monocular video collection. Despite using video data for training, none of these approaches explicitly model the generative distribution of temporal motions of the objects.

Motion Analysis and Synthesis. Modeling motion patterns of dynamic objects has important applications for both behavior analysis and content generation, and is instrumental to our visual perception system [6]. Computational techniques have been used for decades to study and synthesize human motions [7, 56, 74]. In particular, recent works have explored learning generative models for 3D human motions [2, 22, 26, 35, 45, 49, 57, 58, 67], leveraging parametric human shape models, like SMPL [48], and large-scale human pose annotations [5, 29]. In comparison, much less effort is invested in modeling animal motions. Huang *et al.* [27] proposes a hierarchical motion learning framework for animals, but requires costly motion capture data and hardly generalizes to animals in the wild. To sidestep the collection of 3D data, BKinD [70] introduces a self-supervised method for discovering and tracking keypoints from videos, but is limited to a 2D representation. Such 2D keypoints could be lifted to 3D [35, 69], but this requires multi-view videos or ground-truth keypoints for training. Unlike these prior works, our motion learning framework does not require any pose annotations or multi-view videos for training, and is trained simply using raw monocular online videos. Recent success of image diffusion models has also led to promising generic 4D generation models [8, 46, 60, 93, 94]. However, the 3D motions generated by these models are still very limited in terms of quality and diversity, as shown in the comparisons in Section 4.2.

3 Method

Given a collection of raw video clips of an animal category, such as horses, our goal is to learn a generative model of its articulated 3D motions. This allows us to sample 3D motion sequences from a learned latent space, and generate 4D animations of a new animal instance automatically given only a single 2D image at test time. We train this model simply on raw online videos without relying on any external pose annotations. To do so, we design a video photo-geometric auto-encoding framework that decomposes each training video clip into a rest-pose 3D mesh, appearance, camera viewpoint as well as a sequence articulated 3D poses. This allows us to learn a *generative* model over the underlying articulated 3D pose sequences akin to a motion “Variational Auto-Encoder”, but simply using the objective of re-rendering the input frames with a differentiable renderer. Figure 2 gives an overview of the training pipeline.

3.1 Modeling Articulated 3D Animal Motions

Each video clip records a 2D image sequence $\{I_t\}_{t=1}^T$ of the underlying 3D animal motion from one camera trajectory. Since the dataset is obtained from

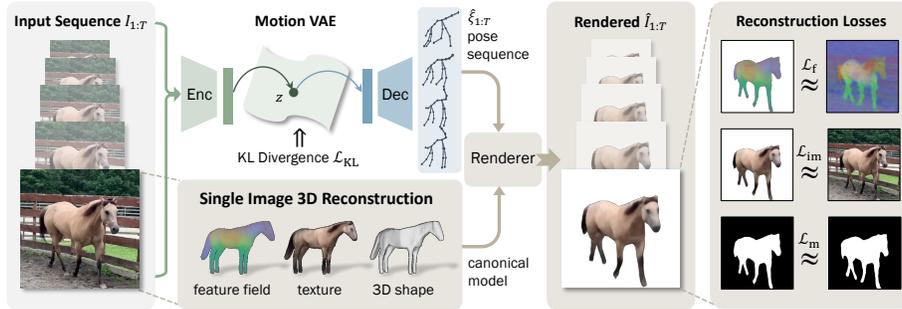


Fig. 2: Training Pipeline. Our method learns a generative model of articulated 3D motion sequences from a collection of unlabeled monocular videos. During training, the model encodes an input video sequence $I_{1:T}$ into a latent code z in the motion VAE, and decodes from it a sequence of articulated 3D poses $\hat{\xi}_{1:T}$. This pose sequence is used to animate the reconstructed 3D shape, allowing the full pipeline to be trained simply using image reconstruction losses with unsupervised image features and object masks obtained from off-the-shelf models, without any external pose annotations.

casually-recorded Internet videos, these training clips have diverse unique motion sequences. In order to learn the distribution of the underlying animal motions from such unstructured video collections, we first need to devise a 3D representation that registers these dynamic 2D sequences onto a canonical 3D model, factoring out the 3D motion of each video instance.

Drawing inspiration from prior work on 3D human motion synthesis [31, 48, 83, 95], we leverage a category-specific skinned model to represent the deformable 3D shape of the animals, and further learn the motion distribution over the articulations of its underlying skeleton. To this end, we follow MagicPony [78] and assume a coarse description of the skeleton, *e.g.*, “quadruped”.

Specifically, we represent the category-specific base 3D shape using a Signed Distance Function (SDF) parametrized by a coordinate Multi-Layer Perceptron (MLP), and extract an explicit mesh on the fly using Differentiable Marching Tetrahedron (DMTet) [64]. Let $V_{\text{base}} \in \mathbb{R}^{K \times 3}$ denote the list of K vertices, and the triangle faces are given by the triplets $F \subset \{1, \dots, K\}^3$. To model the slight shape variation of each animal instance in the canonical pose, we further learn an image-conditioned deformation field $f_{\Delta V}$ parametrized by another MLP that predicts small deformations of each vertex $\Delta V_{\text{ins},i} = f_{\Delta V}(V_{\text{base},i}, \phi)$, where $\phi = f_{\phi}(I)$ is a feature vector obtained from an image I using a pre-trained DINO-ViT [12], and $i \in \{1, \dots, K\}$ denotes the vertex index. Both base shape V_{base} and the instance deformation ΔV_{ins} are enforced to be bilaterally symmetric about yz -plane by mirroring the query locations in the underlying MLPs.

To account for the temporal motions driven by the underlying bone structure, we then instantiate a quadrupedal skeleton in this instance shape using a simple heuristic: a chain of bones going through the two farthest end points along z -axis, and four legs branching out from the body bone to the lowest point in each xz -quadrant. The motion sequence is thus parametrized by a sequence of articulated poses $\xi = \{\xi_t\}_{t=1}^T$, where each pose ξ_t at a timestamp t consists

of a rigid pose $\xi_{t,1} \in SE(3)$ w.r.t. an identity camera pose and the rotation $\xi_{t,b} \in SO(3)$ of each bone $b = 2, \dots, B$ in the skeleton. These articulated poses are applied to the instance mesh V_{ins} to obtain the final posed shape sequence using the widely-used linear blend skinning $g(V_{\text{ins}}, \xi_t)$ [48]. More details are included in the supplementary material.

The appearance of the instance is modeled using a texture field parametrized by an MLP $f_a(\mathbf{x}, \phi) \in [0, 1]^3$ where \mathbf{x} is a 3D location. We then render the posed mesh sequence into a sequence of RGB images using deferred mesh rendering [78], querying f_a at the corresponding 3D locations of the pixels after rasterization.

In the following, we explain the learning formulation to learn the individual components, including V_{base} , $f_{\Delta V}$, f_a , and most importantly, a generative model f_ξ over the motion sequences ξ , purely from an unstructured video collection without external pose annotations.

3.2 Video Photo-Geometric Auto-Encoding

Unlike human motion synthesis, we do not have access to large-scale, high-quality 3D captures or pose annotations for most animal species. Hence, we must instead learn from raw Internet videos, which poses significant challenges. To this end, we design a video *Photo-Geometric Auto-Encoding* framework that deconstructs each training clip into the explicit photometric and geometric factors described in Section 3.1, and train the entire pipeline using the objective of re-rendering the video. At the center of this video auto-encoding pipeline is a generative model of articulated motion sequences, akin to a ‘‘Variational Auto-Encoder’’ (VAE), but learned purely from raw RGB frames. This is very different from simply training a conventional VAE directly in the pose sequence space, which would require explicit pose annotations in the first place.

Video Encoding. To predict the instance shape deformation ΔV_{ins} and appearance of the object, we extract a feature vector ϕ_t for each frame of the video using a pre-trained DINO-ViT [12] with frozen weights, as mentioned previously. We assume the instance shape and appearance remain the same throughout the video, and hence take the average image features across all frames, denoted as $\bar{\phi}$, when querying the MLPs, $f_{\Delta V}$ and f_a .

In order to extract the motion information more effectively from the input video clip, we design a pair of spatial and temporal transformer-based motion encoders, E_s and E_t , that aggregate a set of bone-specific local features first spatially across each frame and then temporally across the entire sequence, eventually obtaining the distribution parameters $\hat{\mu}$ and $\hat{\Sigma}$ of the motion latent VAE.

Specifically, given each frame I_t in the input clip, we first construct a bone-specific feature descriptor $\nu_{t,b} = (\phi_t, \Phi_t(\mathbf{u}_{t,b}), b, \mathbf{J}_b, \mathbf{u}_{t,b})$ for each bone $b = 2, \dots, B$ and each timestamp t . Here, ϕ_t denotes the same global image feature as before. \mathbf{J}_b denotes the 3D location of the center of the bone b at rest-pose, which projects to the pixel location $\mathbf{u}_{t,b}$ in the image space, given the rigid pose $\hat{\xi}_{t,1}$ predicted separately. In addition to the global feature ϕ_t , we also sample an auxiliary bone-specific local feature vector $\Phi_t(\mathbf{u}_{t,b})$ from the DINO-ViT key token map Φ_t at the projected pixel location $\mathbf{u}_{t,b}$.

The spatial transformer encoder E_s then fuses these bone-specific feature descriptors $\{\nu_{t,b}\}_{b=2}^B$ into a single feature vector $\nu_{t,*}$ summarizing the articulated pose of the animal in each frame t :

$$\nu_{t,*} = E_s(\nu_{t,2}, \dots, \nu_{t,B}). \quad (1)$$

In practice, we prepend a learnable token to the list of descriptors, and take the first output token of the transformer as the pose feature $\nu_{t,*}$. We call this E_s a *spatial* transformer as it extracts the spatial geometric features in each input frame that capture the pose information, conditioned on the given skeleton.

Next, we design a second *temporal* transformer encoder E_t , inspired by [57], which operates along the temporal dimension and maps the entire sequence of pose features $\{\nu_{t,*}\}_{t=1}^T$ into the motion latent space. Similarly to the E_s , E_t fuses the pose feature sequence to predict the VAE distribution parameters:

$$(\hat{\mu}, \hat{\Sigma}) = E_t(\nu_{1,*}, \dots, \nu_{T,*}). \quad (2)$$

Using the reparametrization trick [36], we then sample a latent code from the Gaussian distribution $z \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma})$, which will be decoded into a sequence of articulated poses $\{\hat{\xi}_t\}_{t=1}^T$ characterizing the 3D motion of the animal in the clip.

Motion Decoding. Symmetric to the motion encoders, the motion decoder also consists of a temporal decoder D_t that first decodes z into a sequence of pose features $\{z_t\}_{t=1}^T$, and a spatial decoder D_s that further decodes each pose feature z_t to a set of bone rotations $\{\hat{\xi}_{t,b}\}_{b=2}^B$.

Specifically, we query the temporal transformer decoder D_t with a sequence of timestamps \mathcal{T} , and use z as both the key token and the value token to obtain a sequence of pose features:

$$(z_1, \dots, z_T) = D_t(\mathcal{T}, z), \quad \mathcal{T} = (1, \dots, T). \quad (3)$$

Similarly, given each pose feature z_t , we then query the spatial transformer decoder D_s with a sequence of bone indices \mathcal{B} to produce the bone rotations:

$$(\hat{\xi}_{t,2}, \dots, \hat{\xi}_{t,B}) = D_s(\mathcal{B}, z_t), \quad \mathcal{B} = (2, \dots, B). \quad (4)$$

In practice, the rigid pose $\hat{\xi}_{t,1}$ is predicted by a separate network and is not modeled by this motion VAE, since it is entangled with arbitrary camera motions that are difficult to disentangle in dynamic scenes.

We then deform the predicted instance mesh \hat{V}_{ins} using these articulated pose sequence $\{\hat{\xi}_t\}_{t=1}^T$ with the skinning equation $\hat{V}_t = g(\hat{V}_{\text{ins}}, \hat{\xi}_t)$, and render the RGB frames $\{\hat{I}_t\}_{t=1}^T$ and masks $\{\hat{M}_t\}_{t=1}^T$ using a differentiable renderer [41].

3.3 Learning Formulation

Video Re-rendering Losses. We train the entire model by minimizing the reconstruction losses on the object masks \hat{M}_t and RGB frames \hat{I}_t :

$$\mathcal{L}_{\text{m},t} = \|\hat{M}_t - M_t\|_2^2 + \lambda_{\text{dt}} \|\hat{M}_t \odot \text{dt}(M_t)\|_1, \quad \mathcal{L}_{\text{im},t} = \|\hat{M}_t \odot (\hat{I}_t - I_t)\|_1, \quad (5)$$

where distance transform $\text{dt}(\cdot)$ is used in the second term of the mask loss with a weight λ_{dt} for more effective gradients [33, 78, 79], and \odot denotes the Hadamard product. The RGB loss is only computed inside the intersection of the predicted and ground-truth masks $\tilde{M}_t = \hat{M}_t \odot M_t$. To exploit the temporal consistency of the motion in the videos, we further enforce a temporal smoothness constraint between the predicted poses $\hat{\xi}_t$ of consecutive frames: $\mathcal{R}_{\text{temp}} = \sum_{t=2}^T \|\hat{\xi}_t - \hat{\xi}_{t-1}\|_2^2$. We also inherit the multi-hypothesis viewpoint prediction mechanism with the hypothesis loss \mathcal{L}_{hyp} and the shape regularizers $\mathcal{R}_{\text{shape}} = \lambda_{\text{Eik}}\mathcal{R}_{\text{Eik}} + \lambda_{\text{art}}\mathcal{R}_{\text{art}} + \lambda_{\text{def}}\mathcal{R}_{\text{def}}$ [78] with balancing weights λ 's, which include the Eikonal constraint \mathcal{R}_{Eik} on the SDF MLP for the base shape, and magnitude regularizers \mathcal{R}_{art} on the bone rotations $\hat{\xi}_{2:B}$ and \mathcal{R}_{def} on the vertex deformations ΔV_{ins} .

Semantic Correspondences. Instead of relying on external pose annotations or prior shape models to learn the 3D model from monocular videos, we seek a much cheaper alternative solution for establishing correspondences across different instances. We distill semantic correspondences from self-supervised image features, such as DINO [12]. As shown in prior work [4, 78, 90], after a simple PCA reduction, these image features reveal robust part-level correspondences across different instances with varying poses and appearance. To exploit these correspondences, we additionally optimize a feature field in the canonical space using a coordinate MLP $\psi(\mathbf{x}) \in \mathbb{R}^D$, which is rendered into a 2D feature image $\hat{\Phi}_t \in \mathbb{R}^{D \times H \times W}$ given the posed mesh \hat{V}_t , with the same procedure as rendering the appearance of the object described above. We then encourage this rendered feature map $\hat{\Phi}_t$ to match the feature map Φ'_t pre-extracted from the input frame I_t using DINO-ViT with PCA reduction: $\mathcal{L}_{\text{feat},t} = \|\tilde{M}_t \odot (\hat{\Phi}_t - \Phi'_t)\|_2^2$. Intuitively, this enforces the model to establish correspondences across all training video instances through the same canonical feature field, hence disentangling the shape and pose in each monocular frame.

Motion VAE. Similarly to the conventional VAE, we also minimize the Kullback–Leibler (KL) divergence between the learned motion latent distribution and a standard Gaussian distribution:

$$\mathcal{L}_{\text{KL}} = \sum_i -\frac{1}{2} (\log \sigma_i - \sigma_i - \mu_i^2 + 1), \quad (6)$$

where μ_i and σ_i^2 are elements of the predicted distribution parameters $\hat{\mu}$ and $\hat{\Sigma}$.

Training Schedule. As learning 3D articulated motions from unstructured video clips without labels is extremely ill-posed, we devise a two-stage schedule for robust and efficient training. In the *first* stage, we pre-train the monocular 3D reconstruction model using a single-image pose predictor $\tilde{\xi}_t = f_{\tilde{\xi}}^{\text{sim}}(\phi_t)$. Inspired by but unlike [78], we train this model to re-render entire video clips with the temporal smoothness constraint $\mathcal{R}_{\text{temp}}$ and temporal feature averaging $\bar{\phi}$, rather than independent images. The total loss in the first stage is given by:

$$\mathcal{L}_{\text{vid}} = \sum_{t=1}^T (\mathcal{L}_{\text{recon},t} + \lambda_h \mathcal{L}_{\text{hyp},t} + \lambda_s \mathcal{R}_{\text{shape},t}) + \lambda_t \mathcal{R}_{\text{temp}}, \quad (7)$$

Table 1: Statistics of the *AnimalMotion* Dataset. We collect a new animal video dataset containing a total of 82.6k frames for 4 different animal species.

Category	# Sequences	Total Length	# Frames
Horse	640	28'09"	50,682
Zebra	47	5'27"	9,822
Giraffe	60	4'52"	8,768
Cow	69	7'25"	13,359
Total	816	45'54"	82,631

where $\mathcal{L}_{\text{recon},t} = \mathcal{L}_{\text{im},t} + \lambda_m \mathcal{L}_{\text{m},t} + \lambda_f \mathcal{L}_{\text{feat},t}$ summarizes the reconstruction losses on each frame. After this stage, we obtain an accurate monocular 3D reconstruction model, which outperforms the baseline [78] as shown in Table 4, largely owing to the training on videos instead of independent images. More importantly, the model has now learned a reasonable space of articulated poses, on top of which learning a motion generative model is much more efficient.

In the *second* stage, we replace the monocular pose predictor f_{ξ}^{sin} with the spatio-temporal transformer-based motion VAE f_{ξ} detailed in Section 3.2, which encodes the entire video clip and generates the entire sequence of articulated poses at once. Empirically, training the motion VAE from scratch with an expensive rendering step in the loop is inefficient. To facilitate training efficiency, we recycle pose predictions $\hat{\xi}_t$ from the first stage to guide the predictions of the VAE decoder $\hat{\xi}_t$ using a teacher loss $\mathcal{L}_{\text{teacher}} = \sum_{t=1}^T \|\hat{\xi}_t - \xi_t\|_2^2$. The final training objective for the second stage is thus:

$$\mathcal{L} = \mathcal{L}_{\text{vid}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} + \lambda_{\text{teacher}} \mathcal{L}_{\text{teacher}}. \quad (8)$$

3D Motion Generation. During inference time, we can generate diverse 3D motion sequences by sampling from the learned motion VAE latent space. Furthermore, when given a single 2D image of a new animal instance unseen at training, our model can reconstruct its 3D shape and appearance in a feed-forward manner, and generate 4D animations fully automatically within a few seconds, as illustrated in Figure 3.

4 Experiments

4.1 Experimental Setup

Datasets. To train our model, we collected an **AnimalMotion** dataset consisting of video clips of several quadruped animal categories extracted from the Internet. The statistics of the dataset are summarized in Table 1. As pre-processing, we first detect and segment the animal instances in the videos using the off-the-shelf segmentation model of PointRend [37]. To remove occlusion between different instances, we calculated the extent of mask overlap in each frame and exclude crops where two or more masks overlap with each other. We further apply a smoothing kernel to the sequence of bounding boxes to avoid jittering. The non-occluded instances are then cropped and resized to 256×256 . The



Fig. 3: 3D Motion Generation and Animation. During test time, our model generates plausible 3D motion sequences by sampling from the learned motion VAE. It can also reconstruct articulated 3D shapes from a single 2D image in feed-forward fashion, and generate 4D animations fully automatically within seconds. Within each gray box on the right, the first row shows textured animation, and the second row visualizes the corresponding 3D shapes with the generated bone articulations.

original videos are all at 30fps. To ensure sufficient motion in each sequence, we remove frames with minimal motion, measured by the magnitude of optical flows within the instance mask estimated from RAFT [73]. To conduct quantitative evaluations and comparisons, we also use PASCAL VOC [19] which contains 108 images of horses, and APT-36K [88] which contains 81 video clips of horses, each consisting of 15 frames. Both datasets provide 2D keypoint annotations for each animal in the image, allowing us to evaluate the geometric accuracy of the reconstructed shapes and generated motions.

Implementation Details. The encoders and decoders of the motion VAE model (E_s , E_t , D_s , D_t) from Section 3.2 are implemented as stacked trans-

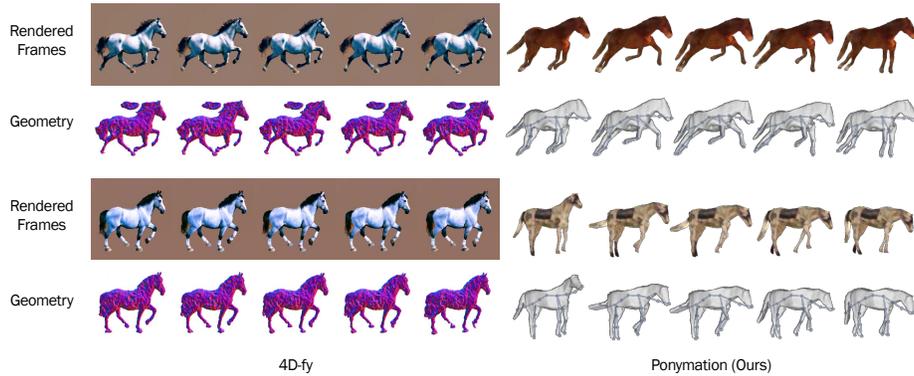


Fig. 4: 4D Generation Comparisons. We compare with 4D-fy [8], a recent text-to-4D generation method distilling from 2D diffusion. Despite heavy prompt engineering and a lengthy training time (12 hours), 4D-fy still fails to produce noticeable motion, whereas our model generates diverse motion sequences in a feed-forward pass within a few seconds, with much better 3D geometry.

formers [75] with 4 transformer blocks and a latent dimension of 256. We use a sinusoidal function for positional encoding following [57]. For the remaining architectures, we base our implementation on top of [78]. We train the model for 120 epochs for the first stage, which takes roughly 10 hours on 8 A6000 GPUs, and another 180 epochs for the second stage, which takes another 48 hours. We use a sequence length of $T = 10$ for training. During inference, we can generate longer sequences by connecting multiple samples and optimizing transition latent codes for smooth interpolation. For visualization, following prior work [78], we finetune (only) the appearance network f_a for 100 iterations on each test image, taking less than 10 seconds, as the model struggles to predict detailed texture in a single feedforward pass. More details are included in the sup. mat.

4.2 3D Motion Generation

Qualitative Results. After training, we can generate 3D motion sequences by sampling the motion latent space VAE, and render 4D animations with the textured mesh reconstructed from a single 2D image, as shown in Figure 3. It also generalizes to horse-like artifacts, such as carousel horses, which the model has never seen during training. The model can be trained on a wide range of animal species besides horses, including giraffes, zebras and cows, capturing category-specific prior distributions of 3D motions, as shown in Figure 5. Because the datasets for these categories are limited in size and diversity, as in [78], in the first stage of the training, we fine-tune from the model trained on horses. Additional animation results are provided in the supplementary video.

Comparison with Existing Methods. Our method is the first to learn a generative model of 3D animal motions from raw videos without pose annotations or prior shape models. We compare with one of the most recent 4D generative models, 4D-fy [8], which has publicly released code. Specifically, we

Table 2: Quantitative Comparison with State-of-the-Art Motion Generative Models.

	Motion Strength	User Preference
4D-fy [8]	0.29	112 (17.0%)
Ponymation (Ours)	4.66	548 (83.0%)

provide the model with a list of prompts, which are enriched by ChatGPT [55] from a list of basic prompts describing horse motions, such as “a horse is running/walking/jumping/eating”¹. We generate 20 4D instances from 4D-fy, and 20 from our method (without text condition). Note that it takes 12 hours to generate one 4D-fy instance on one GPU, whereas our model generates 4D animations within a few seconds in a single forward pass. We first compute the Motion Strength to assess the motion magnitude of the generated videos. We use Flowformer [28] to estimate optical flow strengths between consecutive frames of a generated video, and then compute the average of the largest 5% optical flows as the Motion Strength. We present them in random pairs side by side to 33 participants, and ask them to select one that shows “a more plausible 3D horse motion sequence”. As reported in Table 2, users preferred the 4D instances generated by our method over 4D-fy 83.0% of the time. We show a visual comparison in Figure 4. Notably, 4D-fy produces nearly static animals without perceptible motions despite heavy prompt engineering, whereas our method generates much more plausible motion sequences.

Quantitative Evaluation. Further assessing the quality of the generated 3D motions quantitatively is difficult due to the lack of (1) ground-truth measurements of 3D animal motions, and (2) robust evaluation metrics for generative models. To evaluate and compare different variants of our model, we design a new metric, bi-directional Motion Chamfer Distance (**MCD**), computed between a set of generated motion sequences projected to 2D image space and a set of 2D keypoint sequences annotated from videos in APT-36K [88]. Since the skeleton automatically discovered by our model is different from the 17 keypoints annotated in APT-36K, we first perform 3D reconstruction on all the images in APT-36K, and optimize a linear transformation that maps the 2D projections of the predicted 3D joints to the annotated 2D keypoints following [33]. To compute MCD, we generate 1,400 random motion sequences by sampling from the learned motion VAE, each consisting of 10 frames of 3D articulated poses. We then project these generated 3D poses to 2D using the viewpoints estimated from APT-36K, and apply the previously optimized transformation to align with the annotated keypoints. For each annotated keypoint *sequence* in the test set, we find the closest generated motion *sequence* measured by keypoint MSE averaged across all frames, and vice versa for each generated sequence. We then compute MCD based on the MSE between the closest sequence pairs. In essence, MCD measures the fidelity of generated motions by comparing the sampled distribution to that of the real motion sequences annotated from videos. Table 3 compares the results of our final model with two ablated variants.

¹ The full list of prompts are included in the supplementary material.

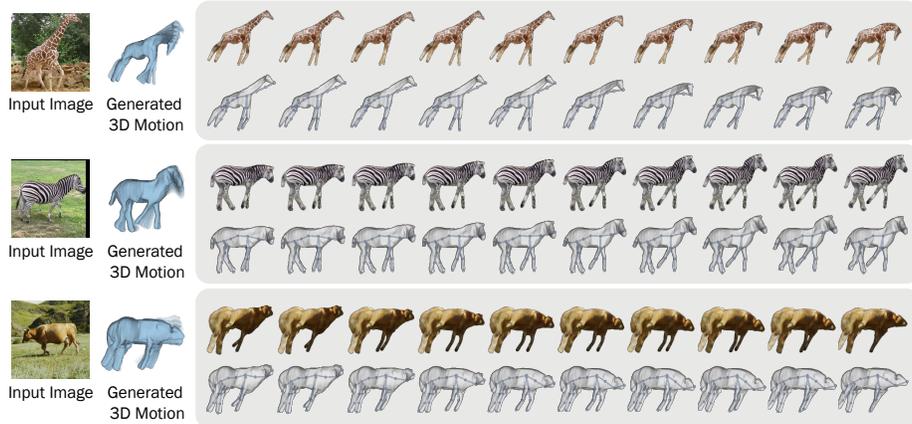


Fig. 5: 3D Motion Generation Results on More Species. Our method can be trained on various animal species, such as corws, zebras, and giraffes illustrated here. The model learns to generate 3D motions and generate plausible motion sequences specific to the animal species, such as the generated neck motion in the first example which is more common in giraffes than others.

Table 3: Motion Chamfer Distance (MCD) on APT-36K [88] for Motion Generation Evaluation. MP: Magicpony, AM: AnimalMotion dataset, TS: temporal smoothness.

Experiment	MCD ↓
MP + VAE	38.77
MP + VAE + AM	38.12
MP + VAE + AM + TS (final)	38.03

4.3 Single-Image 3D Reconstruction

We also quantitatively evaluate the monocular 3D reconstruction results of our model and compare with existing methods [39, 40, 43, 78]. For this purpose, we use PASCAL [19], a widely used benchmarking dataset for 3D reconstruction, as well as the aforementioned APT-36K [88] dataset, both of which come with 2D keypoint annotations. We compute the commonly used keypoint transfer metric measured by Percentage of Correct Keypoints (**PCK**) [33, 43, 78]. Specifically, given a set of annotated visible 2D keypoints on a source image, we identify the closest vertices on the reconstructed 3D mesh, and then project those 3D vertices onto the target 2D image. We calculate the percentage of the re-projected keypoints that land within a small distance from the annotated keypoints in the target image. This margin is set to be 0.1 of the image size following prior work [33, 43, 78]. Another commonly used metric is Mask Intersection over Union (**MIoU**) between the rendered and ground-truth masks, which measures the reconstruction quality in terms of projected 2D silhouettes. In addition, since APT-36K [88] provides keypoint annotations on video sequences, we also measure the temporal consistency across the reconstructions along the video sequences using

Table 4: Comparison of Monocular 3D Reconstruction Results with Different Methods on PASCAL [19] and APT-36K [88]. Our method achieves superior reconstruction accuracy compared to the existing methods, including the recent MagicPony baseline [78].

Method	PASCAL [19]		APT-36K [88]	
	PCK \uparrow	Mask IoU \uparrow	PCK \uparrow	Vel. Err. \downarrow
CSM [40]	31.2%	-	-	-
UMR [43]	24.4%	-	-	-
A-CSM [39]	32.9%	-	-	-
MagicPony [78]	42.8%	64.1%	53.9%	57.3%
Ponymation (ours)	48.0%	71.8%	59.9%	49.1%

a **Velocity Error**, computed as $\frac{1}{T} \sum_{t=1}^T \|\hat{\delta}_t - \delta_t\|/\delta_t$, where $\hat{\delta}_t$ and δ_t are the keypoint displacements between consecutive frame for predicted and GT pose sequences respectively. As the predicted poses are different from the GT keypoints, we use the same procedure described in Section 4.2 to optimize a linear mapping from the predicted poses to the GT keypoints for each method.

The results are summarized in Table 4. The results of MagicPony [78] are computed using the publicly released code and models, and the results of other baselines are taken from A-CSM [39]. Our model outperforms all previous methods. In particular, compared to the MagicPony baseline, our model achieves considerable improvement by learning from videos instead of individual images.

Additional ablation studies on the architecture design, discussions on limitations and more visualizations are included the supplementary material.

5 Conclusions

We have presented a new method for learning generative models of articulated 3D animal motions from raw Internet videos, without relying on any pose annotations or shape templates. To this end, we have proposed a video photo-geometric auto-encoding framework that automatically learns to decompose RGB videos into the underlying 3D shape, articulated motion, and object appearance, simply with the objective of re-rendering the videos. At the core of this pipeline is a transformer-based architecture that effectively extracts the temporal and spatial structure of the video clip into a latent motion VAE, which enables sampling at inference time to generate new 3D motion sequences. Experimental results show that the proposed method learns a reasonable distribution of 3D animal motions for several animal categories. This allows us to instantly turn a single 2D image into 4D animations in a fully automatic fashion, enabling promising downstream applications in game design and movie production.

Acknowledgments. We thank Zizhang Li, Feng Qiu, and Ruining Li for insightful discussions. The work is in part supported by the Stanford Institute for Human-Centered AI (HAI) and Samsung.

References

1. de Aguiar, E., Stoll, C., Theobalt, C., Ahmed, N., Seidel, H.P., Thrun, S.: Performance capture from sparse multi-view video. *ACM TOG* (2008) [3](#)
2. Ahn, H., Ha, T., Choi, Y., Yoo, H., Oh, S.: Text2Action: Generative adversarial synthesis from language to action. In: *ICRA* (2018) [4](#)
3. Akhter, I., Sheikh, Y., Khan, S., Kanade, T.: Nonrigid structure from motion in trajectory space. In: *NeurIPS* (2008) [3](#)
4. Amir, S., Gandselman, Y., Bagon, S., Dekel, T.: Deep vit features as dense visual descriptors. In: *ECCV Workshop on What is Motion For?* (2022) [8](#)
5. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2d human pose estimation: New benchmark and state of the art analysis. In: *CVPR* (2014) [4](#)
6. Badler, N.: Temporal Scene Analysis: Conceptual Descriptions of Object Movements. Ph.D. thesis, Queensland University of Technology (1975) [4](#)
7. Badler, N.I., Phillips, C.B., Webber, B.L.: *Simulating Humans: Computer Graphics, Animation, and Control*. Oxford University Press (09 1993) [4](#)
8. Bahmani, S., Skorokhodov, I., Rong, V., Wetzstein, G., Guibas, L., Wonka, P., Tulyakov, S., Park, J.J., Tagliasacchi, A., Lindell, D.B.: 4D-fy: Text-to-4d generation using hybrid score distillation sampling. In: *CVPR* (2024) [4](#), [11](#), [12](#)
9. Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In: *ECCV* (2016) [2](#)
10. Bregler, C., Hertzmann, A., Biermann, H.: Recovering non-rigid 3d shape from image streams. In: *CVPR* (2000) [2](#), [3](#)
11. Brooks, T., Peebles, B., Holmes, C., DePue, W., Guo, Y., Jing, L., Schnurr, D., Taylor, J., Luhman, T., Luhman, E., Ng, C., Wang, R., Ramesh, A.: Video generation models as world simulators (2024), <https://openai.com/research/video-generation-models-as-world-simulators> [2](#)
12. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: *ICCV* (2021) [2](#), [5](#), [6](#), [8](#)
13. Cashman, T.J., Fitzgibbon, A.W.: What shape are dolphins? building 3d morphable models from 2d images. *IEEE TPAMI* (2012) [3](#)
14. Chan, E., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-GAN: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: *CVPR* (2021) [3](#)
15. Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., Karras, T., Wetzstein, G.: Efficient geometry-aware 3D generative adversarial networks. In: *CVPR* (2022) [3](#)
16. Dai, Y., Li, H., He, M.: A simple prior-free method for non-rigid structure-from-motion factorization. In: *CVPR* (2012) [3](#)
17. Debevec, P.: The light stages and their applications to photoreal digital actors. In: *SIGGRAPH Asia* (2012) [3](#)
18. Duggal, S., Pathak, D.: Topologically-aware deformation fields for single-view 3d reconstruction. *CVPR* (2022) [3](#)
19. Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. *IJCV* (2015) [10](#), [13](#), [14](#)
20. Gao, X., Yang, J., Kim, J., Peng, S., Liu, Z., Tong, X.: Mps-nerf: Generalizable 3d human rendering from multiview images. *IEEE TPAMI* (2022) [2](#)

21. Goel, S., Kanazawa, A., Malik, J.: Shape and viewpoints without keypoints. In: ECCV (2020) [3](#)
22. Guo, C., Zuo, X., Wang, S., Zou, S., Sun, Q., Deng, A., Gong, M., Cheng, L.: Action2motion: Conditioned generation of 3d human motions. In: ACM MM (2020) [2](#), [4](#)
23. Habibie, I., Holden, D., Schwarz, J., Yearsley, J., Komura, T.: A recurrent variational autoencoder for human motion synthesis. In: BMVC (2017) [2](#)
24. Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521540518, second edn. (2004) [3](#)
25. He, Y., Pang, A., Chen, X., Liang, H., Wu, M., Ma, Y., Xu, L.: ChallenCap: Monocular 3d capture of challenging human performances using multi-modal references. In: CVPR (2021) [2](#)
26. Henter, G.E., Alexanderson, S., Beskow, J.: MoGlow: Probabilistic and controllable motion synthesis using normalising flows. ACM TOG (2020) [2](#), [4](#)
27. Huang, K., Han, Y., Chen, K., Pan, H., Zhao, G., Yi, W., Li, X., Liu, S., Wei, P., Wang, L.: A hierarchical 3d-motion learning framework for animal spontaneous behavior mapping. Nature Communications (2021) [4](#)
28. Huang, Z., Shi, X., Zhang, C., Wang, Q., Cheung, K.C., Qin, H., Dai, J., Li, H.: FlowFormer: A transformer architecture for optical flow. ECCV (2022) [12](#)
29. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6M: Large scale datasets and predictive methods for 3d human sensing in natural environments. IEEE TPAMI (2014) [4](#)
30. Jakab, T., Li, R., Wu, S., Rupperecht, C., Vedaldi, A.: Farm3D: Learning articulated 3D animals by distilling 2D diffusion. In: 3DV (2024) [3](#)
31. Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: MotionGPT: Human motion as a foreign language. In: NeurIPS (2024) [2](#), [5](#)
32. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: CVPR (2018) [2](#)
33. Kanazawa, A., Tulsiani, S., Efros, A.A., Malik, J.: Learning category-specific mesh reconstruction from image collections. In: ECCV (2018) [3](#), [8](#), [12](#), [13](#)
34. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. In: CVPR (2019) [2](#)
35. Kapon, R., Tevet, G., Cohen-Or, D., Bermano, A.H.: Mas: Multi-view ancestral sampling for 3d motion generation using 2d diffusion. In: CVPR (2024) [4](#)
36. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: ICLR (2014) [7](#)
37. Kirillov, A., Wu, Y., He, K., Girshick, R.: PointRend: Image segmentation as rendering. In: CVPR (2020) [3](#), [9](#)
38. Kokkinos, F., Kokkinos, I.: To the point: Correspondence-driven monocular 3d category reconstruction. In: NeurIPS (2021) [3](#)
39. Kulkarni, N., Gupta, A., Fouhey, D.F., Tulsiani, S.: Articulation-aware canonical surface mapping. In: CVPR (2020) [3](#), [13](#), [14](#)
40. Kulkarni, N., Gupta, A., Tulsiani, S.: Canonical surface mapping via geometric cycle consistency. In: ICCV (2019) [13](#), [14](#)
41. Laine, S., Hellsten, J., Karras, T., Seol, Y., Lehtinen, J., Aila, T.: Modular primitives for high-performance differentiable rendering. ACM TOG (2020) [7](#)
42. Li, X., Liu, S., De Mello, S., Kim, K., Wang, X., Yang, M., Kautz, J.: Online adaptation for consistent mesh reconstruction in the wild. In: NeurIPS (2020) [3](#)
43. Li, X., Liu, S., Kim, K., De Mello, S., Jampani, V., Yang, M.H., Kautz, J.: Self-supervised single-view 3d reconstruction via semantic consistency. In: ECCV (2020) [3](#), [13](#), [14](#)

44. Li, Z., Dekel, T., Cole, F., Tucker, R., Snavely, N., Liu, C., Freeman, W.T.: Learning the depths of moving people by watching frozen people. In: CVPR (2019) [2](#)
45. Lin, X., Amer, M.R.: Human motion modeling using dvgans. arXiv preprint arXiv:1804.10652 (2018) [4](#)
46. Ling, H., Kim, S.W., Torralba, A., Fidler, S., Kreis, K.: Align your gaussians: Text-to-4d with dynamic 3d gaussians and composed diffusion models. In: CVPR (2024) [4](#)
47. Liu, D., Stathopoulos, A., Zhangli, Q., Gao, Y., Metaxas, D.: LEPARD: Learning explicit part discovery for 3d articulated shape reconstruction. In: NeurIPS (2024) [3](#)
48. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM TOG (2015) [2](#), [4](#), [5](#), [6](#)
49. Minderer, M., Sun, C., Villegas, R., Cole, F., Murphy, K.P., Lee, H.: Unsupervised learning of object structure and dynamics from videos. In: NeurIPS (2019) [4](#)
50. Muybridge, E.: The horse in motion (1887) [1](#)
51. Newcombe, R.A., Fox, D., Seitz, S.M.: DynamicFusion: Reconstruction and tracking of non-rigid scenes in real-time. In: CVPR (2015) [3](#)
52. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: HoloGAN: Unsupervised learning of 3d representations from natural images. In: ICCV (2019) [3](#)
53. Niemeyer, M., Geiger, A.: GIRAFFE: Representing scenes as compositional generative neural feature fields. In: CVPR (2021) [3](#)
54. Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: CVPR (2020) [3](#)
55. OpenAI: ChatGPT (2023), <https://chat.openai.com/> [12](#)
56. Ormoneit, D., Black, M., Hastie, T., Kjellström, H.: Representing cyclic human motion using functional analysis. Image and Vision Computing (2005) [4](#)
57. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3D human motion synthesis with transformer VAE. In: ICCV (2021) [2](#), [4](#), [7](#), [11](#)
58. Petrovich, M., Black, M.J., Varol, G.: Temos: Generating diverse human motions from textual descriptions. In: ECCV (2022) [2](#), [4](#)
59. Piao, J., Sun, K., Wang, Q., Lin, K.Y., Li, H.: Inverting generative adversarial renderer for face reconstruction. In: CVPR (2021) [2](#)
60. 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) [4](#)
61. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: ICCV (2019) [2](#)
62. Saito, S., Simon, T., Saragih, J., Joo, H.: PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: CVPR (2020) [2](#)
63. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: GRAF: Generative radiance fields for 3d-aware image synthesis. In: NeurIPS (2020) [3](#)
64. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In: NeurIPS (2021) [5](#)
65. Singer, U., Polyak, A., Hayes, T., Yin, X., An, J., Zhang, S., Hu, Q., Yang, H., Ashual, O., Gafni, O., Parikh, D., Gupta, S., Taigman, Y.: Make-a-video: Text-to-video generation without text-video data. In: ICLR (2023) [2](#)
66. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. In: NeurIPS (2019) [3](#)

67. Starke, S., Mason, I., Komura, T.: DeepPhase: Periodic autoencoders for learning motion phase manifolds. *ACM TOG* (2022) [2](#), [4](#)
68. Stathopoulos, A., Pavlakos, G., Han, L., Metaxas, D.N.: Learning articulated shape with keypoint pseudo-labels from web images. In: *CVPR* (2023) [3](#)
69. Sun, J.J., Karashchuk, P., Dravid, A., Ryou, S., Fereidooni, S., Tuthill, J., Katsaggelos, A., Brunton, B.W., Gkioxari, G., Kennedy, A., et al.: BKinD-3D: Self-supervised 3d keypoint discovery from multi-view videos. In: *CVPR* (2023) [4](#)
70. Sun, J.J., Ryou, S., Goldshmid, R., Weissbourd, B., Dabiri, J., Anderson, D.J., Kennedy, A., Yue, Y., Perona, P.: Self-supervised keypoint discovery in behavioral videos. In: *CVPR* (2022) [4](#)
71. Sun, K., Wu, S., Huang, Z., Zhang, N., Wang, Q., Li, H.: Controllable 3d face synthesis with conditional generative occupancy fields. In: *NeurIPS* (2022) [3](#)
72. Sun, K., Wu, S., Zhang, N., Huang, Z., Wang, Q., Li, H.: Cgof++: Controllable 3d face synthesis with conditional generative occupancy fields. *IEEE TPAMI* (2023) [3](#)
73. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: *ECCV* (2020) [10](#)
74. Urtasun, R., Fleet, D.J., Lawrence, N.D.: Modeling human locomotion with topologically constrained latent variable models. In: Elgammal, A., Rosenhahn, B., Klette, R. (eds.) *Human Motion – Understanding, Modeling, Capture and Animation*. pp. 104–118. Springer Berlin Heidelberg, Berlin, Heidelberg (2007) [4](#)
75. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *NeurIPS* (2017) [11](#)
76. Wang, Y., Long, M., Wang, J., Gao, Z., Yu, P.S.: Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. In: *NeurIPS* (2017) [2](#)
77. Wu, S., Jakob, T., Rupprecht, C., Vedaldi, A.: DOVE: Learning deformable 3d objects by watching videos. *IJCV* (2023) [4](#)
78. Wu, S., Li, R., Jakob, T., Rupprecht, C., Vedaldi, A.: MagicPony: Learning articulated 3d animals in the wild. In: *CVPR* (2023) [3](#), [5](#), [6](#), [8](#), [9](#), [11](#), [13](#), [14](#)
79. Wu, S., Makadia, A., Wu, J., Snavely, N., Tucker, R., Kanazawa, A.: De-rendering the world’s revolutionary artefacts. In: *CVPR* (2021) [3](#), [8](#)
80. Wu, S., Rupprecht, C., Vedaldi, A.: Unsupervised learning of probably symmetric deformable 3D objects from images in the wild. In: *CVPR* (2020) [3](#)
81. Wu, Y., Chen, Z., Liu, S., Ren, Z., Wang, S.: CASA: Category-agnostic skeletal animal reconstruction. In: *NeurIPS* (2022) [3](#), [4](#)
82. Xiao, J., xiang Chai, J., Kanade, T.: A closed-form solution to non-rigid shape and motion recovery. In: *ECCV* (2004) [3](#)
83. Xie, Y., Jampani, V., Zhong, L., Sun, D., Jiang, H.: OmniControl: Control any joint at any time for human motion generation. In: *ICLR* (2024) [2](#), [5](#)
84. 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* (2021) [4](#)
85. Yang, G., Sun, D., Jampani, V., Vlasic, D., Cole, F., Liu, C., Ramanan, D.: ViSER: Video-specific surface embeddings for articulated 3d shape reconstruction. In: *NeurIPS* (2021) [4](#)
86. Yang, G., Vo, M., Natalia, N., Ramanan, D., Andrea, V., Hanbyul, J.: BANMo: Building animatable 3d neural models from many casual videos. In: *CVPR* (2022) [4](#)
87. Yang, G., Wang, C., Reddy, N.D., Ramanan, D.: Reconstructing animatable categories from videos. In: *CVPR* (2023) [4](#)

88. Yang, Y., Yang, J., Xu, Y., Zhang, J., Lan, L., Tao, D.: APT-36K: A large-scale benchmark for animal pose estimation and tracking. In: *NeurIPS Dataset and Benchmark Track (2022)* [10](#), [12](#), [13](#), [14](#)
89. Yao, C.H., Hung, W.C., Li, Y., Rubinstein, M., Yang, M.H., Jampani, V.: Hi-LASSIE: High-fidelity articulated shape and skeleton discovery from sparse image ensemble. In: *CVPR (2023)* [3](#)
90. Yao, C.H., Hung, W.C., Rubinstein, M., Lee, Y., Jampani, V., Yang, M.H.: LASSIE: Learning articulated shape from sparse image ensemble via 3d part discovery. In: *NeurIPS (2022)* [3](#), [8](#)
91. Yao, C.H., Raj, A., Hung, W.C., Rubinstein, M., Li, Y., Yang, M.H., Jampani, V.: ARTIC3D: Learning robust articulated 3d shapes from noisy web image collections. In: *NeurIPS (2024)* [3](#)
92. Zhang, J.Y., Felsen, P., Kanazawa, A., Malik, J.: Predicting 3d human dynamics from video. In: *ICCV (2019)* [2](#)
93. Zhao, Y., Yan, Z., Xie, E., Hong, L., Li, Z., Lee, G.H.: Animate124: Animating one image to 4d dynamic scene. *arXiv preprint arXiv:2311.14603 (2023)* [4](#)
94. Zheng, Y., Li, X., Nagano, K., Liu, S., Hilliges, O., De Mello, S.: A unified approach for text-and image-guided 4d scene generation. In: *CVPR (2024)* [4](#)
95. Zhou, Z., Wang, B.: UDE: A unified driving engine for human motion generation. In: *CVPR (2023)* [2](#), [5](#)