

# Unsupervised Volumetric Animation

Aliaksandr Siarohin  
Snap Inc.

Willi Menapace\*  
University of Trento

Ivan Skorokhodov\*  
KAUST

Kyle Olszewski  
Snap Inc.

Jian Ren  
Snap Inc.

Hsin-Ying Lee  
Snap Inc.

Menglei Chai  
Snap Inc.

Sergey Tulyakov  
Snap Inc.

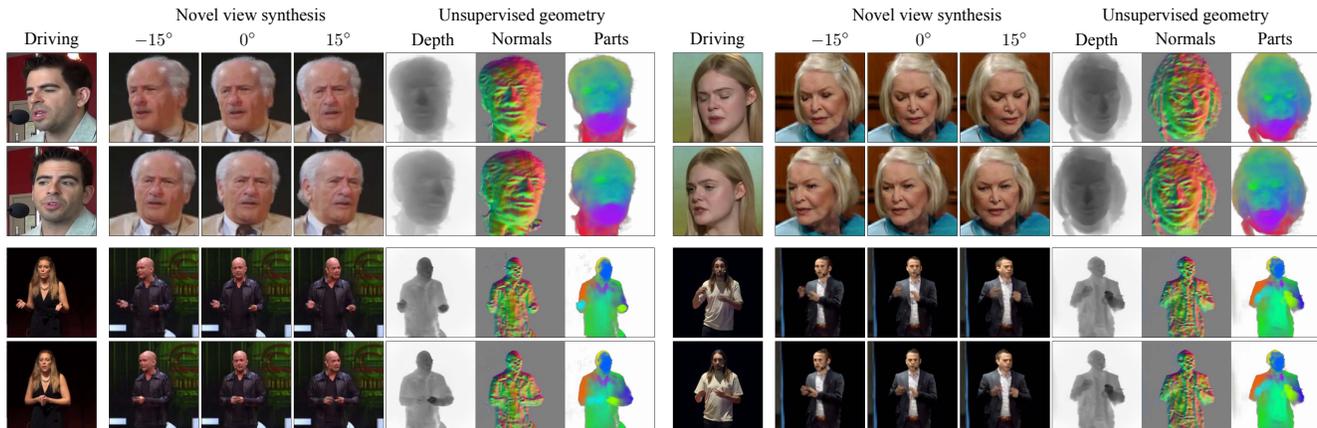


Figure 1. **Unsupervised Volumetric Animation (UVA)**. Selected animation results for faces and bodies. Given a driving image sequence and a source image (not shown), UVA renders realistic animations and *simultaneously* generates novel views of the animated object. With our reconstruction loss, our method also generates high-fidelity depth and normals, and identifies semantically meaningful object parts.

## Abstract

We propose a novel approach for unsupervised 3D animation of non-rigid deformable objects. Our method learns the 3D structure and dynamics of objects solely from single-view RGB videos, and can decompose them into semantically meaningful parts that can be tracked and animated. Using a 3D autoencoder framework, paired with a keypoint estimator via a differentiable PnP algorithm, our model learns the underlying object geometry and parts decomposition in an entirely unsupervised manner. This allows it to perform 3D segmentation, 3D keypoint estimation, novel view synthesis, and animation. We primarily evaluate the framework on two video datasets: VoxCeleb 256<sup>2</sup> and TEDXPeople 256<sup>2</sup>. In addition, on the Cats 256<sup>2</sup> image dataset, we show it even learns compelling 3D geometry from still images. Finally, we show our model can obtain animatable 3D objects from a single or few images<sup>1</sup>.

<sup>1</sup>Code and visual results available on our project website:

<https://snap-research.github.io/unsupervised-volumetric-animation>.

\* Work done while interning at Snap.

## 1. Introduction

The ability to realistically animate a dynamic object seen in a single image enables compelling creative tasks. Such applications range from tractable and cost-effective approaches to visual effects for cinema and television, to more lightweight consumer applications (*e.g.*, enabling arbitrary users to create “performances” by famous modern or historical figures). However, this requires understanding the object’s structure and motion patterns from a single static depiction. Efforts in this field are primarily divided into two approaches: those that outsource this understanding to existing, off-the-shelf models specific to an object category that capture its particular factors of variation; and those that learn the object structure from the raw training data itself. The former group employs *supervision*, and thus requires knowledge about the animated object (*e.g.*, the plausible range of shapes and motions of human faces or bodies). The latter group is *unsupervised*, providing the flexibility needed for a wider range of arbitrary object categories.

Significant progress has been made recently in the domain of unsupervised image animation. Methods in this category typically learn a motion model based on object

parts and the corresponding transformations applied to them. Initially, such transformations were modeled using a simple set of sparse keypoints. Further works improved the motion representation [52, 55], learned latent motion dictionaries [64], kinematic chains [59] or used thin-plate spline transformations [81]. However, broadly speaking, all such works propose *2D motion representations*, warping the pixels or features of the input image such that they correspond to the pose of a given driving image. As such, prior unsupervised animation methods offer means to perform 2D animation only, and are inherently limited in modeling complex, 3D effects, such as occlusions, viewpoint changes, and extreme rotations, which can only be explained and addressed appropriately when considering the 3D nature of the observed objects.

Our work fundamentally differs from prior 2D works in that it is the first to explore unsupervised image animation in 3D. This setting is substantially more challenging compared to classical 2D animation for several reasons. First, as the predicted regions or parts now exist in a 3D space, it is quite challenging to identify and plausibly control them from only 2D videos without extra supervision. Second, this challenge is further compounded by the need to properly model the distribution of the camera in 3D, which is a problem in its own right [40], with multiple 3D generators resorting to existing pose predictors to facilitate the learning of the underlying 3D geometry [5, 58]. Finally, in 3D space, there exists no obvious and tractable counterpart for the bias of 2D CNNs, which are essential for unsupervised keypoint detection frameworks for 2D images [53].

We offer a solution to these challenges. Our framework maps an embedding of each object to a canonical volumetric representation, parameterized with a voxel grid, containing volumetric density and appearance. To allow for non-rigid deformations of the canonical object representation, we assume the object consists of a certain number of rigid parts which are softly assigned to each of the points in the canonical volume. A procedure based on linear blend skinning is employed to produce the deformed volume according to the pose of each part. Rather than directly estimating the poses, we introduce a set of learnable 3D canonical keypoints for each part, and leverage the 2D inductive bias of 2D CNNs to predict a set of corresponding 2D keypoints in the current frame. We propose the use of a differentiable Perspective-n-Point (PnP) algorithm to estimate the corresponding pose, explicitly linking 2D observations to our 3D representation. This framework allows us to propagate the knowledge from 2D images to our 3D representation, thereby learning rich and detailed geometry for diverse object categories using a photometric reconstruction loss as our driving objective. The parts are learned in an unsupervised manner, yet they converge to meaningful volumetric object constituents. For example, for faces, they correspond to the jaw, hair, neck,

and the left and right eyes and cheeks. For bodies, the same approach learns parts to represent the torso, head, and each hand. Examples of these parts are given in Fig. 1.

To simplify the optimization, we introduce a two-stage strategy, in which we start by learning a single part such that the overall geometry is learned, and proceed by allowing the model to discover the remaining parts so that animation is possible. When the object is represented with a single part, the model can perform 3D reconstruction and novel view synthesis. When more parts are used, our method allows us to not only identify meaningful object parts, but to perform non-rigid animation *and* novel view synthesis *at the same time*. Examples of images animated using our Unsupervised Volumetric Animation (UVA) are given in Fig. 1.

We train our framework on three datasets containing images or videos of various objects. We first show that our method learns meaningful 3D geometry when trained on still images of cat faces [79]. We then train our method on the VoxCeleb [38] and TEDXPeople [17] video datasets to evaluate 3D animation. Since our method is the first to consider unsupervised 3D animation, we further introduce evaluation metrics assessing novel view synthesis and animation quality when only single-view data is available.

## 2. Related work

**3D-aware image and video synthesis** experienced substantial progress over the last two years. Early works [40, 41, 51] used Neural Radiance Fields (NeRFs) [37] as a 3D representation to synthesize simple objects and often considered synthetic datasets [51, 70]. They spurred a line of works that scaled the generator and increased its efficiency to attain high-resolution 3D synthesis [5, 18, 45, 58, 72]. These works rely on different types of volumetric representations such as a coordinate-MLP [6], voxel-grids [39], tri-planes [5, 58], generative manifolds [10], multi-plane representations [82], and signed distance functions [42]. Further works combined implicit video synthesis [57, 76] techniques with that of volumetric rendering [18] to generate 3D-aware videos [1]. A common requirement of these methods is access to the ground truth camera distribution (*e.g.*, [6, 18, 45, 51, 80]) or even the known camera poses for each training image [5, 10, 58, 82]. This gives a strong inductive bias towards recovering the proper 3D geometry [58, 82]. Our work shows that it is possible to learn rich geometry and object parts decomposition in a completely unsupervised manner in a non-adversarial framework.

**3D reconstruction from 2D data.** Reconstruction of 3D objects from image or video collections is a long standing problem [23, 32, 67, 68, 73–75]. Initial attempts [23] utilize image collections and try to predict camera, mesh displacement parameters and texture, render the mesh and use reconstruction as the main guiding signal. Later work [68] proposes to further improve this pipeline by incorporating additional knowledge about object symmetry.

However, those works did not model deformations, which was addressed later in works [32, 67, 73–75] that propose to train on video datasets. Most of them [32, 73–75] optimize the object parameters for each frame, and thus can not be trained on a large dataset of videos. On the contrary, Dove [67] infers the articulation parameters from individual frames, which allows training on large video dataset. However, Dove [67] is a mesh based method, thus rendering quality is limited. Moreover, all of these methods utilize additional annotations such as template shapes [32], camera poses [75], 2D keypoints [75], optical flow [73–75] or ground truth object masks [32, 67, 73–75]. Instead, in our method everything, including object masks, was obtained in a purely unsupervised way from video data only.

**Supervised image animation** requires an off-the-shelf keypoint predictor [62, 77, 78] or a 3D morphable model (3DMM) estimator [14, 15] run through the training dataset prior to training. To train such an estimator, one needs to have large amounts of labeled data for the task at a hand. Supervised animation works are typically designed for only one object category, such as bodies [33, 62] or faces [78]. Among them, some support only a single object identity [4], others single- or few-shot cases [46, 62].

Thanks to significant advances in neural rendering and 3D-aware synthesis, several works extended supervised animation to the 3D domain. Initially, a dataset with multiview videos was required to train animatable radiance fields [43]. Later, HumanNeRF [65] and NeuMan [21] showed the feasibility of leveraging only a monocular video of the same subject. However, these models require fitting of a 3D model of human bodies to every frame of a video. With some exceptions [46], such methods do not support multiple identities with the same framework. In contrast, our method features a space in which all objects are represented in their canonical, animation-ready form.

**Unsupervised image animation** is the most related group of works to ours. These works do not require supervision beyond photometric reconstruction loss and, hence, support a variety of object categories with one framework [36, 52, 53, 55]. A key focus area of such works is to design appropriate motion representations for animation [35, 52, 55, 66]. A number of improved representations have been proposed, such as those setting additional constraints on a kinematic tree [59], and thin-plate spline motion modelling [81]. A further work, titled Latent Image Animator [64], learned a latent space for possible motions. Interestingly, a direction in the latent space is found to be responsible for generating novel views of the same subject. As we confirm experimentally, similarly to 2D image generators [24], the direction cannot be reliably used to synthesize the novel views. Several recent works [11, 63], propose to use mixed schemes where pose of the object is supervised and expression is learned, such approaches works well for faces

however did not generalize to other categories.

### 3. Method

This section presents our method for unsupervised 3D animation of non-rigid deformable objects. Our model trains on a set of images  $\{F_i, \alpha_i\}_{i=1}^{N_f}$ , where  $F_i \in \mathbb{R}^{H \times W \times 3}$  is an image frame,  $\alpha_i \in \mathbb{N}$  is an object identifier<sup>2</sup>, and  $N_f$  is the number of frames in a video. The primary training objective of our framework is the reconstruction task. Given a frame  $F_i$  with identity  $\alpha_i$ , we reconstruct it using four core components (see Fig. 2). First, Canonical Voxel Generator  $G$  maps a learnable identity-specific embedding  $\mathbf{e} \in \mathbb{R}^{N_e}$  to an object’s volumetric representation in the canonical pose, parametrized as a voxel grid. Following [52, 53, 55], we assume that each non-rigid object can be represented as a set of moving rigid parts. In this way, our voxel generator segments the volume and assigns each 3D point to its corresponding object’s part (Sec. 3.1). Next, we define 2D keypoint predictor  $C$  with and the differentiable PnP [28] algorithm to estimate each part’s pose (position and orientation) in a given RGB frame  $F_i$  (Sec. 3.2). Subsequently, we employ a method based on linear blend skinning [29] to map the canonical object volume into a deformed one which represents the object in the current frame (Sec. 3.3). Finally, we use volumetric rendering [37] to render the colors to the image space (Sec. 3.4).

#### 3.1. Canonical Voxel Generator

We use a voxel grid  $V$  to parametrize the volume since we found it to provide the best trade-off between generation efficiency, expressivity and rendering speed. Given an object’s embedding  $\mathbf{e} \in \mathbb{R}^{N_e}$ , we use Canonical Voxel Generator  $G$  to produce a volume cube of size  $S$ :

$$G(\mathbf{e}) = V = [V^{\text{Density}} \parallel V^{\text{RGB}} \parallel V^{\text{LBS}}], \quad (1)$$

where  $V^{\text{Density}} \in \mathbb{R}^{S^3}$  is the object’s (discretized) density field in the canonical pose and  $V^{\text{RGB}} \in \mathbb{R}^{S^3 \times 3}$  is its (discretized) RGB radiance field. To animate an object, we assume that it can be modeled as a set of rigid moving parts  $p \in \{1, 2, \dots, N_p\}$  [52, 53, 55], so we use  $V^{\text{LBS}} \in \mathbb{R}^{S^3 \times N_p}$  to model a soft assignment of each point of the volume to one of the  $N_p$  parts. Notably, we do not use any encoder to produce identity embeddings  $\mathbf{e}$  and instead optimize them directly during training [3]. Examples of canonical density, parts, and rendered canonical radiance are shown in Fig. 2.

#### 3.2. Unsupervised Pose Estimation

As described in Sec. 3.1, we assume that an object movement can be factorized into a set of rigid movements of each individual object’s part  $p$ . However, detecting 3D part poses, especially in an *unsupervised* way, is a

<sup>2</sup>We assume that we know which object instance appears in a video. In practice, this assumption is easily satisfied by assigning the same identity to all the frames of a given video.

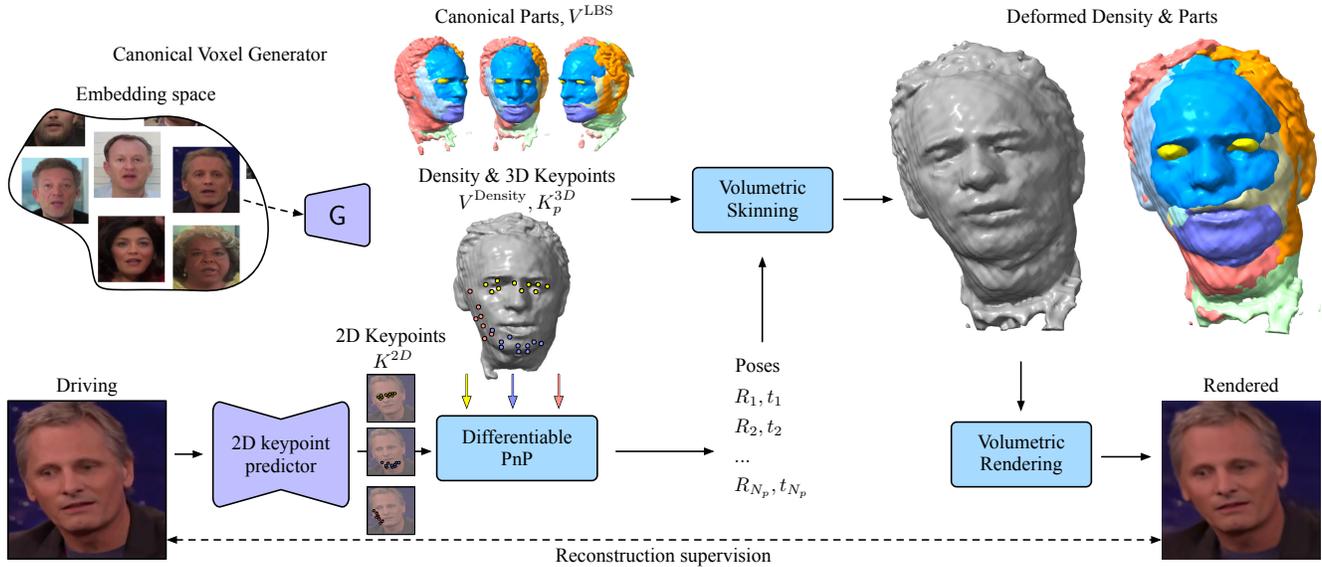


Figure 2. **Unsupervised Volumetric Animation** consists of a canonical voxel generator  $G$  mapping a point in the latent space to the canonical density, radiance and canonical parts. In the embedding space we show canonical shapes rendered under identity camera (faces have the same pose with mouth open). For each part, a set of canonical 3D keypoints  $K^{3D}$  is learnt during training. The 2D keypoint predictor uses a driving image to predict a set of 2D keypoints  $K^{2D}$ , corresponding to  $K^{3D}$ . The differentiable PnP algorithm is used to predict the pose of each part. Canonical density, radiance, poses and parts are then used to compute the deformed density and radiance via volumetric skinning. We then volumetrically render the deformed radiance to produce the rendered image. Note, that our approach does not use any knowledge about the object being animated, and is supervised using the reconstruction loss. Zoom-in for greater detail.

difficult task. MRAA [55] shows that estimating 2D parts and their poses in an unsupervised fashion is an under-constrained problem, which requires specialized inductive biases to guide the pose estimation towards the proper solution. We incorporate such an inductive bias by framing pose prediction as a 2D landmark detection problem which CNNs can solve proficiently due to their natural ability to detect local patterns [20].

To lift this 2D bias into 3D, we estimate the poses of 3D parts by learning a set of 3D keypoints in the canonical space and detecting their 2D projections in the current frame using a 2D CNN. We then use a differentiable Perspective-n-Point (PnP) formulation to recover the pose of each part since we know its corresponding 2D and 3D keypoints. More formally, PnP is a problem where, given a set of the 3D keypoints  $K^{3D} \in \mathbb{R}^{N_k \times 3}$ , a set of corresponding 2D projections  $K^{2D} \in \mathbb{R}^{N_k \times 2}$  and the camera intrinsics parameters, one need to find a camera pose  $T = [R, t] \in \mathbb{R}^{3 \times 4}$ , such that  $K^{3D}$  project to  $K^{2D}$  when viewed from this pose. Note that, while  $T$  represents the pose of the camera with respect to the part, in our framework we consider the camera extrinsics to be constant and equal to the identity matrix, i.e. a part moves while the camera remains fixed. Recovering a part’s pose with respect to the camera is performed by inverting the estimated pose matrix  $T_p = [R_p, t_p] = [R^{-1}, -R^{-1}t]$ .

To implement this idea, we introduce  $N_k$  learnable canonical 3D keypoints  $K_p^{3D}$  for each part, totaling  $N_k \times$

$N_p$ . These 3D keypoints are shared among all the objects in a dataset, which are directly optimized with the rest of the model’s parameters. Then, we define a 2D keypoints prediction network  $C$ , which takes frame  $F_i$  as input and outputs  $N_k$  2D keypoints  $K_p^{2D}$  for each part  $p$ , where each 2D keypoint corresponds to its respective 3D keypoint. The pose of part  $p$  is thus recovered as:

$$T_p^{-1} = \text{PnP}(K_p^{2D}, K_p^{3D}) = \text{PnP}(C(F_i), K_p^{3D}). \quad (2)$$

Crucially, in this formulation  $K_p^{3D}$  are shared for all the objects in the dataset, thus all objects will share the same canonical space for poses. This property is essential for performing cross-subject animations, where poses are estimated on frames depicting a different identity.

We used the EPnP [28] implementation from Pytorch3D [44], since we found it to be faster and more stable than the methods based on declarative layers [8, 16].

### 3.3. Volumetric Skinning

In this section, we describe the procedure to deform the canonical volumetric object representation into its representation in the driving pose. The deformation can be completely described by establishing correspondences between each point  $x_d$  in the deformed space and points  $x_c$  in the canonical space. We establish such correspondence through Linear Blend Skinning (LBS) [29]:

$$x_d = \sum_{p=1}^{N_p} w_p^c(x_c) (R_p x_c + t_p), \quad (3)$$

where  $w_p^c(x)$  is a weight assigned to each part  $p$ . Intuitively, LBS weights segment the object into different parts. As an example, a point with LBS weight equal to 1.0 for the left hand, will always move according to the transformation for the left hand. Unfortunately, during volumetric rendering we typically need to query canonical points using points in the deformed space, requiring solving Eq. (3) for  $x_c$ . This procedure is prohibitively expensive [30], so we rely on the approximate solution introduced in HumanNeRF [65], which defines *inverse* LBS weights  $w_p^d$  such that:

$$x_c = \sum_{p=1}^{N_p} w_p(x_d) (R_p^{-1}x_d - R_p^{-1}t_p), \quad (4)$$

where weights  $w_p^d$  are defined as follows:

$$w_p(x_d) = \frac{w_p^c(R_p^{-1}x_d - R_p^{-1}t_p)}{\sum_{p=1}^{N_p} w_p^c(R_p^{-1}x_d - R_p^{-1}t_p)}. \quad (5)$$

This approximation has an intuitive explanation, i.e. given the deformed point, we project it using the inverse  $T_p$  to the canonical pose and check if it corresponds to the part  $p$  in canonical pose. It is easy to see that if each point has a strict assignment to a *single* part and there is no self-penetration in the deformed space, the approximation is exact. In our work, we parametrize  $w_p^c$  as the channel-wise softmax of  $V^{\text{LBS}}$ . Examples of the parts are given in Figs. 1 & 2.

### 3.4. Volumetric Rendering

We render the deformed object using differentiable volumetric rendering [37]. Given camera intrinsics and extrinsics, we cast a ray  $r$  through each pixel in the image plane and compute the color  $c$  associated to each ray by integration as:

$$c(r) = \int_{t_n}^{t_f} e^{-\int_{t_n}^t \sigma(r(s))ds} \sigma(r(t))c(r(t))dt, \quad (6)$$

where  $\sigma$  and  $c$  are functions mapping each 3D point along each ray  $r(t)$  to the respective volume density and radiance. In our framework, we parametrize  $\sigma$  as  $V^{\text{Density}}$  and  $c$  as  $V^{\text{RGB}}$  which can be efficiently queried using trilinear interpolation. We train the model using a camera with fixed extrinsics initialized to the identity matrix, and fixed intrinsics. Note that, to reduce computational resources, we render images directly from voxels without any additional MLP, nor did we employ any upsampling technique.

We assume that the background is flat and it is not moving. We model it as a plate of fixed, high density. This density is modeled with a single dedicated volume, while the color is obtained from  $V^{\text{RGB}}$ .

### 3.5. Training

Learning a 3D representation of an articulated object from 2D observations without additional supervision is a

highly ambiguous task, prone to spurious solutions with poor underlying geometry that leads to corrupted renderings if the camera is moved away from the origin. We devise a two-stage training strategy that promotes learning of correct 3D representations. First, we train the model with only a single part, i.e.  $N_p = 1$ . This allows the model to obtain meaningful estimation of the object geometry. Thus we name this pretraining a Geometry phase or *G-phase*. During the second phase, we introduce  $N_p = 10$  parts, allowing the model to learn the pose of each part. We copy all the weights from the *G-phase*. Moreover, for C the weight of the final layer is extended such that all the part predictions are the same as in the first stage, while for G, we just add additional weights for  $V^{\text{LBS}}$  initialized to zero.

The model is trained using a range of losses.

**Reconstruction loss.** We use perceptual reconstruction loss [22] as the main driving loss. Similarly to FOMM [52] we use a pyramid of resolutions:

$$\mathcal{L}_r = \sum_l \sum_i \left| \text{VGG}_i(D_l \odot \hat{F}) - \text{VGG}_i(D_l \odot F) \right|, \quad (7)$$

where  $\text{VGG}_i$  is the  $i^{\text{th}}$ -layer of a pretrained VGG-19 [56] network, and  $D_l$  is a downsampling operator corresponding to the current resolution in the pyramid.

**Unsupervised background loss.** Contrary to 2D frameworks for unsupervised animation that use motion cues to separate background from foreground objects, our generator G mostly relies on appearance features, thus it is harder for it to disentangle the background from the foreground. In the first stage, we encourage the model to correctly disentangle the background from the foreground leveraging a coarse background mask  $B$  that we obtain in an unsupervised manner from MRAA [55]. Given the occupancy map  $O$  for the foreground part obtainable by evaluating Eq. (6) excluding the background, we enforce a cross entropy loss:

$$\mathcal{L}_{\text{bkg}} = \sum_i O \log(1 - B) + (1 - O) \log(B), \quad (8)$$

the background mask  $B$  is very coarse and we observe is necessary only in the earliest iterations to avoid degenerate solutions, thus we reduce the contribution of this loss each epoch, we specify the exact schedule in Appx A.

**Pose losses.** Finally, to regularize the PnP-based pose prediction we add two regularization terms: equivariance and projection. First one is a standard technique for regularizing unsupervised keypoints [52, 55]:

$$\mathcal{L}_{\text{eq}} = |A \circ C(F) - C(\mathcal{W}(F, A))|, \quad (9)$$

where  $A$  is a random affine transformation, and  $\mathcal{W}$  is a warping operation. The intuition behind this loss is that when an image is deformed, its keypoints should undergo a similar deformation. Second, we explicitly minimize the

$K^{3D}$  reprojection error with  $K^{2D}$ :

$$\mathcal{L}_{\text{proj}} = \sum_p |K_p^{2D} - \Pi(K_p^{3D}, T_p)|, \quad (10)$$

where  $\Pi_p$  projects the points according to the estimated camera pose  $T_p$ . This loss enforces keypoints to comply with the predicted camera transformation  $T_p$ , improving the stability of the PnP algorithm. The final loss is the sum of all terms with equal weights.

### 3.6. Inference

Despite our model learns the embedding space for the identities in the training set only, it can be used to model previously unseen identities. Given an image of an unseen identity  $F_{\text{test}}$  and a randomly initialized embedding  $e_{\text{test}}$ , we optimize the reconstruction loss  $\mathcal{L}_r$  (Eq. 7) with respect to the embedding  $e_{\text{test}}$ . This procedure produces a volumetric representation with detailed geometry, but imperfect textures. We address this issue by finetuning the generator  $G$ , following the pivotal tuning procedure [49]. In order to avoid significant distortions to the geometry, during this finetuning stage we regularize  $V^{\text{Density}}$  and  $V^{\text{LBS}}$  to stay close to their values prior to finetuning. Note that we only optimize with respect to the appearance and do not modify the 2D keypoint predictor  $C$ , ensuring that motion can be transferred from different objects. Additional details concerning this embedding are provided in Appx A.

## 4. Experiments

Evaluating animation, whether 2D or 3D, is a challenging task as there is no ground truth for the animated images. We are not aware of prior works in unsupervised volumetric animation, hence, in this section we establish an evaluation protocol for this task. Our protocol makes use of established metrics in unsupervised 2D animation, when applicable, and introduces procedures to evaluate the quality of the synthesized 3D geometry and animation under novel views.

**Datasets.** To evaluate our method we use three publicly available datasets: 1) Cats [79], consisting of 9,993 images of cat faces. We used 9,793 for training and 200 for testing. 2) For VoxCeleb [38], we employed the same pre-processing as FOMM [52], using 19522 face videos for training and 100 for testing. 3) TEDXPeople [17] is a video dataset of TEDx speakers. Using timestamps provided by [17], we extract continuous video chunks. More details can be found in Appx B. In total, we employ 40896 videos for training, and retain 100 videos for testing.

### 4.1. Geometry from Image Data

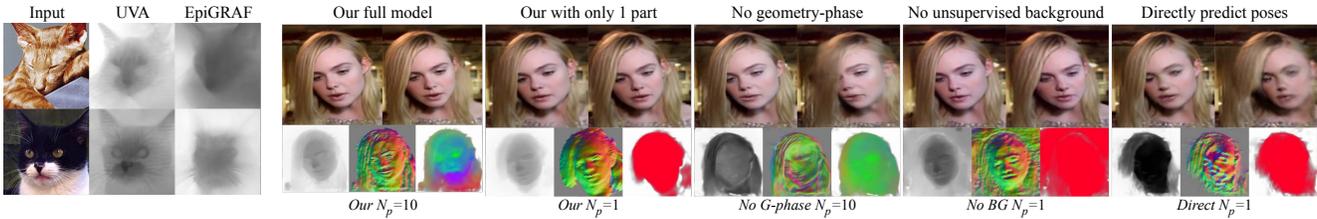
Our method learns high-fidelity geometry from images or videos *without camera or geometry* supervision. This is a challenging setting, even for recent 3D-GANs, as they *require* camera supervision. In this setting, we compare the quality of inferred geometry to a state-of-the-art 3D-GAN, EpiGRAF [58], trained with ground truth camera

poses. As both UVA and EpiGRAF render non-absolute depth, to evaluate its quality, we use the Pearson correlation coefficient. Given a test image, we reconstruct it by inversion, and obtain depth using volumetric rendering. We then measure the correlation between the predicted depth and the depth estimated with an off-the-shell depth estimator [12]. For a fair comparison, during inversion, we do not provide camera poses to EpiGRAF and instead find them during the optimization, in combination with the rest of the parameters. UVA provides higher-quality depth, while not requiring camera supervision during training, reaching a correlation value of 0.63. EpiGRAF reaches only 0.53, often failing to render accurate depth for non-frontal cameras (see Fig. 3a).

### 4.2. Animation Evaluation

Unsupervised animation in 3D is a new task introduced in this work. A key feature of 3D animation is the ability to change the viewpoint from which the object is rendered during animation. Commonly used animation datasets, however, do not typically offer multi-view data. To evaluate viewpoint consistency without access to multi-view data, we introduce three new metrics: Average Yaw Deviation (AYD), Average Shape Consistency (ASC), and Average Pose Consistency (APC). In more detail, given an object, we rotate it along the y-axis using a set of predefined angles. We then fit an SMPL [9] model for humans and a 3DMM [13] for faces to the frontal and rotated views of the objects. These models estimate the root angle, defining how the object is oriented with respect to the camera; a shape parameter, defining the identity of the object; and a parameter defining its pose (in terms of joint rotations for SMPL and facial expression parameters for 3DMM). To evaluate the ability of the model to rotate the object by the required angle to produce novel views, we use AYD. In particular, we compute the y-axis component of the root angle between the rotated and non-rotated object, and compare it with the known yaw of the camera, used to render that view. We use ASC to compare the consistency of the shape parameters between the frontal and the rotated views. A lower ASC indicates that the identity is better preserved during rotation. APC is used to measure how much the pose is altered during rotation, with a lower APC indicating better preservation of the object pose. These metrics enable evaluating the capabilities of competing models in generating view-consistent results. Appx C contains full details on these metrics.

No prior unsupervised animation method [52, 55, 64] offers a built-in ability to generate the data under novel views. Thus, for [52,55] we introduce a simple, depth-based method to generate novel views. First, we predict the depth from a monocular depth predictor [12] and normalize it to make it compatible with our camera intrinsics. Then, for each method, we estimate parts and their affine transforma-



(a) Depth comparisons

(b) Qualitative ablation results of methods in Tab. 2

Figure 3. (a) Typical depth examples of embedded images using our method (UVA) and EpiGRAF [58]. Note, UVA’s depth contains sharper details regardless of the pose. (b) We show a block for each method, with novel views (top), and depth, normals, parts (bottom).



Figure 4. **2D animation.** Example 2D animations on bodies and faces from our method (UVA) and a state-of-the-art work, MRAA [55]. As UVA models objects in canonical 3D space, it better preserves an object’s shapes when animated. Zoom-in for greater detail.

tions. We choose a central 2D keypoint for each part, and augment it with 4 additional keypoints in its neighborhood. Using the depth, we lift the keypoints in 3D and re-project them into the novel view. From these new keypoints, a new affine transformation is estimated and used to drive the view synthesis. We then evaluate against LIA [64], which expresses animation as linear navigation in a latent space. For the VoxCeleb [38] dataset, we found one of the components of its latent space to correlate with the rotation of the head along the y-axis. Exploiting this finding, we fit a linear model mapping the magnitude of the movement along this latent component to the produced head rotation, and use it to generate the head under novel viewpoints.

We also use the standard 2D reconstruction metrics: L1, AKD/MKR [53], AED [53]. However, we emphasize that such metrics favor 2D methods, which can solve the 2D animation problem by copying pixels from the source to the target view, at the cost of limited 3D understanding and consistency. In contrast, UVA renders view-consistent pixel values from a 3D representation, making this shortcut unavailable. A significant gap may also be introduced by

the single-image embedding procedure we adopt. Note, however, that as our embedding procedure seamlessly supports the use of multiple source frames at inference time, a shared representation can be optimized, pooling information from all available frames to improve performance. We demonstrate the results of our model with one and five frames. We also note that, despite a wide range of a viewpoints in different videos, subjects in each individual video in VoxCeleb and TEDXPeople have very limited 3D rotations, as they primarily face towards the camera. Thus, standard reconstruction metrics do not reflect the model’s capacity to perform complex 3D transformations.

We provide the quantitative results in Tab. 1. As the affine transformations of FOMM [52] are mostly based on edge detection that is not very robust, any minor modification of this transformations for novel view synthesis leads to significant movement. Thus, FOMM has the worst AYD among all methods. Affine estimation in MRAA, in contrast, is significantly more robust, and thus it has a significantly lower AYD. However, we observe that MRAA does not have enough knowledge about the 3D

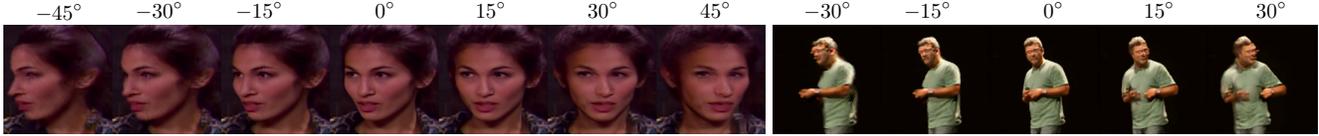


Figure 5. **Novel view synthesis.** We report typical examples of novel views synthesized using a single input image. For bodies, we show a narrower range, as the TEDXPeople dataset is biased towards frontal poses.

Method	VoxCeleb						TEDXPeople					
	AYD↓	ASC↓	APC↓	L1↓	AKD↓	AED↓	AYD↓	ASC↓	APC↓	L1↓	(AKD↓, MKR↓)	AED↓
FOMM [52]	0.655	0.129	0.177	<b>0.0413</b>	1.289	0.134	0.507	0.028	1.07	0.0318	(3.248, 0.009)	0.120
MRAA [55]	0.173	0.123	0.174	0.0424	<b>1.250</b>	0.131	0.181	0.023	0.702	<b>0.0262</b>	<b>(2.282, 0.007)</b>	0.101
LIA [64]	0.207	0.130	0.190	0.0529	1.437	0.138	-	-	-	-	-	-
Our 1 frame	0.051	<b>0.078</b>	0.144	0.0655	1.737	0.226	0.128	<b>0.019</b>	0.635	0.0474	(3.571, 0.017)	0.163
Our 5 frame	<b>0.045</b>	0.091	<b>0.112</b>	0.0418	1.378	<b>0.111</b>	<b>0.107</b>	0.021	<b>0.571</b>	0.029	(2.373, 0.014)	<b>0.086</b>

Table 1. Comparison with 2D animation methods. Novel view synthesis for AYD, ASC & APC from yaw in range  $-45^\circ$  to  $+45^\circ$ .

Method	AYD↓	ASC↓	APC↓	L1↓	AKD↓	AED↓
<i>Direct</i> $N_p = 1$	0.707	0.160	0.239	0.0723	3.582	0.326
<i>No BG</i> $N_p = 1$	0.301	0.117	0.216	0.0702	2.410	0.263
<i>Our</i> $N_p = 1$	0.141	0.113	0.210	0.0637	2.170	0.242
<i>No G-phase</i> $N_p = 10$	1.08	0.145	0.226	<b>0.0620</b>	1.993	0.243
<i>Our</i> $N_p = 10$	<b>0.051</b>	<b>0.078</b>	<b>0.144</b>	0.0655	<b>1.737</b>	<b>0.226</b>

Table 2. Ablation results on the VoxCeleb dataset.

structure of the objects, treating them as planes—while they roughly preserve the shape and pose for small angles, for larger angles objects become thinner, until they eventually disappear. LIA has a rotation direction that is entangled with the other movements, and thus it has the lowest ASC and APC. Finally, our model is the best at preserving shape and expressions, as judged by the ASC and APC. Moreover, our model also provides the most meaningful rotations as judged by the AYD. With respect to standard reconstruction metrics, we hypothesize that they provide unfair advantage to the models that can copy the pixels directly. To quantify the error introduced by the embedding, we embed, reconstruct the same frame and compute L1 error: Our  $-0.0298$ , MRAA  $-0.0015$  (compare with L1 for the full sequence in Tab. 1). Thus the difference in reconstruction performance is mainly due to the embedding procedure. Note however, that the results of our 5 frame model performs on par with the baselines.

However, as previously mentioned, these metrics do not reflect the ability of the model to preform complex 3D movements. This point is further highlighted in Fig. 4, when the pose of the source and driving images differ significantly, MRAA fails to properly reflect that, while our model produces more consistent results. Interesting, we also note that, as our model is based on learning a 3D prior and not copying pixels, it can filter out some occlusions, as seen in the third column in Fig. 4, while MRAA produces artifacts in the occluded region.

### 4.3. Ablation Studies

We evaluate the key design choices made in our framework. First, we compare our PnP-based part pose predictor

with direct part pose prediction (*Direct*). As directly predicting an  $\mathbb{R}^{3 \times 3}$  rotation matrix could produce solutions not corresponding to a rigid rotation, we adopt the 6D rotation parameterization from [83]. The geometry learned by this approach is essentially flat. We compare our method and *Direct* only in the geometry phase of training (e.g., when  $N_p = 1$ ), as it does not produce sufficiently accurate geometry to proceed with the next phase. We also demonstrate the effect of the unsupervised background loss  $\mathcal{L}_{\text{bkg}}$  by training the model without this loss (*No BG*). Finally, we investigate the importance of two-phase training, learning a model with multiple parts without the geometry phase *No G-phase*. We show numerical results in Tab. 2 and qualitative examples in Fig. 3b. Our full model achieves the best scores, and generates higher fidelity novel views and geometric details. The utility of the geometry phase is clearly demonstrated by the scores and qualitative results, which, without this phase, produce corrupted results and do not learn representative parts. While it produces meaningful depth, the model trained without  $\mathcal{L}_{\text{bkg}}$  fails to separate the background and foreground.

## 5. Conclusion

Our approach for unsupervised volumetric animation demonstrates a significant step towards 3D animation of dynamic objects. While trained exclusively on real-world monocular 2D videos, our method obtains high quality geometry, object parts and 3D segmentation and normals. Due to the unsupervised nature of our work, the same approach applies to a variety of object categories without using explicit labels or other cumbersome supervision. This understanding of the underlying geometry and structure of the object, allows our method to perform animation and novel view synthesis at the same time. These properties open exciting possibilities for employing this information for future exploration, e.g. controlling an object’s fine-grained shape, or relighting it for composition into novel environments.

## References

- [1] Sherwin Bahmani, Jeong Joon Park, Despoina Paschalidou, Hao Tang, Gordon Wetzstein, Leonidas Guibas, Luc Van Gool, and Radu Timofte. 3d-aware video generation. *arXiv:2206.14797*, 2022. **2**
- [2] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. **18**
- [3] Piotr Bojanowski, Armand Joulin, David Lopez-Paz, and Arthur Szlam. Optimizing the latent space of generative networks. *arXiv preprint arXiv:1707.05776*, 2017. **3**
- [4] Caroline Chan, Shiry Ginosar, Tinghui Zhou, and Alexei A Efros. Everybody dance now. In *Proceedings of the IEEE International Conference on Computer Vision*, 2019. **3**
- [5] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. **2, 17, 18**
- [6] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. **2**
- [7] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. *arXiv preprint arXiv:2203.09517*, 2022. **17**
- [8] Bo Chen, Alvaro Parra, Jiewei Cao, Nan Li, and Tat-Jun Chin. End-to-end learnable geometric vision by backpropagating pnp optimization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. **4, 17**
- [9] Hongsuk Choi, Gyeongsik Moon, JoonKyu Park, and Kyoung Mu Lee. Learning to estimate robust 3d human mesh from in-the-wild crowded scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. **6, 15**
- [10] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. Gram: Generative radiance manifolds for 3d-aware image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. **2, 14**
- [11] Nikita Drobyshev, Jenya Chelishchev, Taras Khakhulin, Aleksei Ivakhnenko, Victor Lempitsky, and Egor Zakharov. Megaportraits: One-shot megapixel neural head avatars. In *Proceedings of the ACM International Conference on Multimedia*, 2022. **3**
- [12] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *Proceedings of the IEEE International Conference on Computer Vision*, 2021. **6, 15, 17**
- [13] Yao Feng, Haiwen Feng, Michael J. Black, and Timo Bolkart. Learning an animatable detailed 3D face model from in-the-wild images. *ACM Transactions on Graphics*, 2021. **6, 14**
- [14] Zhenglin Geng, Chen Cao, and Sergey Tulyakov. 3d guided fine-grained face manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. **3**
- [15] Zhenglin Geng, Chen Cao, and Sergey Tulyakov. Towards photo-realistic facial expression manipulation. *International Journal of Computer Vision*, 2020. **3**
- [16] Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. **4, 17**
- [17] Artur Grigorev, Karim Isakov, Anastasia Ianina, Renat Bashirov, Ilya Zakharkin, Alexander Vakhitov, and Victor Lempitsky. Stylepeople: A generative model of fullbody human avatars. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. **2, 6, 13, 14, 15, 16, 17, 18**
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *Proceedings of the International Conference on Machine Learning*, 2022. **2**
- [19] Thorsten Hempel, Ahmed A. Abdelrahman, and Ayoub Al-Hamadi. 6d rotation representation for unconstrained head pose estimation. In *Proceedings of the IEEE International Conference on Image Processing*, 2022. **14, 16**
- [20] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *Proceedings of the Neural Information Processing Systems Conference*, 2018. **4**
- [21] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *Proceedings of the European Conference on Computer Vision*, 2022. **3**
- [22] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *Proceedings of the European Conference on Computer Vision*, 2016. **5**
- [23] Angjoo Kanazawa, Shubham Tulsiani, Alexei A. Efros, and Jitendra Malik. Learning category-specific mesh reconstruction from image collections. In *Proceedings of the European Conference on Computer Vision*, 2018. **2**
- [24] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. **3, 14**
- [25] Faisal Khan, Shahid Hussain, Shubhajit Basak, Joseph Lemley, and Peter Corcoran. An efficient encoder-decoder model for portrait depth estimation from single images trained on pixel-accurate synthetic data. *Neural Networks*, 2021. **16, 18**
- [26] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. **18**

- [27] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 13
- [28] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua. Epnnp: An accurate o(n) solution to the pnp problem. *International Journal of Computer Vision*, 2009. 3, 4, 13, 17
- [29] J. P. Lewis, Matt Corder, and Nickson Fong. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In *Special Interest Group on Computer Graphics and Interactive Techniques*, 2000. 3, 4
- [30] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *Proceedings of the European Conference on Computer Vision*, 2022. 5
- [31] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *Special Interest Group on Computer Graphics and Interactive Techniques*, 2017. 15
- [32] 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. In *Proceedings of the Neural Information Processing Systems Conference*, 2020. 2, 3
- [33] Wen Liu, Zhixin Piao, Jie Min, Wenhan Luo, Lin Ma, and Shenghua Gao. Liquid warping gan: A unified framework for human motion imitation, appearance transfer and novel view synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 3
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *Special Interest Group on Computer Graphics and Interactive Techniques*, 2015. 15
- [35] Dominik Lorenz, Leonard Bereska, Timo Milbich, and Bjorn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 3
- [36] Arun Mallya, Ting-Chun Wang, and Ming-Yu Liu. Implicit warping for animation with image sets. In *Proceedings of the Neural Information Processing Systems Conference*, 2022. 3
- [37] 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. In *Proceedings of the European Conference on Computer Vision*, 2020. 2, 3, 5, 13, 17
- [38] Arsha Nagrani, Joon Son Chung, Weidi Xie, and Andrew Senior. Voxceleb: Large-scale speaker verification in the wild. *Computer Science and Language*, 2019. 2, 6, 7, 14, 16, 17, 18
- [39] Thu H Nguyen-Phuoc, Christian Richardt, Long Mai, Yongliang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. In *Proceedings of the Neural Information Processing Systems Conference*, 2020. 2
- [40] Michael Niemeyer and Andreas Geiger. Campari: Camera-aware decomposed generative neural radiance fields. In *Proceedings of the International Conference on 3D Vision*, 2021. 2
- [41] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 2, 13
- [42] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. StyleSDF: High-Resolution 3D-Consistent Image and Geometry Generation. *arXiv:2112.11427*, 2021. 2
- [43] 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 *Proceedings of the IEEE International Conference on Computer Vision*, 2021. 3
- [44] Nikhila Ravi, Jeremy Reizenstein, David Novotny, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020. 4, 13, 17
- [45] Daniel Rebain, Mark Matthews, Kwang Moo Yi, Dmitry Lagun, and Andrea Tagliasacchi. Lolerf: Learn from one look. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [46] Jian Ren, Menglei Chai, Oliver J Woodford, Kyle Olszewski, and Sergey Tulyakov. Flow guided transformable bottleneck networks for motion retargeting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 3
- [47] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Proceedings of the Neural Information Processing Systems Conference*, 2015. 15
- [48] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, Ethan Rublee, and Gary Bradski. Kornia: an open source differentiable computer vision library for pytorch. In *Proceedings of the Winter Conference on Applications of Computer Vision*, 2020. 17
- [49] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Transactions on Graphics*, 2022. 6, 13
- [50] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention*, 2015. 13
- [51] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. In *Proceedings of the Neural Information Processing Systems Conference*, 2020. 2
- [52] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Proceedings of the Neural Information Processing Systems Conference*, 2019. 2, 3, 5, 6, 7, 8, 13, 14, 16
- [53] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2, 3, 7
- [54] Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening and coloring transform for GANs. In *Proceedings*

- of the *International Conference on Machine Learning*, 2019. 13
- [55] Aliaksandr Siarohin, Oliver Woodford, Jian Ren, Menglei Chai, and Sergey Tulyakov. Motion representations for articulated animation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 2, 3, 4, 5, 6, 7, 8, 13, 16
- [56] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 5
- [57] Ivan Skorokhodov, Sergey Tulyakov, and Mohamed Elhoseiny. Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [58] Ivan Skorokhodov, Sergey Tulyakov, Yiqun Wang, and Peter Wonka. Epigraf: Rethinking training of 3d gans. In *Proceedings of the Neural Information Processing Systems Conference*, 2022. 2, 6, 7, 14, 18
- [59] Jiale Tao, Biao Wang, Borun Xu, Tiezheng Ge, Yuning Jiang, Wen Li, and Lixin Duan. Structure-aware motion transfer with deformable anchor model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 3
- [60] Gül Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 16, 18
- [61] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T. Barron, and Pratul P. Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 18
- [62] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. Few-shot video-to-video synthesis. In *Proceedings of the Neural Information Processing Systems Conference*, 2019. 3
- [63] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. One-shot free-view neural talking-head synthesis for video conferencing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 3
- [64] Yaohui Wang, Di Yang, Francois Bremond, and Antitza Dantcheva. Latent image animator: Learning to animate images via latent space navigation. In *Proceedings of the International Conference on Machine Learning*, 2022. 2, 3, 6, 7, 8, 16
- [65] Chung-Yi Weng, Brian Curless, Pratul P Srinivasan, Jonathan T Barron, and Ira Kemelmacher-Shlizerman. Humannerf: Free-viewpoint rendering of moving people from monocular video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 3, 5
- [66] Olivia Wiles, A Koepke, and Andrew Zisserman. X2face: A network for controlling face generation using images, audio, and pose codes. In *Proceedings of the European Conference on Computer Vision*, 2018. 3
- [67] Shangzhe Wu, Tomas Jakab, Christian Rupprecht, and Andrea Vedaldi. DOVE: Learning deformable 3d objects by watching videos. *arXiv preprint arXiv:2107.10844*, 2021. 2, 3
- [68] Shangzhe Wu, Christian Rupprecht, and Andrea Vedaldi. Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020. 2
- [69] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. 14
- [70] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015. 2
- [71] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *arXiv preprint arXiv:2204.12484*, 2022. 15
- [72] Yang Xue, Yuheng Li, Krishna Kumar Singh, and Yong Jae Lee. Giraffe hd: A high-resolution 3d-aware generative model. *arXiv:2203.14954*, 2022. 2
- [73] 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 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021. 2, 3
- [74] 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. In *Proceedings of the Neural Information Processing Systems Conference*, 2021. 2, 3
- [75] Gengshan Yang, Minh Vo, Natalia Neverova, Deva Ramanan, Andrea Vedaldi, and Hanbyul Joo. Banmo: Building animatable 3d neural models from many casual videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 2, 3
- [76] Sihyun Yu, Jihoon Tack, Sangwoo Mo, Hyunsu Kim, Junho Kim, Jung-Woo Ha, and Jinwoo Shin. Generating videos with dynamics-aware implicit generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, 2022. 2
- [77] Egor Zakharov, Aleksei Ivakhnenko, Aliaksandra Shysheya, and Victor Lempitsky. Fast bi-layer neural synthesis of one-shot realistic head avatars. In *Proceedings of the European Conference on Computer Vision*, 2020. 3
- [78] Egor Zakharov, Aliaksandra Shysheya, Egor Burkov, and Victor Lempitsky. Few-shot adversarial learning of realistic neural talking head models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 3
- [79] Weiwei Zhang, Jian Sun, and Xiaoou Tang. Cat head detection-how to effectively exploit shape and texture features. In *Proceedings of the European Conference on Computer Vision*, 2008. 2, 6, 13, 14
- [80] Xuanmeng Zhang, Zhedong Zheng, Daiheng Gao, Bang Zhang, Pan Pan, and Yi Yang. Multi-view consistent generative adversarial networks for 3d-aware image synthesis.

In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [2](#), [18](#)

- [81] Jian Zhao and Hui Zhang. Thin-plate spline motion model for image animation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [2](#), [3](#)
- [82] Xiaoming Zhao, Fangchang Ma, David Güera, Zhile Ren, Alexander G. Schwing, and Alex Colburn. Generative multiplane images: Making a 2d gan 3d-aware. In *Proceedings of the European Conference on Computer Vision*, 2022. [2](#)
- [83] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. [8](#), [17](#)