Unsupervised Learning of Efficient Geometry-Aware Neural Articulated Representations

Atsuhiro Noguchi¹ Xiao Sun² Stephen Lin² Tatsuya Harada^{1,3}

 1 The University of Tokyo $\qquad ^2$ Microsoft Research Asia $\qquad ^3$ RIKEN

Abstract. We propose an unsupervised method for 3D geometry-aware representation learning of articulated objects, in which no image-pose pairs or foreground masks are used for training. Though photorealistic images of articulated objects can be rendered with explicit pose control through existing 3D neural representations, these methods require ground truth 3D pose and foreground masks for training, which are expensive to obtain. We obviate this need by learning the representations with GAN training. The generator is trained to produce realistic images of articulated objects from random poses and latent vectors by adversarial training. To avoid a high computational cost for GAN training, we propose an efficient neural representation for articulated objects based on tri-planes and then present a GAN-based framework for its unsupervised training. Experiments demonstrate the efficiency of our method and show that GAN-based training enables the learning of controllable 3D representations without paired supervision.

Keywords: image synthesis, articulated objects, neural radiance fields, unsupervised learning

1 Introduction

3D models that allow free control over the pose and appearance of articulated objects are essential in various applications, including computer games, media content creation, and augmented/virtual reality. In early work, articulated objects were typically represented by explicit models such as skinned meshes. More recently, the success of learned implicit representations such as neural radiance fields (NeRF) [40] for rendering static 3D scenes has led to extensions for modeling dynamic scenes and articulated objects. Much of this attention has focused on the photorealistic rendering of humans, from novel viewpoints and with controllable poses, by learning from images and videos.

Existing methods for learning explicitly pose-controllable articulated representations, however, require much supervision, such as videos with 3D pose/mesh annotation and a mask for each frame. Preparing such data involves tremendous annotation costs; thus, reducing annotation is very important. In this paper, we propose a novel unsupervised learning framework for 3D pose-aware generative



Fig. 1. Our ENARF-GAN is a geometry-aware, 3D-consistent image generation model that allows independent control of viewpoint, object pose, and appearance information. It is learned from unlabeled images and a prior distribution on object pose.

models of articulated objects, which are learned only from unlabeled images of objects sharing the same structure and a pose prior distribution of the objects.

We exploit recent advances in 3D-aware GAN [56, 9, 44, 19, 8] for unsupervised learning of the articulated representations. They learn 3D-aware image generation models from images without supervision, such as viewpoints or 3D shapes. The generator is based on NeRF [40] and is optimized with a GAN objective to generate realistic images from randomly sampled viewpoints and latent vectors from a prior distribution defined before training. As a result, the generator learns to generate 3D-consistent images without any supervision. We employ the idea for articulated objects by defining a pose prior distribution for the target object and optimizing the GAN objective on randomly generated images from random poses and latent variables. It becomes possible to learn a generative model with free control of poses. We demonstrate this approach by modeling the pose prior as a skeletal distribution [47, 59], while noting that other models like meshes [53, 52] may bring potential performance benefits.

However, the direct application of existing neural articulated representations to GANs is not computationally practical. While NeRF can produce high-quality images, its processing is expensive because it requires network inference for every point in space. Some methods [44, 19] reduce computational cost by volume rendering at low resolution followed by 2D CNN based upsampling. Although this technique achieves high-resolution images with real-time inference speed, it is not geometry-aware (i.e., the surface mesh cannot be extracted). Recently, a method which we call Efficient NeRF [8] overcomes the problem. The method is based on an efficient tri-plane based neural representation and GAN training on it. Thanks to the computational efficiency, it can produce relatively highresolution (128×128) images with volumetric rendering. We extend the tri-plane representation to articulated objects for efficient GAN training. An overview of the method is visualized in Figure 1. The contributions of this work are as follows:

 We propose a novel efficient neural representation for articulated objects based on an efficient tri-plane representation.

- We propose an efficient implementation of deformation fields using tri-planes for dynamic scene training, achieving 4 times faster rendering than NARF [47] with comparable or better performance.
- We propose a novel GAN framework to learn articulated representations without using any 3D pose or mask annotation for each image. The controllable 3D representation can be learned from real unlabeled images.

2 Related Work

Articulated 3D Representations. The traditional approach for modeling pose-controllable 3D representations of articulated objects is by skinned mesh [25, 26,31], where each vertex of the mesh is deformed according to the skeletal pose. Several parametric skinned mesh models have been developed specifically for humans and animals [37, 23, 51, 48, 73]. For humans, the skinned multi-person linear model (SMPL) [37] is commonly used. However, these representations can only handle tight surfaces with no clothing and cannot handle non-rigid or topology-changing objects such as clothing or hair. Some work alleviates the problem by deforming the mesh surface or using a detailed 3D body scan [1, 20]. Recently, implicit 3D shape representations have achieved state-of-the-art performance in pose-conditioned shape reconstruction. These methods learn neural occupancy/indicator functions [11, 39] or signed distance functions [49] of articulated objects [13, 6, 61, 10]. Photorealistic rendering of articulated objects, especially for humans, is also achieved with 3D implicit representations [52, 47, 59, 2, 65, 35]. However, all these models require ground truth 3D pose and/or object shape for training. Very recently, methods have been proposed to reproduce the 3D shape and motion of objects from video data without using 3D shape and pose annotation [67, 68, 46]. However, they either do not allow free control of poses or are limited to optimizing for a single object. An SMPL meshbased generative model [18] and image-to-image translation methods [14, 57, 7, 36] can learn pose controllable image synthesis models for humans. However, the rendering process is completely in 2D and thus is not geometry-aware.

Implicit 3D representations. Implicit 3D representations are memory efficient, continuous, and topology free. They have achieved the state-of-the art in learning 3D shape [11, 39, 49], static [58, 40, 4] and dynamic scenes [54, 33, 50], articulated objects [53, 13, 6, 61, 10, 52, 47, 59, 2, 65, 35], and image synthesis [56, 9]. Although early works rely on ground truth 3D geometry for training [11, 39, 49], developments in differentiable rendering have enabled learning of networks from only photometric reconstruction losses [58, 69, 40]. In particular, neural radiance fields (NeRF) [40] applied volumetric rendering on implicit color and density fields, achieving photorealistic novel view synthesis of complex static scenes using multi-view posed images. Dynamic NeRF [54, 33, 50] extends NeRF to dynamic scenes, but these methods just reenact the motion in the scene and cannot repose objects based on their structure. Recently, articulated representations based on NeRF have been proposed [53, 47, 59, 52, 65, 66, 35]. These methods can render images conditioned on pose configurations. However, all of

them require ground truth skeletal poses, SMPL meshes, or foreground masks for training, which makes them unsuitable for in-the-wild images.

Another NeRF improvement is the reduction of computational complexity: NeRF requires forward computation of MLPs to compute color and density for every point in 3D. Thus, the cost of rendering is very high. Fast NeRF algorithms [16, 70, 55, 41] reduce the computational complexity of neural networks by creating caches or using explicit representations. However, these methods can only be trained on a single static scene. Very recently, a hybrid explicit and implicit representation was proposed [8]. In this representation, the feature field is constructed using a memory-efficient explicit representation called a tri-plane, and color and density are decoded using a lightweight MLP. This method can render images at low cost and is well suited for image generation models.

In this work, we propose an unsupervised learning framework for articulated objects. We extend tri-planes to articulated objects for efficient training.

Generative 3D-aware image synthesis Advances in generative adversarial networks (GANs) [17] have made it possible to generate high-resolution, photorealistic images [29, 30, 28]. In recent years, many 3D-aware image generation models have been proposed by combining GANs with 3D generators that use meshes [60], voxels [64, 72, 15, 22, 42, 43], depth [45], or implicit representations [56, 9, 44, 19, 8]. These methods can learn 3D-aware generators without 3D supervision. Among these, image generation methods using implicit functions, thanks to their continuous and topology-free properties, have been successful in producing 3D-consistent and high-quality images. However, fully implicit models [56,9] are computationally expensive, making the training of GANs inefficient. Therefore, several innovations have been proposed to reduce the rendering cost of generators. Neural rendering-based methods [44, 19] reduce computation by performing volumetric rendering at low resolution and upsampling the rendered feature images using a 2D CNN. Though this enables the generation of high-resolution images at a faster rate, 2D upsampling does not consider 3D consistency and cannot generate detailed 3D geometry. Very recently, a hybrid of explicit and implicit methods [8] has been developed for 3d geometry-aware image generation. Instead of using a coordinate-based implicit representation, this method uses tri-planes, which are explicit 3D feature representations, to reduce the number of forward computations of the network and to achieve volumetric rendering at high resolution.

The existing research is specific to scenes of static objects or objects that exist independently of each other, and do not allow free control of the skeletal pose of the generated object. Therefore, we propose a novel GAN framework for articulated objects.

3 Method

Recent advances in implicit neural rendering [59, 47] have made it possible to generate 3D-aware pose-controllable images of articulated objects from images with accurate 3D pose and foreground mask annotations. However, training such

models from only in-the-wild images remains challenging since accurate 3D pose annotations are generally difficult to obtain for them. In the following, we first briefly review Neural Articulated Radiance Field (NARF) [47], then propose an adversarial-based framework, named ENARF-GAN, to efficiently train the NARF model without any paired image-pose and foreground mask annotations.

3.1 Neural Articulated Radiance Field Revisited

NARF is an implicit 3D representation for articulated objects. It takes a kinematic 3D pose configuration of an articulated object $o = \{l_k, R_k, \mathbf{t}_k\}_{k=1:K}$ as input and predicts the color and the density of any 3D location \mathbf{x} , where l_k is the length of the k^{th} part, and R_k and \mathbf{t}_k are its rotation and translation matrices, respectively. Given the pose configuration o, NARF first transforms a global 3D position \mathbf{x} into several local coordinate systems defined by the rigid parts of the articulated object. Specifically, the transformed local location \mathbf{x}_k^l for the k^{th} part is computed as $\mathbf{x}_k^l = (R^k)^{-1}(\mathbf{x} - \mathbf{t}^k)$ for $k \in \{1, ..., K\}$.

NARF first trains an extra lightweight selector S in the local space to decide which part a global 3D location \mathbf{x} belongs to. Specifically, it outputs the probability p^k of \mathbf{x} belonging to the k^{th} part. Then NARF computes color c and density σ at the location \mathbf{x} from a concatenation of local locations masked by the corresponding part probability p^k .

$$c, \sigma = G(\operatorname{Cat}(\{\gamma(\mathbf{x}_k^l) * p^k\}_{k=1:K})), \tag{1}$$

where γ is a positional encoding [40], Cat is the concatenation operation, and G is an MLP network. The RGB color **C** and foreground mask value **M** for each pixel are generated by volumetric rendering [40]. The network is trained with a reconstruction loss between the generated and ground truth color $\hat{\mathbf{C}}$ and mask $\hat{\mathbf{M}}$,

$$\mathcal{L}_{\text{supervised}} = \sum_{r \in \mathcal{R}} \left(||\mathbf{C} - \hat{\mathbf{C}}||_2^2 + ||\mathbf{M} - \hat{\mathbf{M}}||_2^2 \right),$$
(2)

where \mathcal{R} is the set of rays in each batch. Please refer to the original NARF paper [47] for more details.

3.2 Unsupervised Learning by Adversarial Training

In this work, we propose a method for efficient and unsupervised training of the NARF model from unposed image collections. Without loss of generality, we consider humans as the articulated objects here. To this end, we first define a human pose distribution \mathcal{O} . For one training iteration, our NARF based generator G takes a latent vector \mathbf{z} and a sampled human pose instance o from \mathcal{O} as input and predicts a synthesized image \mathbf{C} . Following standard adversarial training of GANs, a discriminator \mathcal{D} is used to distinguish the synthesized image \mathbf{C} from real ones $\tilde{\mathbf{C}}$. Formally, the training objectives of the generator \mathcal{L}_{ady}^{G} and



Fig. 2. Overview of (a) Efficient NARF (ENARF) and (b) GAN training.

discriminator \mathcal{L}_{adv}^D are defined as follows,

$$\mathcal{L}_{\mathrm{adv}}^{G} = -\mathbb{E}\left[\log(D(G(\mathbf{z}, o)))\right], \mathcal{L}_{\mathrm{adv}}^{D} = -\mathbb{E}\left[\log(D(\tilde{\mathbf{C}})) + \log(1 - D(G(\mathbf{z}, o)))\right].$$
(3)

An overview of this method is illustrated in Figure 2 (b).

However, this training would be computationally expensive. The rendering cost of NARF is heavy because computation is performed for many 3D locations in the viewed space. Even though in supervised training, the time and memory cost of computing the reconstruction loss could be reduced by evaluating it over just a small proportion of the pixels [47], the adversarial loss in Equation 3 requires the generation of the full image for evaluation. As a result, the amount of computation becomes impractical.

In the following, we propose a series of changes in feature computation and the selector to address this issue. Note that these changes not only enable the GAN training but also greatly improve the efficiency of the original NARF.

3.3 Efficiency Improvements on NARF

Recently, Chan et al. [8] proposed a hybrid explicit-implicit 3D-aware network that uses a memory-efficient tri-plane representation to explicitly store features on axis-aligned planes. With this representation, the efficiency of feature extraction for a 3D location is greatly improved. Instead of forwarding all the sampled 3D points through the network, the intermediate features of arbitrary 3D points can be obtained via simple lookups on the tri-planes. The tri-plane representation can be more efficiently generated with Convolutional Neural Networks (CNNs) instead of MLPs. The intermediate features of 3D points are then transformed into the color and density using a lightweight decoder MLP. The decoder significantly increases the non-linearity between features at different positions, thus greatly enhancing the expressiveness of the model.

Here, we adapt the tri-plane representation to NARF for more efficient training. Similar to [8], we first divide the original NARF network G into an intermediate feature generator (the first linear layer of G) W and a decoder network G_{dec} . Then, Equation 1 is rewritten as follows.

$$c, \sigma = G_{\text{dec}}(\mathbf{f})$$
, where $\mathbf{f} = W(\text{Cat}(\{\gamma(\mathbf{x}_k^l) * p^k | k \in \{1, ..., K\}\})),$ (4)

where **f** is an intermediate feature vector of input 3D location **x**. However, reimplementing the feature generator W to produce the tri-plane representation is not straightforward because its input, a weighted concatenation of \mathbf{x}_k^l , does not form a valid location in a specific 3D space. We make two important changes to address this issue. First, we decompose W into K sub-matrices $\{W_k\}$, one for each part, where each one takes the corresponding local position \mathbf{x}_k^l as input and outputs an intermediate feature for the k^{th} part. Then, the intermediate feature in Equation 4 can be equivalently rewritten as follows.

$$\mathbf{f} = \sum_{k=1}^{K} p^{k} * \mathbf{f}^{k} , \text{ where } \mathbf{f}^{k} = W_{k}(\gamma(\mathbf{x}_{k}^{l})),$$
 (5)

where \mathbf{f}^k is a feature generated in the local coordinate system of the k^{th} part. Now, $W_k(\gamma(\mathbf{x}_k^l))$ can be directly re-implemented by tri-planes F. However, the computational complexity of this implementation is still proportional to K. In order to train a single tri-plane for all parts, the second change is to further transform the local coordinates \mathbf{x}_k^l into a canonical space defined by a canonical pose o^c , similar to Animatable NeRF [52].

$$\mathbf{x}_k^c = R_k^c \mathbf{x}_k^l + \mathbf{t}_k^c = R_k^c (R^k)^{-1} (\mathbf{x} - \mathbf{t}^k) + \mathbf{t}_k^c, \tag{6}$$

where R_k^c and \mathbf{t}_k^c are the rotation and translation matrices of the canonical pose. Intuitively, \mathbf{x}_k^c is the corresponding point location of \mathbf{x} transformed into the canonical space when \mathbf{x} is considered to belong to the k^{th} part. Finally, the tri-plane feature F is learned in the canonical space. The feature extraction for location \mathbf{x} is achieved by retrieving the 32-dimensional feature vector \mathbf{f}^k on the tri-plane F at \mathbf{x}_k^c for all parts, then taking a weighted sum of those features as in Equation 5,

$$\mathbf{f} = \sum_{k=1}^{K} p^{k} * \mathbf{f}^{k} , \text{ where } \mathbf{f}^{k} = \sum_{ij \in \{xy, yz, xz\}} F_{ij}(\mathbf{x}_{k}^{c}).$$
(7)

 $F_{**}(\mathbf{a})$ is a retrieved feature vector from each axis-aligned plane at location \mathbf{a} .

T.2

We estimate the RGB color c and density σ from \mathbf{f} using a lightweight decoder network G_{dec} consisting of three FC layers with a hidden dimension of 64 and output dimension of 4. We apply volume rendering [40] on the color and density to output an RGB image \mathbf{C} and a foreground mask \mathbf{M} .

Although we efficiently parameterize the intermediate features, the probability p^k needs to be computed for every 3D point and every part. In the original NARF, though lightweight MLPs are used to estimate the probabilities, they are still computationally infeasible.

Therefore, we propose an efficient selector network using tri-planes. Since p^k is used to mask out features of irrelevant parts, this probability can be a rough

approximation of the shape of each part. Thus the tri-plane representation is expressive enough to model the probability. We use K separate 1-channel tri-planes to represent P^k , the part probability projected to each axis-aligned plane. We retrieve the three probability values $(p_{xy}^k, p_{xz}^k, p_{yz}^k) = (P_{xy}^k(\mathbf{x}_k^c), P_{xz}^k(\mathbf{x}_k^c), P_{yz}^k(\mathbf{x}_k^c))$ of the k^{th} part by querying the 3D location in the canonical space \mathbf{x}_k^c . The probability p^k that \mathbf{x}_k^c belongs to the k^{th} part is approximated as $p^k = p_{xy}^k p_{xz}^k p_{yz}^k$.

In this way, the features F and part probabilities P are modeled efficiently with a single tri-plane representation. The tri-plane is represented by a $(32 + K) \times 3$ channel image. The first 96 channels represent the tri-plane features in the canonical space. The remaining 3K channels represent the tri-plane probability maps for each of the K parts. We call this approach Efficient NARF, or ENARF.

3.4 GAN

To condition the generator on latent vectors, we utilize a StyleGAN2 [30] based generator to produce tri-plane features. We condition each layer of the proposed ENARF by the latent vector with a modulated convolution [30]. Since the proposed tri-plane based generator can only represent the foreground object, we use an additional StyleGAN2 based generator for the background.

We randomly sample latent vectors for our tri-plane generator and background generator: $\mathbf{z} = (\mathbf{z}_{\text{tri}}, \mathbf{z}_{\text{ENARF}}, \mathbf{z}_b) \sim \mathcal{N}(0, I)$, where $\mathbf{z}_{\text{tri}}, \mathbf{z}_{\text{ENARF}}$, and \mathbf{z}_b are latent vectors for the tri-plane generator, ENARF, and background generator, respectively. The tri-plane generator G_{tri} generates tri-plane feature F and part probability P from randomly sampled \mathbf{z}_{tri} and bone length $\{l_k\}_{k=1:K}$. G_{tri} takes l_k as inputs to account for the diversity of bone lengths.

$$F, P = G_{\text{tri}}(\mathbf{z}_{\text{tri}}, \{l_k\}_{k=1:K})$$
(8)

The ENARF based foreground generator G_{ENARF} generates the foreground RGB image \mathbf{C}_f and mask \mathbf{M}_f from the generated tri-planes, and the background generator G_{b} generates background RGB image \mathbf{C}_b .

$$\mathbf{C}_{f}, \mathbf{M}_{f} = G_{\text{ENARF}}(\mathbf{z}_{\text{ENARF}}, \{l_{k}, R_{k}, \mathbf{t}_{k}\}_{k=1:K}), \mathbf{C}_{b} = G_{b}(\mathbf{z}_{b})$$
(9)

The final output RGB image is $\mathbf{C} = \mathbf{C}_f + \mathbf{C}_b * (1 - \mathbf{M}_f)$, which is a composite of \mathbf{C}_f and \mathbf{C}_b . To handle the diversity of bone lengths, we replace Equation 6 with one normalized by the length of the bone: $\mathbf{x}_k^c = \frac{l_k^c}{l_k} R_k^c \mathbf{x}_k + \mathbf{t}_k^c$, where l_k^c is the bone length of the k^{th} part in the canonical space.

We optimize these generator networks with GAN training. We use a bone loss in addition to an adversarial loss on images, R1 regularization on the discriminator [38], and L2 regularization on the tri-planes. The bone loss ensures that an object is generated in the foreground. Based on the input pose of the object, a skeletal image *B* is created, where pixels with skeletons are 1 and others are 0, and the generated mask *M* at pixels with skeletons is made close to 1: $\mathcal{L}_{\text{bone}} = \frac{\sum_{r \in \mathcal{R}} (1-M)^2 B}{\sum_{r \in \mathcal{R}} B}$. Additional details are provided in the supplement. The final loss is the linear combination of these losses.

3.5 Dynamic Scene Overfitting

Since ENARF is an improved version of NARF, we can directly use it for single dynamic scene overfitting. For training, we use the ground truth 3D pose and foreground mask of each frame and optimize the reconstruction loss in Equation 2.

If the object shape is strictly determined by the poses that comprise the kinematic motion, we can use the same tri-plane features for the entire sequence and directly optimize them. However, real-world objects have time or pose-dependent non-rigid deformation such as clothing and facial expression change in a single sequence. Therefore, the tri-plane features should change depending on time and pose. We use a technique based on deformation fields [54, 50] proposed in time-dependent NeRF, also known as Dynamic NeRF. Deformation field based methods learn a mapping network from observation space to canonical space and learn the NeRF in the canonical frame. Since learning the deformation field with an MLP is expensive, we also approximate it with tri-planes. We approximate the deformation in 3D space by independent 2D deformations in each tri-plane. First, a StyleGAN2 [30] generator takes positionally encoded time t and a rotation matrix of each part R_k and generates 6-channel images representing the relative 2D deformation from the canonical space of each tri-plane feature. We deform the tri-plane feature based on the generated deformation. Please refer to the supplement for more details. We use constant tri-plane probabilities P for all frames since the object shares the same coarse part shape throughout the entire sequence. The remaining networks are the same. We refer to this method as D-ENARF.

4 Experiments

Our experimental results are presented in two parts. First, in Section 4.1, we compare the proposed Efficient NARF (ENARF) with the state-of-the-art methods [47, 52] in terms of both efficiency and effectiveness, and we conduct ablation studies on the deformation modeling and the design choices for the selector. Second, in Section 4.2, we present our results of using adversarial training on ENARF, namely, ENARF-GAN, and compare it with baselines. Then, we discuss the effectiveness of the pose prior and the generalization ability of ENARF-GAN.

4.1 Training on a Dynamic Scene

Following the training setting in Animatable NeRF [52], we train our ENARF model on synchronized multi-view videos of a single moving articulated object. The ZJU mocap dataset [53] consisting of three subjects (313, 315, 386) is used for training. We use the same pre-processed data provided by the official implementation of Animatable NeRF. All images are resized to 512×512 . We use 4 views and the first 80% of the frames for training, and the remaining views or frames for testing. In this setting, the ground truth camera and articulated

	Cost			Novel view			Novel pose		
	#Memory	$\# {\rm FLOPS}$	$\operatorname{Time}(s)$	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$
Animatable NeRF [52]	-	-	0.42	28.28	0.9484	0.05818	29.09	0.9507	0.05706
NARF [47].	283.9 GB	15.9T	2.17	30.62	0.9625	0.05228	29.51	0.959	0.05208
ENARF	<u>27.0GB</u>	<u>71.7G</u>	0.47	31.94	0.9655	0.04792	29.66	0.953	0.05702
D-ENARF	27.6GB	354G	0.49	32.93	0.9713	0.03718	30.06	0.9396	0.05205
ENARF w/o selector	23.0GB	70.2G	0.43	29.16	0.9493	0.07234	27.9	0.9377	0.08316
ENARF w/ MLP selector	83.2 GB	337G	1.13	<u>32.27</u>	0.9684	0.04633	<u>29.74</u>	$\underline{0.9573}$	0.05228

 Table 1. Quantitative comparison on dynamic scenes.

object poses, as well as the ground truth foreground mask, are given for each frame. More implementation details can be found in the supplement.

First, we compare our method with the state-of-the-art supervised methods NARF [47] and Animatable NeRF [52]. Our comparison with Neural Actor [35] is provided in the supplement. Note that our method and NARF [47] take ground truth kinematic pose parameters (joint angles and bone lengths) as inputs, while Animatable NeRF needs the ground truth SMPL mesh parameters. In addition, Animatable NeRF requires additional training on novel poses to render novel-pose images, which is not necessary for our model.

Table 1 shows the quantitative results. To compare the efficiency between models, we examine the GPU memory, FLOPS, and the running time used to render an entire image of resolution 512×512 on a single A100 GPU as evaluation metrics. To compare the quality of synthesized images under novel view and novel pose settings, PSNR, SSIM [63], and LPIPS [71] are used as evaluation metrics. Table 1 shows that the proposed Efficient NARF achieves competitive or even better performance compared to existing methods with far fewer FLOPS and 4.6 times the speed of the original NARF. Although the runtime of ENARF is a bit slower than Animatable NeRF (0.05s), its performance is superior under both novel view and novel pose settings. In addition, it does not need extra training on novel poses. Our dynamic model D-ENARF further improves the performance of ENARF with little increased overhead in inference time, and outperforms the state-of-the-arts Animatable NeRF and NARF by a large margin. Qualitative results for novel view and pose synthesis are shown in Figure 3. ENARF produces much sharper images than NARF due to the more efficient explicit-implicit tri-plane representation. D-ENARF, which utilizes deformation fields, further improves the rendering quality. In summary, the proposed D-ENARF method achieves better performance in both image quality and computational efficiency. Ablation Study To evaluate the effectiveness of the tri-plane based selector, we compare our method against models using an MLP selector or without a selector. A quantitative comparison is provided in Table 1, and a qualitative comparison is provided in the supplement. Although an MLP based selector improves the metrics a bit, it results in a significant increase in testing time. In contrast, the model without a selector is unable to learn clean/sharp part shapes and textures, because a 3D location will inappropriately be affected by all parts without a selector. These results indicate that our tri-plane based selector is efficient and effective.



Fig. 3. Qualitative comparison of novel view and pose synthesis.

4.2 Unsupervised Learning with GAN

In this section, we train the proposed efficient NARF using GAN objectives without any image-pose pairs or mask annotations.

Comparison with Baselines Since this is the first work to learn an articulated representation without image-pose pairs or mesh shape priors, no exact competitor exists. We thus compare our method against two baselines. The first is a supervised baseline called ENARF-VAE, inspired by the original NARF [47]. Here, a ResNet50 [21] based encoder estimates the latent vectors \mathbf{z} from images, and the efficient NARF based decoder decodes the original images from estimated latent vectors \mathbf{z} and ground truth pose configurations o. These networks are trained with the reconstruction loss defined in Equation 2 and the KL divergence loss on the latent vector \mathbf{z} . Following [47], ENARF-VAE is trained with images with a black background. The second model is called StyleNARF, which is a combination of the original NARF and the state-of-the-art high-resolution 3D-aware image generation model called StyleNeRF [19]. To reduce the computational cost, the original NARF first generates low-resolution features using volumetric rendering. Subsequently, a 2D CNN-based network upsamples them into final images. Additional details are provided in the supplement. Please note that ENARF-VAE is a *supervised* method and cannot handle the background, and StyleNARF loses 3D consistency and thus cannot generate high-resolution geometry.

We use the SURREAL dataset [62] for comparison. It is a synthetic human image dataset with a resolution of 128×128 . Dataset details are given in the supplement. For the pose prior, we use the ground truth pose distribution of the training dataset, where we randomly sample poses from the entire dataset. Please note that we do not use image-pose pairs for these unsupervised methods.

Quantitative results are shown in Table 2. We measure image quality with the Fréchet Inception Distance (FID) [24]. To better evaluate the quality of foreground, we use an extra metric called FG-FID that replaces the background with

Table 2. Quantitative comparison on generative models. * indicates that the methods are modified from the cited papers.



Fig. 4. Learned geometry and disentangled representations on the SURREAL dataset by ENARF-GAN. For each of the generated results, the leftmost three columns show the generated images with background, foreground, and corresponding geometry. The rightmost three images show the results of changing only the viewpoint, object pose, and latent variables for the same results, respectively.

black color, using the generated or ground truth mask. We measure depth plausibility by comparing the real and generated depth map. Although there is no ground truth depth for the generated images, the depth generated from a pose would have a similar depth to the real depth that arises from the same pose. We compare the L2 norm between the inverse depth generated from poses sampled from the dataset and the real inverse depth of them. Finally, we measure the correspondence between pose and appearance following the contemporary work named GNARF [5]. We apply an off-the-shelf 2D human keypoint estimator [12] to both generated and real images with the same poses and compute the Percentage of Correct Keypoints (PCK) between them, which is commonly used for evaluating 2D pose estimators. We report the averaged PCKh@0.5 metric [3] for all keypoints. Details are provided in the supplement. Qualitative results are shown in Figure 5. Not surprisingly, ENARF-VAE produces the most plausible depth/geometry and learns the most accurate pose conditioning among the three since it uses image-pose pairs for supervised training. However, compared to styleNARF, its FID is worse and the images lack photorealism. styleNARF achieves the best FID among the three, thanks to the effective CNN renderer. However, it cannot explicitly render the foreground only or generate accurate geometry of the generated images. In contrast, our method performs volumetric rendering at the output resolution, and the generated geometry perfectly matches the generated foreground image.

Using Different Pose Distribution Obtaining a ground truth pose distribution of the training images is not feasible for in-the-wild images or new categories.



Fig. 5. Qualitative comparison on generative models.

Thus, we train our model with a pose distribution different from the training images. Here, we consider two pose prior distributions. The first uses poses from CMU Panoptic [27] as a prior, which we call the CMU prior. During training, we randomly sample poses from the entire dataset. In addition, to show that our method works without collecting actual human motion capture data, we also create a much simpler pose prior. We fit a multi-variate Gaussian distribution on each joint angle of the CMU Panoptic dataset, and we use randomly sampled poses from the distribution for training. Each Gaussian distribution only defines the rotation angle of each part, which can be easily constructed for novel objects. We call this the random prior.

Quantitative and qualitative results are shown in Table 2 and Figure 5. We can confirm that even when using the CMU prior, our model learns posecontrollable 3D representations with just a slight sacrifice in image quality. When using the random prior, the plausibility of the generated images and the quality of the generated geometry are worse. This may be because the distribution of the random prior is so far from the distribution of poses in the dataset that the learned space of latent vectors too often falls outside the distribution of the actual data. Therefore, we used the truncation trick [29] to restrict the diversity of the latent space, and the results are shown in the bottom row of Table 2. By using the truncation trick, even with a simple prior, we can eliminate latent variables outside the distribution and improve the quality of the generated images and geometry. Further experimental results on truncation are given in the supplement.

Additional Results on Real Images To show the generalization ability of the proposed framework, we train our model on two real image datasets, namely AIST++ [32] and MSCOCO [34]. AIST++ is a dataset of dancing persons with relatively simple backgrounds. We use the ground truth pose distribution for training. MSCOCO is a large scale in-the-wild image dataset. We choose images capturing roughly the whole human body and crop them around the persons. Since 3D pose annotations are not available for MSCOCO, we use poses in CMU Panoptic as the pose prior. Note that we do not use any image-pose or mask supervisions for training. Qualitative results are shown in Figure 6. Experimental results with AIST++, which has a simple background, show that it is possible to generate detailed geometry and images with independent control of viewpoint,



Fig. 6. Qualitative results on AIST++ and MSCOCO.

pose, and appearance. For MSCOCO, two successful and two unsuccessful results are shown in Figure 6. MSCOCO is a very challenging dataset because of the complex background, the lack of clear separation between foreground objects and background, and the many occlusions. Although our model does not always produce plausible results, it is possible to generate geometry and control each element independently. As an initial attempt, the results are promising.

5 Conclusion

In this work, we propose a novel unsupervised learning framework for 3D geometryaware articulated representations. We showed that our framework is able to learn representations with controllable viewpoint and pose. We first propose a computationally efficient neural 3D representation for articulated objects by adapting the tri-plane representation to NARF, then show it can be trained with GAN objectives without using ground truth image-pose pairs or mask supervision. However, the resolution and the quality of the generated images are still limited compared to recent NeRF-based GAN methods; meanwhile, we assume that a prior distribution of the object's pose is available, which may not be easily obtained for other object categories. Future work includes incorporating the neural rendering techniques proposed in 3D-aware GANs to generate photorealistic high-quality images while preserving the 3D consistency and estimating the pose prior distribution directly from the training data.

Acknowledgements. This work was supported by D-CORE Grant from Microsoft Research Asia and partially supported by JST AIP Acceleration Research JPMJCR20U3, Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015, JSPS KAKENHI Grant Number JP19H01115, and JP20H05556 and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo. We would like to thank Haruo Fujiwara, Lin Gu, Yuki Kawana, and the authors of [5] for helpful discussions.

References

- Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Detailed human avatars from monocular video. In: 2018 International Conference on 3D Vision (3DV) (2018)
- 2. Alldieck, T., Xu, H., Sminchisescu, C.: imghum: Implicit generative models of 3d human shape and articulated pose. In: ICCV (2021)
- 3. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2D human pose estimation: New benchmark and state of the art analysis. In: CVPR (2014)
- Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. In: ICCV (2021)
- Bergman, A.W., Kellnhofer, P., Wang, Y., Chan, E.R., Lindell, D.B., Wetzstein, G.: Generative neural articulated radiance fields. arXiv preprint arXiv:2206.14314 (2022)
- Bozic, A., Palafox, P., Zollhofer, M., Thies, J., Dai, A., Nießner, M.: Neural deformation graphs for globally-consistent non-rigid reconstruction. In: CVPR (2021)
- Chan, C., Ginosar, S., Zhou, T., Efros, A.A.: Everybody dance now. In: ICCV (2019)
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3D generative adversarial networks. In: CVPR (2022)
- Chan, E.R., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-GAN: Periodic implicit generative adversarial networks for 3D-aware image synthesis. In: CVPR (2021)
- Chen, X., Zheng, Y., Black, M.J., Hilliges, O., Geiger, A.: SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In: ICCV (2021)
- Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: CVPR (2019)
- 12. Contributors, M.: Openmmlab pose estimation toolbox and benchmark. https://github.com/open-mmlab/mmpose (2020)
- Deng, B., Lewis, J.P., Jeruzalski, T., Pons-Moll, G., Hinton, G., Norouzi, M., Tagliasacchi, A.: NASA neural articulated shape approximation. In: ECCV (2020)
- 14. Esser, P., Sutter, E., Ommer, B.: A variational u-net for conditional appearance and shape generation. In: CVPR (2018)
- Gadelha, M., Maji, S., Wang, R.: 3d shape induction from 2D views of multiple objects. In: 2017 International Conference on 3D Vision (3DV) (2017)
- Garbin, S.J., Kowalski, M., Johnson, M., Shotton, J., Valentin, J.: Fastnerf: Highfidelity neural rendering at 200fps. In: ICCV (2021)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NeurIPS (2014)
- Grigorev, A., Iskakov, K., Ianina, A., Bashirov, R., Zakharkin, I., Vakhitov, A., Lempitsky, V.: Stylepeople: A generative model of fullbody human avatars. In: CVPR (2021)
- 19. Gu, J., Liu, L., Wang, P., Theobalt, C.: StyleNeRF: A style-based 3D aware generator for high-resolution image synthesis. In: ICLR (2022)
- Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Realtime deep dynamic characters. ACM Transactions on Graphics (TOG) 40(4), 1–16 (2021)

- 16 A. Noguchi et al.
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- 22. Henzler, P., Mitra, N.J., Ritschel, T.: Escaping plato's cave: 3D shape from adversarial rendering. In: ICCV (2019)
- Hesse, N., Pujades, S., Black, M.J., Arens, M., Hofmann, U.G., Schroeder, A.S.: Learning and tracking the 3D body shape of freely moving infants from rgb-d sequences. IEEE transactions on pattern analysis and machine intelligence 42(10), 2540–2551 (2019)
- 24. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. In: NeurIPS (2017)
- Jacobson, A., Deng, Z., Kavan, L., Lewis, J.P.: Skinning: Real-time shape deformation (full text not available). In: ACM SIGGRAPH 2014 Courses (2014)
- James, D.L., Twigg, C.D.: Skinning mesh animations. ACM Transactions on Graphics (TOG) 24(3), 399–407 (2005)
- Joo, H., Liu, H., Tan, L., Gui, L., Nabbe, B., Matthews, I., Kanade, T., Nobuhara, S., Sheikh, Y.: Panoptic studio: A massively multiview system for social motion capture. In: ICCV (2015)
- Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. In: NeurIPS (2021)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T.: Analyzing and improving the image quality of stylegan. In: CVPR (2020)
- 31. Lewis, J.P., Cordner, M., Fong, N.: Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In: Proceedings of the 27th annual conference on Computer graphics and interactive techniques (2000)
- 32. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++. In: ICCV (2021)
- 33. Li, Z., Niklaus, S., Snavely, N., Wang, O.: Neural scene flow fields for space-time view synthesis of dynamic scenes. In: CVPR (2021)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
- Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. In: ACM SIGGRAPH Asia (2021)
- Liu, L., Xu, W., Habermann, M., Zollhöfer, M., Bernard, F., Kim, H., Wang, W., Theobalt, C.: Neural human video rendering by learning dynamic textures and rendering-to-video translation. IEEE Transactions on Visualization and Computer Graphics (2020)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM transactions on graphics (TOG) 34(6), 1–16 (2015)
- Mescheder, L., Geiger, A., Nowozin, S.: Which training methods for gans do actually converge? In: ICML (2018)
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3D reconstruction in function space. In: CVPR (2019)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: NeRF: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)

- Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Trans. Graph. 41(4), 102:1–102:15 (2022)
- 42. Nguyen-Phuoc, T., Li, C., Theis, L., Richardt, C., Yang, Y.L.: Hologan: Unsupervised learning of 3d representations from natural images. In: ICCV (2019)
- Nguyen-Phuoc, T.H., Richardt, C., Mai, L., Yang, Y., Mitra, N.: Blockgan: Learning 3d object-aware scene representations from unlabelled images. In: NeurIPS (2020)
- 44. Niemeyer, M., Geiger, A.: Giraffe: Representing scenes as compositional generative neural feature fields. In: CVPR (2021)
- 45. Noguchi, A., Harada, T.: RGBD-GAN: Unsupervised 3D representation learning from natural image datasets via rgbd image synthesis. In: ICLR (2020)
- Noguchi, A., Iqbal, U., Tremblay, J., Harada, T., Gallo, O.: Watch it move: Unsupervised discovery of 3D joints for re-posing of articulated objects. In: CVPR (2022)
- Noguchi, A., Sun, X., Lin, S., Harada, T.: Neural articulated radiance field. In: ICCV (2021)
- 48. Osman, A.A., Bolkart, T., Black, M.J.: Star: Sparse trained articulated human body regressor. In: ECCV. Springer (2020)
- 49. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
- Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: ICCV (2021)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: CVPR (2019)
- 52. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: ICCV (2021)
- 53. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural Body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: CVPR (2021)
- 54. Pumarola, A., Corona, E., Pons-Moll, G., Moreno-Noguer, F.: D-nerf: Neural radiance fields for dynamic scenes. In: CVPR (2021)
- 55. Sara Fridovich-Keil and Alex Yu, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022)
- 56. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: GRAF: Generative radiance fields for 3D-aware image synthesis. In: NeurIPS (2020)
- Shysheya, A., Zakharov, E., Aliev, K.A., Bashirov, R., Burkov, E., Iskakov, K., Ivakhnenko, A., Malkov, Y., Pasechnik, I., Ulyanov, D., et al.: Textured neural avatars. In: CVPR (2019)
- 58. Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3D-structure-aware neural scene representations. In: NeurIPS (2019)
- 59. Su, S.Y., Yu, F., Zollhöfer, M., Rhodin, H.: A-NeRF: Articulated neural radiance fields for learning human shape, appearance, and pose. In: NeurIPS (2021)
- Szabó, A., Meishvili, G., Favaro, P.: Unsupervised generative 3D shape learning from natural images. arXiv preprint arXiv:1910.00287 (2019)
- Tiwari, G., Sarafianos, N., Tung, T., Pons-Moll, G.: Neural-GIF: Neural generalized implicit functions for animating people in clothing. In: ICCV (2021)
- Varol, G., Romero, J., Martin, X., Mahmood, N., Black, M.J., Laptev, I., Schmid, C.: Learning from synthetic humans. In: CVPR (2017)

- 18 A. Noguchi et al.
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13(4), 600–612 (2004)
- Wu, J., Zhang, C., Xue, T., Freeman, B., Tenenbaum, J.: Learning a probabilistic latent space of object shapes via 3D generative-adversarial modeling. In: NeurIPS (2016)
- 65. Xu, H., Alldieck, T., Sminchisescu, C.: H-NeRF: Neural radiance fields for rendering and temporal reconstruction of humans in motion. In: NeurIPS (2021)
- Xu, T., Fujita, Y., Matsumoto, E.: Surface-aligned neural radiance fields for controllable 3d human synthesis. In: CVPR (2022)
- 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)
- 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)
- Yariv, L., Kasten, Y., Moran, D., Galun, M., Atzmon, M., Ronen, B., Lipman, Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. In: NeurIPS (2020)
- Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: PlenOctrees for real-time rendering of neural radiance fields. In: ICCV (2021)
- 71. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
- 72. Zhu, J.Y., Zhang, Z., Zhang, C., Wu, J., Torralba, A., Tenenbaum, J., Freeman, B.: Visual object networks: Image generation with disentangled 3D representations. In: NeurIPS (2018)
- Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3D menagerie: Modeling the 3D shape and pose of animals. In: CVPR (2017)