

SNARF: Differentiable Forward Skinning for Animating Non-Rigid Neural Implicit Shapes

Xu Chen^{1,3} Yufeng Zheng^{1,3} Michael J. Black³ Otmar Hilliges¹ Andreas Geiger^{2,3}

¹ETH Zürich, Department of Computer Science ²University of Tübingen

³Max Planck Institute for Intelligent Systems, Tübingen

Abstract

Neural implicit surface representations have emerged as a promising paradigm to capture 3D shapes in a continuous and resolution-independent manner. However, adapting them to articulated shapes is non-trivial. Existing approaches learn a backward warp field that maps deformed to canonical points. However, this is problematic since the backward warp field is pose dependent and thus requires large amounts of data to learn. To address this, we introduce SNARF, which combines the advantages of linear blend skinning (LBS) for polygonal meshes with those of neural implicit surfaces by learning a forward deformation field without direct supervision. This deformation field is defined in canonical, pose-independent, space, enabling generalization to unseen poses. Learning the deformation field from posed meshes alone is challenging since the correspondences of deformed points are defined implicitly and may not be unique under changes of topology. We propose a forward skinning model that finds all canonical correspondences of any deformed point using iterative root finding. We derive analytical gradients via implicit differentiation, enabling end-to-end training from 3D meshes with bone transformations. Compared to state-of-the-art neural implicit representations, our approach generalizes better to unseen poses while preserving accuracy. We demonstrate our method in challenging scenarios on (clothed) 3D humans in diverse and unseen poses.

1. Introduction

Modeling the shape and deformation of articulated 3D objects has traditionally been achieved by deforming a polygonal mesh via linear blend skinning (LBS) with pose-correctives. However, meshes are inherently limited by their resolution-to-memory ratio and their fixed topology. Therefore, neural implicit surface representations [10, 30, 31, 38] have recently attracted much attention because they provide a resolution-independent, smooth and continuous alternative to discrete meshes. However, updating an im-

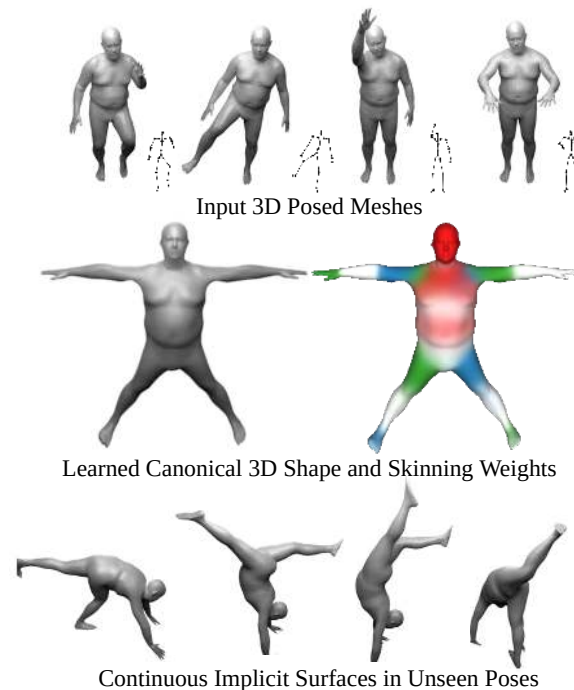


Figure 1: **SNARF**: From a sequence of posed meshes (top), we learn a neural implicit 3D shape and a skinning field in canonical pose (middle) without supervision of skinning weights or part correspondences. Learned forward skinning enables generalization to unseen poses (bottom) while capturing local details via pose conditioning.

PLICIT surface representation as a function of the underlying pose changes is challenging since it requires modifying a continuous function rather than a discrete set of points.

To address this, we propose *SNARF* (*Skinned Neural Articulated Representations with Forward skinning*), a novel approach to learning articulated 3D shapes represented by neural implicit surfaces directly from 3D watertight meshes and corresponding bone transformations with no need for supervision via pre-defined skinning weights. SNARF combines the simplicity of skeletal-driven deformation of LBS with the fidelity and topological flexibility of implicit sur-

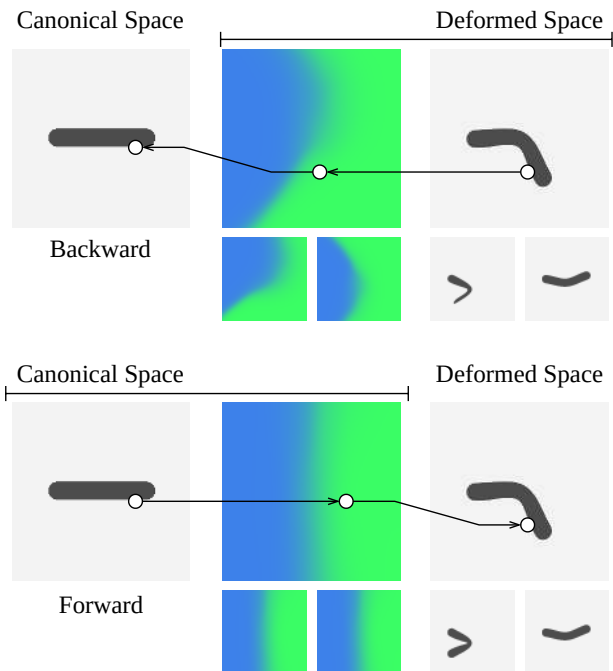


Figure 2: **Forward vs. Backward Skinning.** Forward skinning weights are defined in pose-independent canonical space and therefore naturally generalize to unseen poses as the one in the bottom left panel. In contrast, backward skinning weights are defined in pose-dependent deformed space and thus struggle with unseen poses.

faces, enabling animation of complex human bodies as shown in Fig. 1. Moreover, SNARF goes beyond LBS by conditioning the neural shape on poses to capture pose-dependent non-linear deformations. The main challenge is to express the mapping between surface points in canonical pose and their deformed counterparts. Existing approaches attempt to learn shape in the canonical pose and a *backward* deformation field, transforming deformed points to the canonical pose [12, 34, 39, 43]. However, as illustrated in Fig. 2, backward skinning is problematic since the deformation field depends on the pose of the deformed object, limiting generalization to unseen poses.

To tackle this problem, we devise a method that learns a dense *forward* skinning weight field without requiring direct supervision. Once learned, this skinning field can be leveraged to generate shape deformations even for poses outside of the training set. However, to jointly learn the forward skinning field and the object shape from posed meshes alone, we must establish the correspondence of any 3D point in deformed space to the undeformed space. Yet, this requires the availability of the *backward* mapping which is only implicitly defined and has no analytical solution.

To overcome this issue, we propose a forward skinning model that exploits an iterative root finding algorithm to find the corresponding canonical point for any deformed

point. Our approach is able to retrieve multiple correspondences for any deformed point and therefore naturally handles topology changes. We further derive the gradients of our forward skinning module, hence making it differentiable and enabling end-to-end learning of the canonical shape and skinning weights jointly from deformed observations. Importantly, and in contrast to prior work, our method does not require any a priori skinning weights or pose correctives defined on the surface and hence can be applied in scenarios where pre-rigged mesh models are not available.

We experimentally demonstrate that our method is able to generate high-quality shapes with arbitrary desired bone transformations, even those far beyond the training distribution, where other recent methods like NASA [12] fail. Since our approach operates in continuous space, it enables reconstruction of fine geometric details. By conditioning the neural implicit function on poses, our method faithfully models local pose-dependent deformations, *e.g.*, the movement of clothing or soft tissue. Our code is available at github.com/xuchen-ethz/snarf.

2. Related Work

Skinning Polygonal Meshes: Modeling the deformation of non-rigid and articulated 3D objects is a fundamental problem in computer vision and graphics with many applications. Traditionally, this problem is formulated for polygonal meshes and is referred to as skinning. Skinning enables deformation of a high-resolution surface mesh with low-order control primitives such as skeletal bones. The most common approach is linear blend skinning (LBS), which models each mesh vertex’s deformation as a convex combination of input bone transformations as defined by skinning weights. These skinning weights are typically defined by an artist or learned from data. LBS produces well-known artifacts that many methods attempt to address, *e.g.* with dual quaternion blend skinning [20] or multi-weight enveloping [29, 54]. The key concept is to define pose-dependent “corrective blend shapes” that are added to a shape such that, when it is posed, the LBS errors are minimized [21, 46]. Classically, these “pose correctives” are artist defined, though they can also be learned [24]. Here we extend the concept of LBS and pose correctives to neural implicit surface representations.

Learning both blend weights and rigs from examples has a long history, starting with James and Twigg [17]. Specifically for human bodies, numerous learning methods have been proposed, many of which learn the LBS weights [15, 24, 37, 56]. Recent methods attempt to disentangle shape and pose in an unsupervised fashion given registered training meshes [19, 60]. RigNet [57] uses a deep network to learn both articulated rigs and skinning weights jointly. NeuroSkinning [23] also uses a deep network to learn blend weights and can cope with complex surface topology. In

contrast to us, these methods require a large dataset of rigged models with hand-painted skinning weights and do not consider implicit surface representations.

Neural Implicit Shapes: Neural implicit shape representations can model complex shapes with arbitrary topology in a continuous fashion. Given a 3D location, these networks regress the distance to the surface [38], occupancy probability [30], color [36] or radiance [33] of a 3D point. Conditioning on local information such as 2D image features or 3D point cloud features has been shown to yield more detailed reconstructions [11, 16, 41, 47, 48]. While early methods require watertight meshes for training, several recent approaches have demonstrated unsupervised training from raw 3D points clouds [5, 13, 49] or images [33, 35, 51, 58]. A current limitation of most existing implicit models is that they do not support high-quality skeletal deformation. Our method addresses this key limitation, enabling learning and generation of realistic skeletal deformations of neural implicit surfaces.

Deformable Neural Shapes: Compared to meshes, deforming neural implicit shapes is more challenging as one needs to deform continuous space rather than a fixed set of vertex points. Very recently, various approaches have been proposed to model backward deformation fields [18, 34, 39, 43]. These fields map points in deformed space to canonical ones, where geometric properties (e.g. occupancy) are queried from a canonical shape network. The deformation field is modeled as a neural network that outputs velocity [34], translation [43] or rigid transformation [39] and is jointly trained with the canonical occupancy network using observations in deformed space. NiLBS [18] learns skinning weights for each point and then derives the deformation via LBS according to the bone transformations. An inherent limitation of learned backward deformation, however, is poor generalization to unseen poses. As illustrated in Fig. 2, backward deformation fields are defined in deformed space and, hence, inherently deform with pose. Thus, the network must memorize deformation fields for different spatial configurations, making it difficult to generate deformations that have not been seen during training.

Part-based Models: In recent work, NASA [12], proposes to represent a 3D human body model as a combination of independent parts, each of which is represented by an occupancy network [30]. Rigidly transforming these parts according to the input bone transformations produces deformed shapes. While such a formulation preserves the global structure after articulation, the continuity of surface deformations is violated, causing artifacts at intersections of body parts. Although each part can learn to deform itself to partially compensate for this undesired effect, noticeable artifacts remain, particularly for poses that are beyond the training distribution. Moreover, NASA requires

ground-truth surface skinning weights to learn correct part assignments. In contrast to NASA, our method learns forward skinning weights without such supervision and captures pose-dependent deformations.

More generally, the previous approaches suffer from artifacts due to overly simple assumptions about deformation or do not generalize well to unseen poses as shown in Fig. 4. In contrast, SNARF generates continuous shapes in arbitrary poses, even those far beyond the training distribution, by learning pose-independent forward skinning weights and pose-dependent correctives in canonical space.

3D Human Avatars: While more general, we demonstrate our approach on the problem of learning and animating realistic 3D human avatars. Recent [2, 3, 4, 6, 59] and concurrent works [9, 14, 22, 25, 26, 32, 40, 42, 44, 49, 50, 52, 53, 55] on learning 3D human models typically require a template mesh model with fixed topology, e.g. SMPL [24], or are limited in resolution due to the underlying 3D representation [2, 3, 55, 59]. In contrast, our method is able to represent articulated shapes at high fidelity without strong prior assumptions about the object’s shape. This allows us to better model deformations of objects with more flexible topology, e.g., humans in clothing.

3. Method

In this section, we first define our representation for the canonical shape and forward skinning weights. Next, we introduce our forward mapping and derive the gradients for learning the canonical shape representation and skinning weights in an end-to-end manner.

3.1. Representation

We represent an articulated object by its shape and skinning weights in canonical space. Similar to classical approaches like SMPL, we split the problem into LBS with pose-independent skinning weights and pose-dependent non-linear deformations. LBS captures many important aspects of the shape change, thus the pose-dependent model only has to learn a corrective. This makes training with limited data feasible and aids generalization to unseen poses.

Shape: We use a neural network to predict the occupancy probability for any input 3D point \mathbf{x} in canonical space. To model pose-dependent local deformations such as wrinkles or soft tissue, we inject the object pose \mathbf{p} as additional input:

$$f_{\sigma_f} : \mathbb{R}^3 \times \mathbb{R}^{n_p} \rightarrow [0, 1]. \quad (1)$$

Here, σ_f are the network parameters and n_p is the dimensionality of the pose condition $\mathbf{p} \in \mathbb{R}^{n_p}$ which we specify in terms of joint angles. The canonical shape is implicitly defined as the 0.5 level set of the neural function \mathcal{S} :

$$\mathcal{S} = \{\mathbf{x} \mid f_{\sigma_f}(\mathbf{x}, \mathbf{p}) = 0.5\}. \quad (2)$$

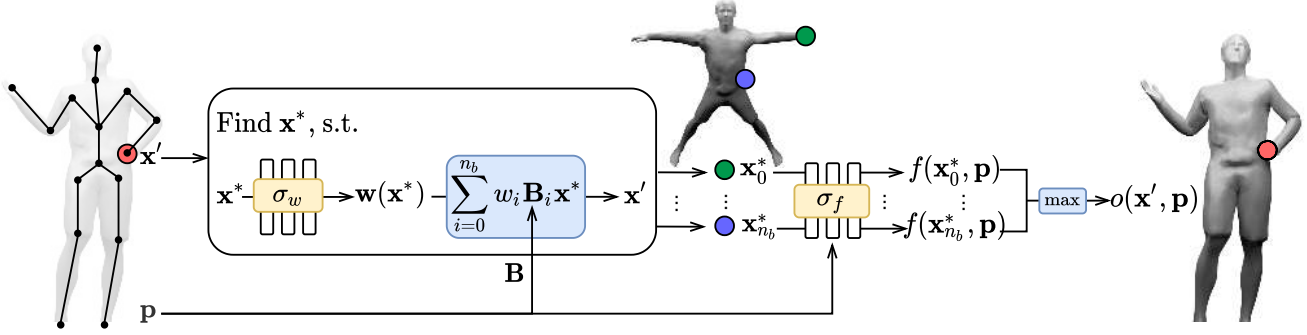


Figure 3: **Generating Deformed Shapes with Forward Skinning.** Given a query point in deformed space \mathbf{x}' , our method first finds its canonical correspondences \mathbf{x}^* which satisfy the forward skinning equation (4) via iterative root finding. Multiple correspondences may exist due to topological changes, which can be reliably found by initializing the root finding algorithm with multiple starting points derived from the bone transformations. The canonical occupancy network f_{σ_f} then predicts the occupancy probabilities at $\{\mathbf{x}^*\}$ which are finally aggregated to yield the occupancy probability of the query point \mathbf{x}' .

Neural Blend Skinning: We model the non-rigid deformation induced by skeleton changes using linear blend skinning (LBS). Towards this goal, we represent an LBS weight field in canonical space using a second neural network:

$$\mathbf{w}_{\sigma_w} : \mathbb{R}^3 \rightarrow \mathbb{R}^{n_b}, \quad (3)$$

where σ_w are the network parameters and n_b denotes the number of bones. Following traditional LBS, we enforce the weights $\mathbf{w} = \{w_1, \dots, w_{n_b}\}$ of each point \mathbf{x} to satisfy $w_i \geq 0$ and $\sum_i w_i = 1$ using a softmax activation function. Note that \mathbf{w}_{σ_w} does not depend on the pose \mathbf{p} .

Given the LBS weights \mathbf{w} of a 3D point \mathbf{x} and the bone transformations $\mathbf{B} = \{\mathbf{B}_1, \dots, \mathbf{B}_{n_b}\}$ corresponding to a particular body pose \mathbf{p} , the deformed point \mathbf{x}' is determined by the following convex combination:

$$\mathbf{x}' = \mathbf{d}_{\sigma_w}(\mathbf{x}, \mathbf{B}) = \sum_{i=1}^{n_b} w_{\sigma_w, i}(\mathbf{x}) \cdot \mathbf{B}_i \cdot \mathbf{x}. \quad (4)$$

Note that the canonical pose is a free hyper-parameter. Empirically, we found the canonical pose shown in Fig. 1 to work well and used it for all experiments on human shapes.

3.2. Differentiable Forward Skinning

To predict the occupancy probability $o'_{\mathbf{x}}$ of an observed 3D point \mathbf{x}' in deformed space, we must first determine the canonical correspondence \mathbf{x}^* of the deformed query \mathbf{x}' in order to evaluate the occupancy $o(\mathbf{x}', \mathbf{p}) = f(\mathbf{x}^*, \mathbf{p})$ with the canonical occupancy network.

At the core of our forward skinning approach lies the problem of finding canonical correspondence \mathbf{x}^* of any query point \mathbf{x}' . This is non-trivial because (i) their relationship is defined implicitly via Eq. (4) without an analytical inverse form, and (ii) multiple canonical points might correspond to the same deformed point as space can overlap after warping (cf. Fig. 3). To address this problem, we propose a

procedure that is able to retrieve all potential canonical correspondences $\{\mathbf{x}_i^*\}$ of any deformed point \mathbf{x}' from the implicitly defined relationship and then composite these correspondences using standard operations for implicit shape composition. An overview is provided in Fig. 3.

Correspondence Search: Unlike backward skinning, forward skinning defines the canonical correspondence \mathbf{x}^* of \mathbf{x}' implicitly as the root of the following equation

$$\mathbf{d}_{\sigma_w}(\mathbf{x}, \mathbf{B}) - \mathbf{x}' = \mathbf{0}, \quad (5)$$

which cannot be solved in closed form. The solution of Eq. (4) can be attained numerically via standard Newton or quasi-Newton methods:

$$\mathbf{x}^{k+1} = \mathbf{x}^k - (\mathbf{J}^k)^{-1} \cdot (\mathbf{d}_{\sigma_w}(\mathbf{x}^k, \mathbf{B}) - \mathbf{x}'), \quad (6)$$

where \mathbf{J} is the Jacobian matrix of $\mathbf{d}_{\sigma_w}(\mathbf{x}^k, \mathbf{B}) - \mathbf{x}'$. To prevent computing the Jacobian at each iteration, we apply Broyden's method [8] using a low-rank approximation of \mathbf{J} .

Handling Multiple Correspondences: We find multiple roots $\{\mathbf{x}_i^*\}$ by initializing the optimization procedure with different starting locations and exploiting the local convergence of iterative root finding. The initial states $\{\mathbf{x}_i^0\}$ are thereby obtained by transforming the deformed point \mathbf{x}' rigidly to the canonical space for each of the n_b bones, and the initial Jacobian matrices $\{\mathbf{J}_i^0\}$ are the spatial gradients of the LBS weight field at the corresponding initial states:

$$\mathbf{x}_i^0 = \mathbf{B}_i^{-1} \cdot \mathbf{x}' \quad \mathbf{J}_i^0 = \left. \frac{\partial \mathbf{d}_{\sigma_w}(\mathbf{x}, \mathbf{B})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}_i^0} \quad (7)$$

Initial states that are far from the optima lead to either convergence to one of the optima and can be safely included for further computation, or divergence, and can therefore be easily discarded by thresholding. Consequently, we define the final set of correspondences as:

$$\mathcal{X}^* = \{\mathbf{x}_i^* \mid \|\mathbf{d}_{\sigma_w}(\mathbf{x}_i^*, \mathbf{B}) - \mathbf{x}'\|_2 < \epsilon\}, \quad (8)$$

where ϵ is the convergence threshold which we set to 10^{-5} in our experiments. This allows us to retrieve all canonical correspondences of any deformed point \mathbf{x}' even under topological changes which induce one-to-many mappings.

Note that if any of the canonical correspondences is occupied, the deformed point \mathbf{x}' is occupied as well. Thus, the maximum over the occupancy probabilities of all canonical correspondences gives the final occupancy prediction:

$$o(\mathbf{x}', \mathbf{p}) = \max_{\mathbf{x}^* \in \mathcal{X}^*} \{f_{\sigma_f}(\mathbf{x}^*, \mathbf{p})\}. \quad (9)$$

This union operator is commonly used to composite independent shapes [45]. Similar to NASA [12], in practice we use softmax instead of a hard maximum to allow gradients to back-propagate to all canonical correspondences.

3.3. Training Losses

Our model is trained via minimizing the binary cross entropy loss $\mathcal{L}_{BCE}(o(\mathbf{x}', \mathbf{p}), o_{gt}(\mathbf{x}'))$ between the predicted occupancy of the deformed points $o(\mathbf{x}', \mathbf{p})$ and the corresponding ground-truth $o_{gt}(\mathbf{x}')$ for all posed 3D meshes of a single subject. In addition, we apply two auxiliary losses during the first epoch of bootstrap training. We randomly sample points along the bones that connect joints in canonical space and encourage their occupancy probabilities to be one. Moreover, we encourage the skinning weights of all joints to be equal to 0.5 for their respective two neighboring bones. No ground truth skinning weights or part segmentations are required by our method.

3.4. Gradients

During training, we must determine the gradient of the overall loss \mathcal{L} w.r.t. the network parameters $\sigma = \{\sigma_f, \sigma_w\}$. For the occupancy network f_{σ_f} , the gradient is given by

$$\frac{\partial \mathcal{L}}{\partial \sigma_f} = \frac{\partial \mathcal{L}}{\partial o} \cdot \frac{\partial o}{\partial f_{\sigma_f}} \cdot \frac{\partial f_{\sigma_f}}{\partial \sigma_f} \quad (10)$$

which can be easily obtained by backpropagating gradients through the corresponding computation graph. For the LBS weight field \mathbf{w}_{σ_w} , the gradient is given by

$$\frac{\partial \mathcal{L}}{\partial \sigma_w} = \frac{\partial \mathcal{L}}{\partial o} \cdot \frac{\partial o}{\partial f_{\sigma_f}} \cdot \frac{\partial f_{\sigma_f}(\mathbf{x}^*)}{\partial \mathbf{x}^*} \cdot \frac{\partial \mathbf{x}^*}{\partial \sigma_w} \quad (11)$$

where \mathbf{x}^* is the root as defined in Eq. (8) and the last term can be analytically obtained via implicit differentiation:

$$\mathbf{d}_{\sigma_w}(\mathbf{x}^*, \mathbf{B}) - \mathbf{x}' = \mathbf{0} \quad (12)$$

$$\Leftrightarrow \frac{\partial \mathbf{d}_{\sigma_w}(\mathbf{x}^*, \mathbf{B})}{\partial \sigma_w} + \frac{\partial \mathbf{d}_{\sigma_w}(\mathbf{x}^*, \mathbf{B})}{\partial \mathbf{x}^*} \cdot \frac{\partial \mathbf{x}^*}{\partial \sigma_w} = \mathbf{0} \quad (13)$$

$$\Leftrightarrow \frac{\partial \mathbf{x}^*}{\partial \sigma_w} = - \left(\frac{\partial \mathbf{d}_{\sigma_w}(\mathbf{x}^*, \mathbf{B})}{\partial \mathbf{x}^*} \right)^{-1} \cdot \frac{\partial \mathbf{d}_{\sigma_w}(\mathbf{x}^*, \mathbf{B})}{\partial \sigma_w}. \quad (14)$$

4. Experiments

We first conduct toy experiments on synthetic 2D data to analyze different methods and model design choices in a controlled setting. Next, we apply our approach to model minimally clothed human bodies and compare it to NASA [12] and other self-implemented baselines. Finally, we demonstrate that our method can handle clothed humans, generalizing well to unseen poses.

4.1. Datasets

We use the following datasets in our experiments:

2D Stick: We simulate a 2D stick articulated by two bones. We set the true skinning weights of each point as the the inverse of its distance to each bone. To simulate topology changes, we include a further rigid object. While this object is separate in canonical space, the two may intersect in posed space and therefore cause topology changes to simulate human self-contact or object interaction.

Minimally Clothed Humans: Following NASA [12], we use the DFaust [7] subset of AMASS [28] for training and evaluating our model on SMPL meshes of people in minimal clothing. This dataset covers 10 subjects of varying body shapes. For each subject, we use 10 sequences, from which we randomly select one sequence for validation, using the rest for training. For each frame in a sequence, 20K points are sampled, among which, half are sampled uniformly in space and half are sampled in near-surface regions by first applying Poisson disk sampling on the mesh surface, followed by adding isotropic Gaussian noise with $\sigma = 0.01$ to the sampled point locations. Besides the “within distribution” evaluation on DFaust, we also include another subset named PosePrior [1] from AMASS for an “out of distribution” evaluation. This dataset contains natural, more challenging, poses beyond those in DFaust.

Clothed Humans: We use the registered meshes from CAPE [27] and corresponding joints and bone transformations derived from the accompanied SMPL model registration to train our model. We use 8 subjects from the dataset with different clothing types including short/long lower body clothing and short/long upper body clothing. We train a model for each subject and clothing condition.

4.2. Baselines

We consider the following baselines in our evaluation. For “Back-LBS”, “Back-D” and “Pose-ONet” we use the same training losses and hyperparameters as in our approach.

Pose-Conditioned Occupancy Networks (Pose-ONet): This baseline extends Occupancy Networks [30] by directly concatenating the pose input to the occupancy network.

Backward Skinning (Back-LBS): This baseline imple-




Figure 4: **Qualitative Results on 2D Toy Experiment.** Row 1 (2 bones): Our deformed shape appears similar to the ground-truth. In contrast, Back-LBS and Pose-ONet produce distorted shapes. The piecewise rigid model (Piecewise) leads to artifacts around bone intersections. Row 2 (2 bones + 1 rigid object): Our forward skinning algorithm can handle topology changes while artifacts at the intersection are noticeable in the result of the backward skinning baseline (Back-LBS).

ments the concept of backward skinning similar to [18]. A network takes a deformed point and pose condition as input and outputs the skinning weights of the deformed point. The deformed point is then warped back to canonical space via LBS and the canonical correspondence is fed into the canonical shape network to query occupancy.

Backward Displacement (Back-D): This baseline directly predicts the displacement from deformed space to canonical space, similar to D-NeRF [18].

NASA: NASA [12] models articulated human bodies as a composition of multiple parts, each of which transforms rigidly and deforms according to the pose. Note that in contrast to us, NASA requires ground-truth skinning weights for surface points as supervision. We use the official NASA implementation provided by the authors.

Piecewise: For evaluation on the 2D toy dataset, we created a variant of NASA for 2D which we refer to as “Piecewise”.

4.3. Results on 2D Stick Dataset

For our results on the simple 2D stick dataset, we do not use local pose-conditioning as the shape does not locally deform with pose. We consider the following three settings:

Extrapolation: An essential requirement for articulated models is the ability to deform into arbitrary poses. In this setting, we generate training data using the articulated 2D stick with joint angles from the interval $[-60, 60]^\circ$. At test time, the models are tasked to generate deformed shapes with larger joint angles in $[-120, 60]^\circ \cup [60, 120]^\circ$. Fig. 4 and Fig. 5 (left) show our results. While our forward skinning model follows the ground truth closely, Pose-ONet fails to generate a meaningful structure as it learns a direct mapping from poses to shapes and thus cannot produce un-

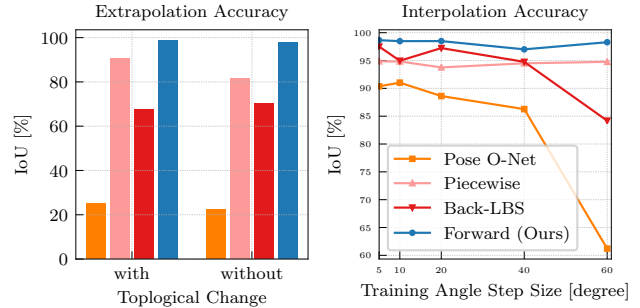


Figure 5: **Quantitative Results on 2D Toy Experiment.** Left: For pose extrapolation, our method outperforms all baselines on both test cases, with and without topological changes. Right: When interpolating, the performance gap increases as training angles are sampled more sparsely.

seen shapes. By disentangling deformations from shapes, Back-LBS preserves the structure better, but the learned pose-dependent skinning weights do not generalize. The piecewise model (Piecewise) generates the correct global pose configuration but exhibits visible artifacts as the rigidity assumption is violated at the joint.

Topological Changes: To simulate topological changes, we include a rigid object but otherwise keep the setting the same as in the previous experiment. Changing topology is challenging for Back-LBS since it is not able to model one-to-many backward correspondences. To compensate for this, the occupancy field gets distorted as shown in Fig. 4. In contrast, our model gracefully handles topological changes, as also shown quantitatively in Fig. 5 (left).

Interpolation: To assess interpolation performance, we evaluate the accuracy of the generated shapes with angles sampled continuously from $[-60, 60]^\circ$ while increasing the sampling step size of the training poses. As shown in Fig. 5 (right), with increasing difficulty, the gap between the baseline methods (Pose O-Net and Back-LBS) and ours becomes larger. An exception is the piecewise model (Piecewise), whose performance is invariant to the training sample density, but instead exhibits artifacts at part intersections.

4.4. Results on Minimally Clothed Humans

Following NASA [12], we now consider the more challenging case of modeling articulated 3D human bodies. Human bodies are challenging due to their complex skeletal structure and local deformations that are non-linearly dependent on the bone transformations. While NASA requires ground-truth skinning weights as additional supervision, our method does not require such knowledge.

Within Distribution: Overall, all methods perform well in this relatively simple setting, as shown in Tab. 1. However, our method still provides an improvement over all baselines. In particular, compared to NASA [12], we im-

| Subject | Within Distribution | | | | | | | | | | Out of Distribution | | | | | | | | | | | |
|---------|---------------------|--------|---------------|--------|---------------|--------|--------|---------------|-------------|---------------|---------------------|--------|--------|--------|---------------|--------|--------|--------|-------------|---------------|--------|---------------|
| | P-ONet | Back-D | IoU bbox | | NASA | Ours | P-ONet | Back-D | IoU surface | | NASA | Ours | P-ONet | Back-D | IoU bbox | | NASA | Ours | IoU surface | | NASA | Ours |
| 50002 | 84.80% | 87.89% | 47.34% | 96.56% | 97.50% | 63.86% | 66.42% | 85.41% | 84.02% | 89.57% | 60.61% | 70.02% | 73.42% | 87.71% | 94.51% | 31.94% | 39.84% | 71.01% | 60.25% | 79.75% | 60.25% | 79.75% |
| 50004 | 80.09% | 84.52% | 93.53% | 96.31% | 97.84% | 57.79% | 59.93% | 88.07% | 85.45% | 91.16% | 55.44% | 64.63% | 65.17% | 86.01% | 95.61% | 34.26% | 38.62% | 69.43% | 62.53% | 83.34% | 62.53% | 83.34% |
| 50007 | 88.31% | 89.09% | 50.13% | 96.72% | 97.96% | 67.14% | 68.02% | 83.46% | 86.28% | 91.02% | 40.53% | 59.68% | 62.66% | 80.22% | 93.99% | 17.80% | 34.76% | 59.53% | 51.82% | 77.08% | 51.82% | 77.08% |
| 50009 | 71.67% | 74.75% | 65.36% | 94.94% | 96.68% | 50.87% | 53.96% | 85.38% | 84.52% | 89.19% | 38.17% | 50.18% | 63.34% | 78.15% | 91.22% | 23.24% | 30.85% | 64.40% | 55.86% | 75.84% | 55.86% | 75.84% |
| 50020 | 69.21% | 73.37% | 93.04% | 95.75% | 96.27% | 48.73% | 53.72% | 86.03% | 87.57% | 88.81% | 42.66% | 52.43% | 64.98% | 83.06% | 93.57% | 26.56% | 33.62% | 68.24% | 62.01% | 81.37% | 62.01% | 81.37% |
| 50021 | 79.30% | 79.48% | 96.86% | 95.92% | 96.86% | 57.80% | 64.02% | 89.96% | 87.01% | 90.16% | 45.50% | 58.99% | 69.89% | 81.80% | 93.76% | 29.07% | 37.19% | 61.69% | 65.49% | 81.49% | 65.49% | 81.49% |
| 50022 | 86.60% | 90.59% | 97.60% | 97.94% | 97.96% | 66.82% | 74.27% | 93.51% | 91.91% | 92.06% | 52.17% | 60.41% | 67.83% | 87.54% | 94.67% | 33.00% | 34.71% | 73.46% | 70.23% | 83.37% | 70.23% | 83.37% |
| 50025 | 80.14% | 79.81% | 95.28% | 95.50% | 97.54% | 59.47% | 60.37% | 87.33% | 86.19% | 91.25% | 52.78% | 56.93% | 68.91% | 83.14% | 94.48% | 31.37% | 34.49% | 70.60% | 60.88% | 82.48% | 60.88% | 82.48% |
| 50026 | 79.39% | 84.58% | 97.32% | 96.65% | 97.64% | 60.52% | 64.07% | 90.17% | 87.72% | 91.09% | 56.09% | 64.33% | 65.20% | 84.58% | 94.13% | 32.07% | 37.71% | 71.85% | 59.78% | 80.01% | 59.78% | 80.01% |
| 50027 | 73.91% | 76.71% | 80.33% | 95.53% | 96.80% | 53.91% | 57.46% | 85.04% | 86.13% | 89.47% | 48.22% | 57.00% | 67.86% | 83.97% | 93.76% | 27.56% | 32.56% | 70.55% | 61.82% | 81.81% | 61.82% | 81.81% |
| Avg. | 79.34% | 82.08% | 81.68% | 96.14% | 97.31% | 58.61% | 62.22% | 87.44% | 86.98% | 90.38% | 49.21% | 59.46% | 66.93% | 83.16% | 93.97% | 28.69% | 35.44% | 68.93% | 60.21% | 80.65% | 60.21% | 80.65% |

Table 1: **Quantitative Results on Minimally Clothed Humans.** The mean IoU of uniformly sampled points in space (IoU bbox) and points near the surface (IoU surface) are reported. Our method outperforms all baselines including NASA [12]. Improvements are more pronounced for points near the surface, and for poses outside the training distribution.




Figure 6: **Qualitative Results on Minimally Clothed Humans.** Our method produces results similar to the ground-truth with correct body pose and plausible local details, both for mild poses within the training distribution and more extreme poses. In contrast, the baseline methods suffer from various artifacts including incorrect poses (Pose-ONet), degenerate shapes (Pose-ONet, Back-D, LBS), and discontinuities near joints (NASA) which become more severe for unseen poses.



Figure 7: **Qualitative Results for Clothed Humans.** Our method can model 3D clothed humans in various clothing types, with rich details including wrinkles, and in novel poses. Moreover, our method faithfully learns the non-linear relationship between cloth deformations and body poses. On the right, we show a failure case where the cloth does not fall naturally for an extreme, unseen pose. However, note how our method still degrades gracefully in this situation.

prove the IoU of uniformly sampled points by 1.2% and the IoU of near-surface points by 4.6%. This improvement can also be observed in the qualitative results Fig. 6. Our method produces bodies with smooth surfaces and correct poses. In contrast, NASA suffers from discontinuous artifacts near joints. Back-D, Back-LBS and Pose-ONet suffer from missing body parts.

Out of Distribution: In this setting, we test the trained models on a different dataset, PosePrior [1], to assess the performance in more realistic settings, where poses can be far from those in the training set. Similar to the observations in the 2D toy setting, unseen poses may cause drastic performance degradation to the baseline methods as shown in Tab. 1. In contrast, our method degrades gracefully despite test poses being drastically different from training poses and very challenging. Hence, the performance gap on IoU surface between our method and NASA increases from 4.6% to 20.4%. As can be seen in Fig. 6, our method generates natural shapes for the given poses while NASA fails to generate correctives at bone intersections for unseen poses, leading to noticeable artifacts. Pose-ONet and Back-D fail to generate meaningful shapes and Back-LBS produces distorted bodies due to incorrect skinning weights.

Learned Skinning Weights: We demonstrate our learned skinning weights in Fig. 1. Our model learns plausible skinning weights with smooth transitions for all moving body parts, reflecting the correct body part assignment. More results can be found in the supplementary material.

4.5. Results on Clothed Humans

Our method can also be applied to modeling clothed humans. We train SNARF using meshes from the CAPE dataset. The results are shown in Fig. 7. Our method is able to model different clothing types with flexible topology and generates realistic results in novel poses with plausible local details, such as wrinkles. The clothing deforms naturally with the body pose, except for very extreme poses where prediction quality degrades gracefully. Fig. 8 shows the canonical shapes corresponding to different body poses.



Figure 8: **Pose-dependent non-linear deformations (correctives) in canonical space.** The heatmaps show the differences (yellow=large, zoom in for more details) between the canonical shape for the current pose and the one for the canonical pose, demonstrating the flexibility of the deformations that can be captured by our model.

5. Conclusion

In this paper, we proposed a differentiable forward skinning model for articulating neural implicit surfaces. Our method learns continuous pose-conditioned shapes and skinning weights from meshes and is able to generate plausible shapes in nearly arbitrary poses. We obtain state-of-the-art results on articulated neural implicit representations for 3D human bodies and demonstrate significantly better generalization to unseen poses than the baselines. We show SotA results on challenging cases of (clothed) 3D humans with diverse shapes and poses. In future work, we plan to extend our method to learn across subjects and from images only using differentiable rendering [35].

Acknowledgements: Xu Chen and Yufeng Zheng were supported by the Max Planck ETH Center for Learning Systems. Andreas Geiger was supported by the DFG EXC number 2064/1 - project number 390727645.

Disclosure: MJB has received research gift funds from Adobe, Intel, Nvidia, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his research was performed solely at, and funded solely by, Max Planck. MJB has financial interests in Amazon, Datagen Technologies, and Meshcapade GmbH.

References

- [1] Ijaz Akhter and Michael J. Black. Pose-conditioned joint angle limits for 3D human pose reconstruction. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015. 5, 8
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *Proc. of the International Conf. on 3D Vision (3DV)*, 2018. 3
- [3] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2Shape: Detailed full human body geometry from a single image. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 3
- [4] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape completion and animation of people. *ACM Trans. on Graphics*, 24(3):408–416, 2005. 3
- [5] Matan Atzmon and Yaron Lipman. SAL: Sign agnostic learning of shapes from raw data. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [6] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Loopreg: Self-supervised learning of implicit surface correspondences, pose and shape for 3D human mesh registration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [7] Federica Bogo, Javier Romero, Gerard Pons-Moll, and Michael J. Black. Dynamic FAUST: Registering human bodies in motion. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [8] Charles G Broyden. A class of methods for solving non-linear simultaneous equations. *Mathematics of computation*, 19(92):577–593, 1965. 4
- [9] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [11] Julian Chibane, Thiemo Alldieck, and Gerard Pons-Moll. Implicit functions in feature space for 3D shape reconstruction and completion. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [12] Boyang Deng, JP Lewis, Timothy Jeruzalski, Gerard Pons-Moll, Geoffrey Hinton, Mohammad Norouzi, and Andrea Tagliasacchi. Neural articulated shape approximation. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 2, 3, 5, 6, 7
- [13] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proc. of the International Conf. on Machine Learning (ICML)*, 2020. 3
- [14] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Trans. on Graphics*, 40(4), aug 2021. 3
- [15] Nils Hasler, Thorsten Thormählen, Bodo Rosenhahn, and Hans-Peter Seidel. Learning skeletons for shape and pose. *ACM Trans. on Graphics*, page 23–30, 2010. 2
- [16] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-PIFu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [17] Doug L. James and Christopher D. Twigg. Skinning mesh animations. *ACM Trans. on Graphics*, 24(3):399, 2005. 2
- [18] Timothy Jeruzalski, David IW Levin, Alec Jacobson, Paul Lalonde, Mohammad Norouzi, and Andrea Tagliasacchi. NiLBS: Neural inverse linear blend skinning. *arXiv.org*, 2004.05980, 2020. 3, 6
- [19] B. Jiang, J. Zhang, J. Cai, and J. Zheng. Disentangled human body embedding based on deep hierarchical neural network. *IEEE Transactions on Visualization and Computer Graphics*, 26(8):2560–2575, 2020. 2
- [20] Ladislav Kavan, Steven Collins, Jiri Zara, and Carol O’Sullivan. Geometric skinning with approximate dual quaternion blending. *ACM Trans. on Graphics*, 27(4):105, 2008. 2
- [21] J. P. Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. *ACM Trans. on Graphics*, pages 165–172, 2000. 2
- [22] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *arXiv.org*, 2004.05980, 2021. 3
- [23] Lijuan Liu, Youyi Zheng, Di Tang, Yi Yuan, Changjie Fan, and Kun Zhou. Neuroskinning: Automatic skin binding for production characters with deep graph networks. *ACM Trans. on Graphics*, 38(4), July 2019. 2
- [24] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. on Graphics*, 2015. 2, 3
- [25] Qianli Ma, Shunsuke Saito, Jinlong Yang, Siyu Tang, and Michael J. Black. SCALE: Modeling clothed humans with a surface codec of articulated local elements. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [26] Qianli Ma, Jinlong Yang, , Siyu Tang, and Michael J Black. The power of points for modeling humans in clothing. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [27] Qianli Ma, Jinlong Yang, Anurag Ranjan, Sergi Pujades, Gerard Pons-Moll, Siyu Tang, and Michael J. Black. Learning to dress 3D people in generative clothing. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 5
- [28] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 5
- [29] Bruce Merry, Patrick Marais, and James Gain. Animation space: A truly linear framework for character animation. *ACM Trans. on Graphics*, 25(4):1400–1423, 2006. 2
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3D reconstruction in function space. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 3, 5

- [31] Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Deep level sets: Implicit surface representations for 3D shape inference. *arXiv.org*, 1901.06802, 2019. 1
- [32] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [33] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 3
- [34] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Occupancy flow: 4D reconstruction by learning particle dynamics. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 2, 3
- [35] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3D representations without 3D supervision. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3, 8
- [36] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 3
- [37] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. Star: Sparse trained articulated human body regressor. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 2
- [38] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 3
- [39] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Deformable neural radiance fields. *arXiv.org*, 2011.12948, 2020. 2, 3
- [40] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2021. 3
- [41] Songyu Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 3
- [42] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [43] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 3
- [44] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. ANR-articulated neural rendering for virtual avatars. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [45] Antonio Ricci. A constructive geometry for computer graphics. *The Computer Journal*, 16(2):157–160, 1973. 5
- [46] C. Rouet and J. Lewis. Method and apparatus for creating lifelike digital representations of computer animated objects by providing corrective enveloping, Mar. 16 1999. US Patent 5,883,638. 2
- [47] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 3
- [48] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [49] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [50] Aliaksandra Shysheya, Egor Zakharov, Kara-Ali Aliiev, Renat Bashirov, Egor Burkov, Karim Iskakov, Aleksei Ivakhnenko, Yury Malkov, Igor Pasechnik, Dmitry Ulyanov, et al. Textured neural avatars. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [51] Vincent Sitzmann, Michael Zollhöfer, and Gordon Wetzstein. Scene representation networks: Continuous 3D-structure-aware neural scene representations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 3
- [52] Shaofei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3D human mesh registration. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [53] Shaofei Wang, Marko Mihajlovic, Qianli Ma, Andreas Geiger, and Siyu Tang. MetaAvatar: Learning animatable clothed human models from few depth images. *arXiv.org*, 2004.05980, 2021. 3
- [54] Xiaohuan Corina Wang and Cary Phillips. Multi-weight enveloping: least-squares approximation techniques for skin animation. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 129–138, 2002. 2
- [55] Chung-Yi Weng, Brian Curless, and Ira Kemelmacher-Shlizerman. Vid2actor: Free-viewpoint animatable person synthesis from video in the wild. *arXiv.org*, 2012.12884, 2020. 3
- [56] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: Generative 3D human shape and articulated pose models. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [57] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. RigNet: Neural rigging for articulated characters. *ACM Trans. on Graphics*, 39(4), 2020. 2
- [58] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan

- Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [59] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3D human reconstruction from a single image. In *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*, 2019. 3
- [60] Keyang Zhou, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Unsupervised shape and pose disentanglement for 3D meshes. In *Proc. of the European Conf. on Computer Vision (ECCV)*, 2020. 2