

# VINECS: Video-based Neural Character Skinning

Zhouyingcheng Liao<sup>1,2</sup>

Vladislav Golyanik<sup>1</sup>

Marc Habermann<sup>1,3</sup>

Christian Theobalt<sup>1,3</sup>

<sup>1</sup>Max Planck Institute for Informatics, Saarland Informatics Campus

<sup>2</sup>The University of Hong Kong

<sup>3</sup>Saarbrücken Research Center for Visual Computing, Interaction and AI

## Abstract

Rigging and skinning clothed human avatars is a challenging task and traditionally requires a lot of manual work and expertise. Recent methods addressing it either generalize across different characters or focus on capturing the dynamics of a single character observed under different pose configurations. However, the former methods typically predict solely static skinning weights, which perform poorly for highly articulated poses, and the latter ones either require dense 3D character scans in different poses or cannot generate an explicit mesh with vertex correspondence over time. To address these challenges, we propose a fully automated approach for creating a fully rigged character with pose-dependent skinning weights, which can be solely learned from multi-view video. Therefore, we first acquire a rigged template, which is then statically skinned. Next, a coordinate-based MLP learns a skinning weights field parameterized over the position in a canonical pose space and the respective pose. Moreover, we introduce our pose- and view-dependent appearance field allowing us to differentially render and supervise the posed mesh using multi-view imagery. We show that our approach outperforms state-of-the-art while not relying on dense 4D scans. More details can be found on our project page<sup>1</sup>.

## 1. Introduction

Animating 3D characters is a long-standing and challenging problem in computer graphics and vision. Traditionally, rigging, the process of positioning a sparse skeletal structure inside a dense 3D mesh, and skinning, the process of associating vertices with the skeleton, require large amounts of manual effort by experienced artists. This is not only expensive and time-consuming but also inherently difficult since one set of skinning weights is typically not ideal for all sorts of articulated character poses.

There are also methods for automated rigging and skinning. Some recent works [33, 34, 41, 59] propose a learning-based solution for obtaining skinning and rigging




Figure 1. We propose the first end-to-end trainable method for generating a dense and rigged 3D character mesh including learned pose-dependent skinning weights solely from multi-view videos. The results demonstrate that our model can generate visually-plausible geometry without requiring manual user input.

given a single 3D character mesh. However, they do not consider pose-dependent skinning correctives and lead to the typical appearance artifacts. Therefore, other techniques focus on pose-dependent skinning formulations accounting for pose-dependent variations and deformations [3, 30, 40, 52, 57]. Noteworthy, none of these methods tried to learn pose-dependent skinning weight correctives purely from 2D image data. Instead, they typically assume a 3D mesh is given, and most of them solely work on a fixed mesh resolution, which cannot be easily adjusted.

To overcome these limitations, we propose VINECS, *i.e.*, Video-based NEural Character Skinning. Given solely a multi-view video of a human actor, our method creates a fully rigged, skinned, and textured 3D template (including hands) along with pose-dependent skinning weights while *no manual editing* is required. Moreover, VINECS enables multi-resolution character skinning since skinning weights can be sampled over a continuous 3D canonical space, which is not limited to any specific mesh resolution.

For each character, we select one canonical pose frame, for which we extract a dense and textured 3D template mesh using implicit surface reconstruction techniques [55]. Given the template and the respective pose obtained from an off-the-shelf multi-view markerless capture system [53],

<sup>1</sup><https://people.mpi-inf.mpg.de/~mhaberma/projects/2023-Vinecs>

we leverage an auto-rigging and skinning method, Pinocchio [7], for obtaining an initial set of skinning weights. Note that Pinocchio produces only a *static* set of skinning weights, which leads to the typical geometry artifacts when poses deviate too much from the canonical pose. Therefore, we propose a coordinate-based multi-layer perceptron (MLP)—which takes as input a 3D point on the surface of the template in canonical pose and the skeletal pose—and predicts the pose-dependent skinning weights at that point. Since we want to learn the skinning weights solely from multi-view imagery, we design dedicated losses consisting of a silhouette loss and a rendering loss. Specifically, we find that obtaining a well-behaved convergence is challenging only using a static texture for the rendering loss. Thus, we adopt the rendering formulation with an albedo network and a shadow network, which predict pose- and view-dependent appearance of our 3D character resulting in an improved convergence and skinning weights prediction. In summary, our technical contributions are as follows:

- An end-to-end trainable method for animation-ready explicit character creation, involving static template generation, rigging, and pose-dependent skinning solely using 2D multi-view videos (as few as 10 views).
- A coordinate-based and pose-dependent skinning formulation, which enables multi-resolution skinning.
- A dedicated character appearance formulation and differentiable rendering enabling weakly supervised learning based on multi-view videos.

We compare VINECS to the state-of-the-art and our results confirm that our method is a clear step toward accurate and automated creation of animatable 3D characters.

**Potential Impact.** The focus of this work is *not* to model deformable geometry [13, 14, 20, 32] or photorealistic appearance [15, 16, 25, 35, 44, 47, 48] of humans but on learning 3D characters with pose-dependent skinning from 2D observations only. We believe our method can greatly benefit many downstream tasks as many existing works leverage naked body models [36] for space canonicalization and feature assignment, fundamentally limiting them to clothing types of the same topology as the human body. As our rigging and skinning are accurate, easy-to-use, and robust to any type of clothing, our approach can stimulate research on the modeling of geometry and appearance of loose apparel.

## 2. Related Work

**Static Skinning.** Linear Blend Skinning (LBS) [28, 38] linearly blends rigid transformations of each bone to obtain a posed geometry. However, this formulation can lead to a collapse of geometry near the skeleton joints, also called *candy-wrappers effect*. To overcome this, several follow-up works have been proposed using spherical representations [22], sweep surfaces [17], curve-based skeleton parameterizations [61], Dual Quaternions [23], or by optimiz-

ing the centers of rotation [27]. Other works focused more on computing the skinning weights itself by modeling the weight distribution as a heat diffusion process [7], as illumination [58], or by voxelizing the mesh to increase robustness [9, 10]. Later on, deformation editing was also modeled using a Laplacian energy formulation [18], and user-defined splines on the mesh enable intuitive editing of skinning-based deformation behaviors [6]. They all only assume a single geometry and skeleton are given.

In contrast, data-driven methods assume multiple instances of the same character under different poses or a large collection of rigged and skinned characters. Earlier works [19, 26] requires meshes under different poses and then automatically detect skeleton joints and compute the skinning weights. Recently, NeuroSkinning [34] was proposed as the first learning-based method predicting a static set of skinning weights given the mesh and a corresponding skeleton. RigNet [59] instead jointly predicts the skeleton as well as a set of skinning weights solely given a 3D mesh as input. Therefore, they collected a large dataset of 3D characters including artist-designed skeletons and skinning weights. SkinningNet [41] proposed a two-stream graph convolutional network architecture for improved skinning weight prediction given a mesh and the skeleton. Yang et al. [60] proposed to predict the 3D shape together with the rigging from a single image. Liao et al. [33] introduced a learning-based method which jointly learns to predict skinning weights and transfer poses.

What all the above methods have in common is that they do not explicitly account for dynamic or pose-dependent effects, such as muscle bulging or coarse cloth deformations. In contrast, our proposed method learns pose-dependent skinning weights, which alleviate typical artifacts that arise when using a static set of skinning weights. Moreover, it can also represent coarse deformations induced by the skeletal pose, such as coarse cloth deformations.

**Dynamic Methods.** More closely related to our method are approaches that account for dynamic or pose-dependent effects, which are discussed next. There are works [11] that allow an artist to add pose-dependent effects by painting onto specialized textures or by adding so-called helper bones [42]. Alternatively, Kavan et al. [21] proposed to find the optimal skinning weights by minimizing an elastic energy and adding joint-based deformer. Interestingly, those methods do not assume any dynamic observations of the character while still modeling dynamic effects. However, they either require manual editing [11, 42] or can mostly model elasticity-based dynamics [21].

Similar to the static methods, several dynamic methods assume a (large) dataset of character meshes is given. Earlier works [3, 40, 46, 57] account for dynamic effects by capturing a set of range scans depicting the subject in different poses. Given a new pose, a set of pose-dependent blend

weights is estimated, which linearly combines the set of example poses in the dataset or alternatively uses linear and radial basis functions [52]. However, since the complexity of computation typically linearly scales with the number of example poses in the dataset, those methods cannot scale up to a very large number of examples (e.g., 20000 pose frames), which is in stark contrast to our approach that can deal with such an amount of data. There are also parametric naked human body models [4, 5, 36] learned by applying PCA over a large set of human scans, which can account for different shape- and pose-dependent deformations. In contrast, we are building a highly-detailed rigged and skinned character that can wear (loose) clothing in a completely automated fashion. Li et al. [30] predict static skinning weights and pose-dependent neural blend shapes. They assume a dataset of dense meshes is provided, whereas we purely learn pose-dependent skinning weights from videos. SCANimate [49] also requires dense point clouds, and learns a static set of skinning weights, while pose-dependent shapes are predicted by a pose-aware implicit field. SNARF [8] represents the character as a pose-dependent neural implicit surface in the canonical pose space while the learned skinning weights are pose-agnostic. Again, they require point cloud data while our method can be supervised directly on the video. In summary, none of the above methods that consider pose-dependent effects investigated learning solely from 2D imagery. Thus, we are the first method demonstrating that dynamic effects in the form of pose-dependent skinning weights can be learned only from videos.

Recently, some works were proposed to skin and rig implicit surface representations learn an implicit human avatar. ARAH [56] learns an implicit human model using neural rendering. However, as it relies on a SMPL template and an inverse root finding to convert from view space to canonical space, it easily fails when the garment deviates much from the body. Importantly, it also does not have the vertex correspondence over time. TAVA [31] avoids these problems by getting rid of a template, which, however, leads to poor geometry recovery mainly due to their density-based geometry representation. While these methods can reconstruct an implicit avatar, they do not learn a high-quality skinning weight, which is important to traditional graphics pipeline.

### 3. Method

Given synchronized and calibrated multi-view RGB videos of a human character in motion, our goal is to learn an animatable 3D human model with dynamic skinning, without any manual process, such as meshing, rigging, and skinning. To this end, we first extract the static geometry of the actor from one canonical pose frame (typically depicting the actor in a T-pose) of the recording and automatically compute initial skinning weights (Sec. 3.1). However, the initial skinning weights are calculated only based on

the static geometry of the human character, and, as studied in earlier works, one set of weights might be ideal for one pose while it can lead to strong artifacts in other poses. Thus, we propose a skinning network, called SkinNet, to predict the dynamic skinning weights as a function of the skeletal pose and 3D coordinates in the canonical space (Sec. 3.2). Given the skeletal pose, the static geometry, and the pose-dependent skinning weights, we leverage Linear Blend Skinning (LBS) to obtain the posed geometry. We then use a differentiable renderer to weakly supervise the skinning weights solely based on multi-view imagery. Thereafter, we propose a novel appearance field (Sec. 3.3), which is composed of an albedo component and a pose- and view-dependent shadow/shading component. In addition, we propose several regularization losses ensuring that the skinning weights are also robust to out-of-distribution poses and that the final animated geometry looks visually plausible (Sec. 3.4). An overview is shown in Fig. 2.

**Data Preparation.** For a human subject, we record a multi-view video of the performance using  $C$  calibrated and synchronized cameras with a total length of  $F$  frames. For all frames  $\mathcal{I}_{c,f}$ , we calculate the foreground segmentation using color keying or background matting [51] and compute respective distance transform images  $\mathcal{D}_{c,f}$ . Here,  $c$  and  $f$  refer to the respective camera and frame index. We capture the size of the human skeleton and the motion using markerless motion capture [53]. The skeleton motion is parameterized by a root rotation  $\alpha_f \in \mathbb{R}^3$ , a root translation  $\mathbf{t}_f \in \mathbb{R}^3$ , and the joint angles  $\rho_f \in \mathbb{R}^D$ . The full pose vector is defined as  $\theta_f = [\alpha_f, \mathbf{t}_f, \rho_f]^T \in \mathbb{R}^{3+3+D}$ .

**Notation.** To represent a variable in the canonical space, we use a bar over the letter, i.e.,  $\bar{\mathbf{x}}$ . For any point in canonical and global space, we use  $\bar{\mathbf{x}}$  and  $\mathbf{x}$ , respectively, and mesh vertices are denoted as  $\bar{\mathbf{v}}_i$  and  $\mathbf{v}_i$  where  $i$  is the  $i$ th vertex. Without loss of generality, we assume the frame  $f$  is fixed in the following (if not explicitly mentioned otherwise) and omit the subscript for better readability.

#### 3.1. Canonical Character Model

As the first step, we acquire a canonical model of the actor. Recently, implicit surface and coordinate-based representations [45] have gained considerable attention due to their flexibility and comparably simple integration into learning frameworks, and they were also leveraged in the context of pose-dependent human deformation modeling [8]. Moreover, we compared the reconstruction accuracy on the first frame  $f = 0$  between classical multi-view stereo and implicit reconstruction method called NeuS [55] in the supplemental material. Interestingly, NeuS has a much higher level of detail and comparably less noise on the surface. Thus, we use NeuS [55] to reconstruct the mesh in the first frame, also called the canonical frame/space.

**Implicit vs. Explicit Canonical Model.** In stark con-

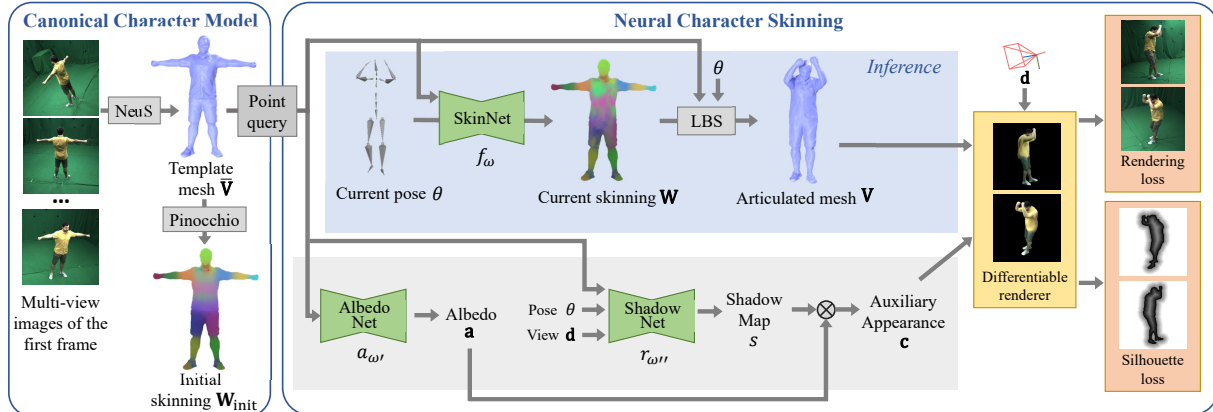


Figure 2. **Method overview.** Our method, VINECS, is a fully automated solution to template generation, rigging, and pose-dependent skinning solely using multi-view videos. We first recover the skeletal poses using markerless motion capture [53] and a static template mesh using implicit surface reconstruction [55]. An initial skinning is obtained by a heat diffusion process [7]. Since one fixed set of skinning weights can lead to artifacts when significantly changing the character pose, we learn pose-dependent skinning weights using a coordinate-based MLP. An auxiliary appearance field is applied to better supervise the learning guide the training of pose-dependent skinning. Note that we do not intend to learn a detailed appearance and only the components in the blue box are used during inference.

trast to Chen et al. [8], we run Marching Cubes [37] to convert the implicit surface into an explicit mesh and obtain per-vertex colors from NeuS by setting the viewing direction to the normal direction of the mesh. The reason for choosing an explicit mesh in canonical space is three-fold: 1.) It avoids running the rather slow Marching Cubes algorithm for every potentially posed mesh. 2.) We avoid the backward skinning [8], which boils down to a correspondence search that has to be solved using iterative optimization. 3.) We can maintain correspondence on the surface across poses, which is especially useful when texturing and other editing have to be applied. Thus, we aim for a hybrid strategy for template acquisition, i.e., using the state-of-the-art method for implicit surface reconstruction and then converting it to an explicit mesh.

**Meshing and Initial Skinning.** During training, we downsample the template mesh to  $N$  vertices where  $\bar{v}_i \in \mathbb{R}^3$  denotes the  $i$ th vertex. Given the template mesh and the corresponding posed skeleton, i.e.  $\theta_f$ , we utilize Pinocchio [7] to obtain an initial set of skinning weights  $\mathbf{w}_{\text{init},i} \in \mathbb{R}^J$  for each vertex  $i$  (see Fig. 4). Here,  $J$  denotes the number of skinning joints. Given these weights and a skeletal pose  $\theta$ , we can transform the vertex  $\bar{v}_i$  from the canonical space to the posed space using LBS [28]:

$$\mathbf{v}_i = \text{LBS}(\bar{v}_i, \mathbf{w}_{\text{init},i}, \theta) : \mathbb{R}^3 \times \mathbb{R}^J \times \mathbb{R}^D \rightarrow \mathbb{R}^3. \quad (1)$$

**Mesh Parsing.** Since different body parts and materials (skin vs. clothing) have different deformation behaviors, we compute per-vertex human parsing labels. We render the 3D model into all camera views and apply a 2D human parsing approach [29] on each view. We then perform max-voting to obtain the per-vertex human parsing labels. Note that so far, we obtained a segmented, rigged, and statically skinned character model *without* any manual user input.

### 3.2. Pose-dependent Skinning Field

So far, the initial skinning weights obtained in Sec. 3.1 have two main limitations: 1.) They are solely computed on a single static pose leading to artifacts such as candy wrappers and mesh distortions if the new pose significantly differs from the canonical one. 2.) The set of weights is computed per-vertex. Thus, the resolution of the mesh has to be fixed beforehand and any change in the mesh requires a new computation of skinning weights or some complicated weight transfer from one mesh to the other.

To resolve this, we propose SkinNet  $f_\omega$ , which is a coordinate-based MLP where  $\omega$  denotes the trainable network weights. Given any 3D point  $\bar{\mathbf{x}}$  in the canonical space,  $f_\omega$  predicts its skinning weight  $\mathbf{w}$  conditioned on the normalized human pose  $\tilde{\theta}$ , which can be formulated as:

$$\mathbf{w}_{\tilde{\theta}} = f_\omega(\bar{\mathbf{x}}, \tilde{\theta}) : \mathbb{R}^3 \times \mathbb{R}^D \rightarrow \mathbb{R}^J, \quad (2)$$

where  $\bar{\mathbf{x}}$  and  $\tilde{\theta}$  are concatenated before being fed into the MLP. Here, the normalized pose  $\tilde{\theta}$  is obtained by neglecting the global translation as well as the yaw angle, i.e., the rotation around the spine. The intuition behind is that the skinning weights do not depend on where in global 3D space the person performs a respective pose, and, similarly, the skinning weight prediction should be agnostic to which direction the person is facing. Thus, we mask those degrees of freedom before feeding the pose vector into the network. To obtain the transformed point  $\mathbf{x}$  in posed space, one can insert Eq. (2) into Eq. (1) resulting in:

$$\mathbf{x} = \text{LBS}(\bar{\mathbf{x}}, f_\omega(\bar{\mathbf{x}}, \tilde{\theta}), \theta). \quad (3)$$

Revisiting the drawbacks of initial skinning discussed above, the SkinNet now supports pose-dependent skinning

weights such that the artifacts of static skinning can be reduced. Moreover, since SkinNet is a coordinate-based MLP, it supports arbitrary mesh resolution since each vertex is queried independently, which is in stark contrast to an architecture that outputs all skinning weights at once and where the final tensor shape is fixed to a specific mesh resolution.

### 3.3. Auxiliary Appearance Field

We can now pose the geometry with our pose-dependent skinning weight formulation. However, since we want to supervise SkinNet solely on multi-view images, we also have to model the appearance of the actor. A naive choice would be to just leverage the static vertex colors acquired from NeuS. However, we found that this texture contains baked-in shadows and wrinkles, and cannot account for the pose- and view-dependent appearance changes harming the training of the SkinNet. We also refer to our ablation studies in Sec. 4.3. To address this issue, we propose an auxiliary appearance field, which consists of an albedo field predicting a pose- and view-agnostic color for a given point  $\bar{\mathbf{x}}$  in the canonical space and a shadow field predicting the pose- and view-dependent shadow/shading properties of  $\bar{\mathbf{x}}$ . Important to note here is that we are *not* interested in creating a highly photorealistic appearance model of the human, but we mainly use it as an auxiliary tool to better supervise the learning of the skinning weight network.

More specifically, the albedo field, referred to as *AlbedoNet*, is an MLP that predicts a static albedo value  $a_{\omega'}(\bar{\mathbf{x}}) = \mathbf{a} \in \mathbb{R}^3$  for a given 3D sample point in the canonical space. Here,  $\omega'$  are the trainable weights of the albedo network. Moreover, the shadow field is also parameterized by an MLP  $r_{\omega''}(\bar{\mathbf{x}}, \theta, \mathbf{n}, \mathbf{d}) = s \in \mathbb{R}^+$ , called *ShadowNet*, which predicts a scalar multiplier accounting for color changes due to shadows and shading. Note that we assume that such effects happen uniformly across all color channels and, thus, we are only predicting a single scalar instead of a scaling factor per color channel. In addition to taking the point in canonical space, the shadow field also takes the skeletal pose  $\theta$  as input as well as the surface normal  $\mathbf{n}$  of  $\mathbf{x}$  in global space and the viewing direction  $\mathbf{d}$ , which allows to potentially model pose- and view-dependent lighting effects such as shadows and shading.  $\omega''$  denotes the respective network weights. Details about the architectures are provided in the supplemental document. The final color  $\mathbf{c} \in \mathbb{R}^3$  for point  $\mathbf{x}$  is computed as

$$\mathbf{c} = a_{\omega'}(\bar{\mathbf{x}})r_{\omega''}(\bar{\mathbf{x}}, \theta, \mathbf{n}, \mathbf{d}). \quad (4)$$

### 3.4. Multi-view Video-based Supervision

Next, we describe our dedicated supervision strategy that solely requires multi-view imagery. We first introduce the individual loss terms and then describe our training strategy. Importantly, we can obtain a posed template mesh

$\mathbf{V} \in \mathbb{R}^{N \times 3}$  and the pose- and view-dependent per-vertex colors  $\mathbf{C}_c \in \mathbb{R}^{N \times 3}$  by evaluating Eq. (3) and Eq. (4) for all vertices  $\bar{\mathbf{v}}_i$  and views  $c$ . Moreover, the following losses that are applied to the deformed geometry and its appearance can directly backpropagate into the respective networks since our formulation is fully differentiable.

**Silhouette Loss.** This loss [15] ensures that the projected posed geometry projects onto the foreground masks for all views (see supp. material). Since our pose-dependent skinning field Eq. (3) is differentiable with respect to the SkinNet weights  $\omega$ , the silhouette loss can provide supervision and learn pose-dependent skinning by matching the posed model silhouette and the image silhouettes.

**Rendering Loss.** Although the silhouette loss ensures that the posed geometry matches the multi-view silhouettes, it cannot supervise drifts on the visual hull. Supervision directly from the RGB images can help with resolving this issue since appearance features can constrain drifts in the image plane. Thus, we introduce the rendering loss

$$\mathcal{L}_{\text{rend}} = \sum_{c=1}^C \|\Pi_c(\mathbf{V}, \mathbf{C}_c) - \mathcal{I}_c\|_1, \quad (5)$$

which leverages a differentiable renderer  $\Pi_c$  based on the one of Habermann et al. [15] to render the posed and colored model into all camera views  $c$  that is then compared to the ground truth image  $\mathcal{I}_c$ . Importantly, this loss can supervise both the pose-dependent skinning field as well as the pose- and view-dependent appearance field.

**Laplacian Loss.** Since the data terms alone can still lead to degenerate results, we also introduce some regularization on the geometry and the skinning weights. First, we regularize the posed geometry with a Laplacian loss

$$\mathcal{L}_{\text{lap}} = \sum_{i=1}^N \left\| \mathbf{v}_i - \frac{1}{|\mathcal{N}_i|} \sum_{k \in \mathcal{N}_i} \mathbf{v}_k \right\|_2^2, \quad (6)$$

where  $\mathcal{N}_i$  is the indices of one-ring neighborhood of vertex  $i$ . This regularizer ensures locally smooth geometry.

**Skinning Regularization.** We found that SkinNet can overfit the training poses and respective deformations by predicting non-local skinning weights, e.g., the foot joint has a non-zero skinning weight for vertices on the shoulder. This, however, harms generalization to novel poses at test time. To prevent this, we regularize the skinning weights by constraining their values based on the geodesic distance  $d_{\text{geo}}$  between the vertex and the initial body part it belongs to. This can be formalized with the following loss

$$\mathcal{L}_{\text{skin}} = \sum_{i=1}^N \sum_{j=1}^J \mathbf{w}_{i,j} \left( \frac{\min_{k \in \mathcal{A}_j} d_{\text{geo}}(\bar{\mathbf{v}}_i, \bar{\mathbf{v}}_k)}{d_{\text{geomax}}} \right)^r, \quad (7)$$

$$\mathcal{A}_j = \{k | j = \underset{t}{\operatorname{argmax}} \mathbf{w}_{\text{init},k,t}\}, \quad (8)$$

where  $\mathcal{A}_j$  is the set of vertices, which have the highest skinning weight with respect to joint  $j$ .  $r = 3$  is a hyperparameter which controls how penalization increases as the geodesic distance increases.  $d_{\text{geomax}}$  is the maximum geodesic distance between any of the vertex pairs.

**Part-wise Regularization.** In contrast to the clothed areas, the skin part has rather static skinning. Thus, we use the initial skinning to regularize the predicted skinning weights of the skin part, i.e., the set  $\mathcal{G}$  of vertices, which have the parsing label *skin*. However, even for skin vertices, we found that the initial skinning weights around the joints are not accurate while the rigid parts (i.e., the vertices away from joints) usually have correct initial skinning. We identify the set of rigid vertices as  $\mathcal{R} = \{k | \max(\mathbf{w}_{\text{init},k}) > u\}$  where  $u = 0.95$ . Thus, our part loss

$$\mathcal{L}_{\text{part}} = \sum_{i \in \mathcal{R} \cap \mathcal{G}} \|f_{\omega}(\bar{\mathbf{v}}_i, \tilde{\theta}) - \mathbf{w}_{\text{init},i}\|_2^2, \quad (9)$$

enforces that the predicted skinning weights are close to the initial ones *iff* the human parsing label of a vertex is the skin ( $\mathcal{G}$ ) and the vertex belongs to the rigid part ( $\mathcal{R}$ ).

**Training Strategy.** Our training proceeds in four stages. First, we train SkinNet *without* using the rendering loss  $\mathcal{L}_{\text{rend}}$ . This step is required to ensure that the posed mesh roughly overlaps with the foreground masks. Next, the AlbedoNet is trained using only the rendering loss  $\mathcal{L}_{\text{rend}}$  while the SkinNet weights  $\omega$  are fixed. Importantly, the AlbedoNet is *not* pose-dependent, but we optimize the weights  $\omega'$  across all frames. Afterward, the AlbedoNet and the SkinNet are fixed and we train the ShadowNet using the rendering loss  $\mathcal{L}_{\text{rend}}$ . Again we train on all training frames. Lastly, since we now have a pose- and view-dependent appearance model available, we refine the SkinNet weights including the rendering loss. We found this gives the best performance, as validated by our results.

## 4. Results

**Dataset.** We evaluate our method on 3 subjects wearing tight and loose clothing, i.e., pants or skirts. Two of them are from the DynaCap dataset [15] and they are referred as *D2* and *D5*. *D2* is wearing short pants and a T-shirt while *D5* is in a skirt. The total numbers of cameras are 94 and 101 for those two sequences and our method can be trained using as few as 10 cameras. There are around 19000 training frames per subject depicting the actor in various articulated poses. For evaluation, there is a separate testing set of around 7000 frames per subject. For convenience, we keep 1 out of 10 frames for training and evaluation. Additionally, we captured a new subject, *V6*, wearing a T-shirt and trousers using 116 cameras. We captured a separate training (17730 frames) and testing sequence (5000 frames). Importantly, for this sequence, we recover hand gestures and a respective 3D model including hands.

**Metrics.** For evaluation, we reconstruct the dense 3D geometry for the testing frames using multi-view stereo reconstruction [2]. We do not use NeuS here as it is very time-consuming. Those dense meshes serve as ground truth scans and we compare the posed meshes recovered by our method as well as competing methods in terms of the Chamfer distance and the Hausdorff distance going from the recovered mesh to the ground truth scan (M2S) and vice versa (S2M). We report the average results over all testing frames. All metrics are reported in centimeter scale.

**Competing Methods.** We compare our work to our reimplementation of Pinocchio [7], which proposed a heat diffusion process to obtain a static set of skinning weights. Moreover, we compare to SCANimate [49] and SNARF [8], which, in stark contrast to our work, require dense point clouds for training. They learn static skinning weights and a pose-aware implicit shape. Since our approach mostly focuses on skinning, we compare to SCANimate with and without pose-aware shape, referred to as SCANimate and SCANimate\* in the following. Finally, we compare to RigNet [59], which can jointly predict a rigging skeleton and static skinning weights given a template. Unfortunately, SkinningNet [41] and NeuroSkinning [34] do not provide their code and, thus, we cannot compare to them. For more details and comparisons to other (slightly less) related works [31, 56], we refer to the supplemental document.

### 4.1. Qualitative Results

In Fig. 3, we qualitatively evaluate our approach. In the first column, we show the reconstructed static geometry for different subjects. In the following columns, we show the characters in novel poses from the test set overlaid onto the respective images. Note that our posed character nicely matches the reference image, which confirms that our pose-dependent skinning indeed generalizes to novel poses and results in accurately posed meshes. In Fig. 4, we further visualize the effect of the pose-dependent skinning weights obtained by our approach. The initial skinning weights are fixed across poses leading to the typical skinning artifacts and wrong deformation behavior. In contrast, our proposed pose-dependent skinning weights adapt according to the pose and as a result the deformations look more natural.

### 4.2. Comparison

In Tab. 1, we compare our results to previous works. Note that we achieve the most accurate results when compared to works only focusing on character skinning. This validates the necessity of pose-dependent skinning weights since other works rely on static ones. In this table, SCANimate and SNARF predict pose-aware shapes while assuming dense point clouds are given. Since it is not the scope of this work to predict pose-dependent non-rigid deformations and our method solely learns from video data, we placed

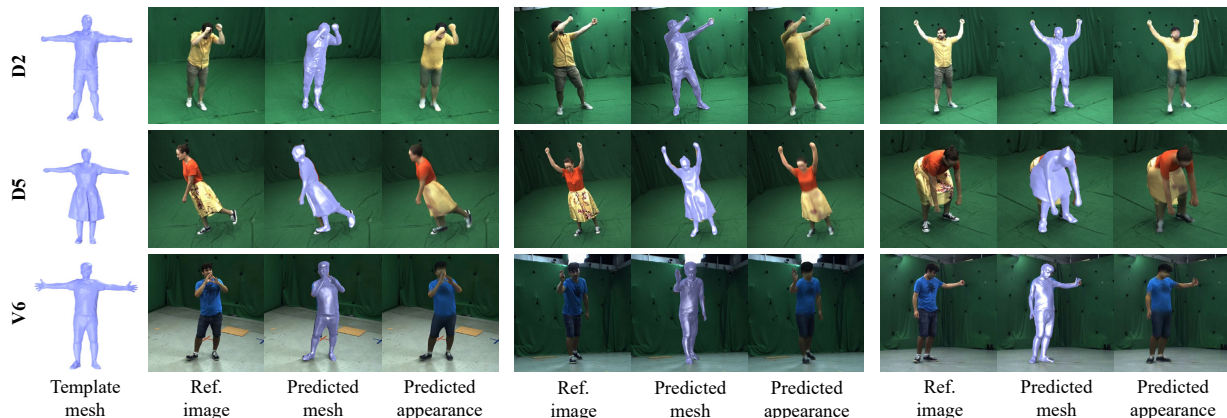


Figure 3. Qualitative results. From left to right: Recovered static character template. Reference image showing a test pose. Our posed character overlaid onto the reference image. The precise overlay of the posed geometry confirms the effectiveness of our approach.

Subject	D2			D5			V6		
Method	Chamfer↓	M2S↓	S2M↓	Chamfer↓	M2S↓	S2M↓	Chamfer↓	M2S↓	S2M↓
Pinnocchio [7]	3.760	2.162	1.599	5.077	2.811	2.267	3.358	1.871	1.487
SCANimate* [49]	3.750	2.099	1.650	5.453	2.965	2.488	3.502	1.890	1.612
RigNet [59]	3.599	2.078	1.521	4.989	2.836	2.153	3.369	1.894	1.475
<b>Ours</b>	<b>3.034</b>	<b>1.746</b>	<b>1.288</b>	<b>4.512</b>	<b>2.442</b>	<b>2.070</b>	<b>2.993</b>	<b>1.719</b>	<b>1.274</b>
SCANimate [49]	2.842	1.646	1.196	4.982	2.284	2.698	3.154	1.714	1.440
SNARF [8]	2.591	1.452	1.139	4.320	2.288	2.032	2.760	1.496	1.263

Table 1. We compare our method to previous works on 3 subjects. Our method obtains better results than other skinning-based approaches. It is even close to SCANimate and SNARF, which require dense pointclouds while we solely use multi-view video. Importantly, they focus on learning pose-aware shapes while our goal is to obtain a fully rigged and skinned character in a completely automated fashion.




Figure 4. Qualitative results showing the pose-dependent skinning weights. As a reference, we show the result using the initial static skinning weights. Notably, using static skinning weights leads to artifacts and worse deformation. When using our pose-dependent skinning, one can see a clear improvement in the final result and we also show how the weights change over different poses.

them separately. Nonetheless, it is worth noting that our approach is close to SCANimate and SNARF in terms of accuracy and better than SCANimate for subject *D5* and *V6* even though our training does not require point clouds.

We also visually compare to those works in Fig. 5. We can see qualitatively our method achieves a lower error especially in joint regions where most of the skinning artifacts

Method	Chamfer↓	M2S↓	S2M↓
Initial weights	3.761	2.162	1.599
Static weights	3.354	1.935	1.418
w/o $\mathcal{L}_{\text{rend}}$	3.116	1.783	1.333
w/ albedo-only	3.047	1.750	1.297
w/ NeuS appearance	3.152	1.786	1.365
w/ NeuS albedo	3.137	1.795	1.342
w/ displacements	3.191	1.805	1.386
Ours (10 views)	3.129	1.791	1.337
<b>Ours</b>	<b>3.034</b>	<b>1.746</b>	<b>1.288</b>

Table 2. We compare our method against baselines. All designs improve the accuracy confirming the effectiveness of VINECS.

become visible. This can be nicely seen in the error maps, which visualize the per-vertex error. Although SCANimate and SNARF perform well on naked body and tight clothes regions, they are worse than ours around loose clothing, as they rely on SMPL while our method does not have such a restriction. Further, the posed meshes clearly show fewer artifacts for our approach compared to others. Note that the entire character including the geometry, rigged skeleton, and the pose-dependent skinning can be obtained solely from multi-view video without any manual user input.

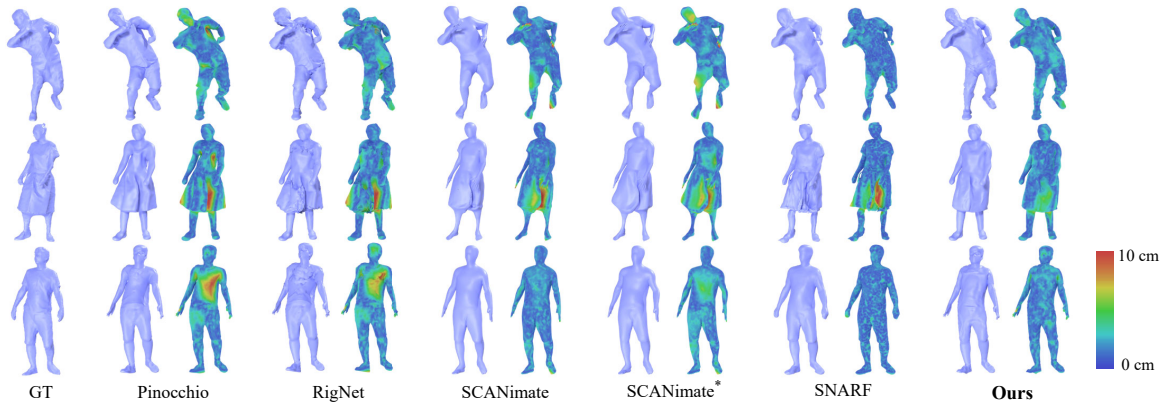


Figure 5. Qualitative comparison. For each method, we visualize the recovered posed geometry as well as the per-vertex error map when comparing the ground truth in terms of M2S. Note that our method consistently shows the lowest error and also has the least visual artifacts.




Figure 6. We compare the initial skinning, the static skinning optimized using our supervision scheme, ablating the rendering loss, and our pose-dependent skinning. We outperform all baselines.

### 4.3. Ablation

Next, we evaluate our design choices in Tab. 2. First, we evaluate the effect of pose-dependent skinning weights. Therefore, we compare our final result to the initial skinning, which we obtain by leveraging Pinocchio [7] and a version of our method where we use our supervision scheme, but optimize a static set of skinning weights, i.e., the weights are fixed across all poses. From the results, we can see that pose-dependent weights significantly improve the accuracy over the two baselines proofing the effectiveness of our formulation. This is also visualized in Fig. 6.

Another key aspect when learning those weights solely from video data is our proposed supervision scheme. When we do not employ the dense rendering loss or solely use the albedo without shadow estimation, the performance drops. Moreover, we also evaluate the accuracy when using a dense rendering loss with the static appearance from NeuS and using NeuS appearance as albedo while the shadow network is learned. We found that our proposed rendering scheme performs best. To validate the effectiveness of training pose-dependent skinning, we replace the SkinNet by a network predicting per-vertex displacement, which generalizes worse than SkinNet. We train our model with sparse cameras (10 views), and the performance is still satisfying.

## 5. Discussion and Conclusion

**Limitations.** Although we propose the first method for learning pose-dependent skinning solely from multi-view video, our method still has some limitations, which opens up interesting directions for future work. Currently, we query the SkinNet for every vertex in canonical space, which may be inefficient when many points are sampled. Thus, future work could involve exploring more efficient architectures like hashgrids [43]. While hand and body can be faithfully animated, there is no control over the facial expression. Thus, extending our automated character creation pipeline with some parametric model for facial expressions might be a promising direction. Also, a joint consideration of rigging, skinning, and pose tracking may further improve the quality of the 3D animation-ready characters. Lastly, some artifacts still remain since skinning alone cannot account for high-frequency deformations. Combining VINECS with downstream tasks such as human surface tracking/modeling could alleviate such artifacts.

**Conclusion.** We proposed the first method for creating a fully rigged and skinned character solely from multi-view video without any manual processing. To this end, we first acquire a canonical template using an implicit surface reconstruction, automated initial skinning and mesh parsing. Most importantly, we learn a pose-dependent skinning field from a multi-view video, which greatly reduces artifacts. We show how such a skinning field can be learned in a weakly supervised manner by an appearance model and differentiable rendering. Our work demonstrates that animatable character creation can be fully automated while maintaining high-quality geometry, rigging, and skinning.

**Acknowledgements.** We thank Heming Zhu and Diogo Luvizon for their help. This project was also supported by Saarbrücken Research Center for Visual Computing, Interaction, and AI. Christian Theobalt was supported by ERC Consolidator Grant 4DRely (770784).



## References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. **2**
- [2] Agisoft. PhotoScan. <http://www.agisoft.com>, 2016. **6, 1, 3**
- [3] Brett Allen, Brian Curless, and Zoran Popović. Articulated body deformation from range scan data. *ACM Trans. Graph.*, 21(3):612–619, 2002. **1, 2**
- [4] Brett Allen, Brian Curless, Zoran Popovic, and Aaron Hertzmann. Learning a Correlated Model of Identity and Pose-Dependent Body Shape Variation for Real-Time Synthesis. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation*. The Eurographics Association, 2006. **3**
- [5] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape Completion and Animation of People. *ACM Transactions on Graphics*, 24(3):408–416, 2005. **3**
- [6] Seungbae Bang and Sung-Hee Lee. Spline interface for intuitive skinning weight editing. *ACM Trans. Graph.*, 37(5), 2018. **2**
- [7] Ilya Baran and Jovan Popović. Automatic rigging and animation of 3d characters. *ACM Trans. Graph.*, 26(3), 2007. **2, 4, 6, 7, 8, 3**
- [8] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *International Conference on Computer Vision (ICCV)*, 2021. **3, 4, 6, 7**
- [9] Olivier Dionne and Martin de Lasa. Geodesic voxel binding for production character meshes. In *Proceedings of the 12th ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 173–180, New York, NY, USA, 2013. Association for Computing Machinery. **2**
- [10] Olivier Dionne and Martin de Lasa. Geodesic binding for degenerate character geometry using sparse voxelization. *IEEE Transactions on Visualization and Computer Graphics*, 20(10):1367–1378, 2014. **2**
- [11] Sven Forstmann, Jun Ohya, Artus Krohn-Grimberghe, and Ryan McDougall. Deformation styles for spline-based skeletal animation. In *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, page 141–150, Goslar, DEU, 2007. Eurographics Association. **2**
- [12] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *Proceedings of Machine Learning and Systems 2020*, pages 3569–3579. 2020. **1**
- [13] Marc Habermann, Weipeng Xu, Michael Zollhöfer, Gerard Pons-Moll, and Christian Theobalt. Livecap: Real-time human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 38(2):14:1–14:17, 2019. **2**
- [14] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR)*, 1:1, 2020. **2**
- [15] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Trans. Graph.*, 40(4), 2021. **2, 5, 6, 1**
- [16] Marc Habermann, Lingjie Liu, Weipeng Xu, Gerard Pons-Moll, Michael Zollhoefer, and Christian Theobalt. Hdhumans: A hybrid approach for high-fidelity digital humans. *Proc. ACM Comput. Graph. Interact. Tech.*, 6(3), 2023. **2**
- [17] Dae-Eun Hyun, Seung-Hyun Yoon, Jung-Woo Chang, Joon-Kyung Seong, Myung-Soo Kim, and Bert Jüttler. Sweep-based human deformation. *The Visual Computer*, 21:542–550, 2005. **2**
- [18] Alec Jacobson, Ilya Baran, Jovan Popović, and Olga Sorkine-Hornung. Bounded biharmonic weights for real-time deformation. *Commun. ACM*, 57(4):99–106, 2014. **2**
- [19] Doug L. James and Christopher D. Twigg. Skinning mesh animations. *ACM Trans. Graph.*, 24(3):399–407, 2005. **2**
- [20] Yue Jiang, Marc Habermann, Vladislav Golyanik, and Christian Theobalt. Hifecap: Monocular high-fidelity and expressive capture of human performances. In *BMVC*, 2022. **2**
- [21] Ladislav Kavan and Olga Sorkine. Elasticity-inspired deformers for character articulation. *ACM Trans. Graph.*, 31(6), 2012. **2**
- [22] Ladislav Kavan and Jiří Žára. Spherical blend skinning: A real-time deformation of articulated models. In *Proceedings of the 2005 Symposium on Interactive 3D Graphics and Games*, page 9–16, New York, NY, USA, 2005. Association for Computing Machinery. **2**
- [23] Ladislav Kavan, Steven Collins, Jiří Žára, and Carol O’Sullivan. Skinning with dual quaternions. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games*, pages 39–46. ACM, 2007. **2**
- [24] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2014. **3**
- [25] Youngjoong Kwon, Lingjie Liu, Henry Fuchs, Marc Habermann, and Christian Theobalt. Deliffas: Deformable light fields for fast avatar synthesis. *Advances in Neural Information Processing Systems*, 2023. **2**
- [26] Binh Huy Le and Zhigang Deng. Robust and accurate skeletal rigging from mesh sequences. *ACM Trans. Graph.*, 33(4), 2014. **2**
- [27] Binh Huy Le and Jessica K. Hodgins. Real-time skeletal skinning with optimized centers of rotation. *ACM Trans. Graph.*, 35(4), 2016. **2**

- [28] J. P. Lewis, Matt Corder, and Nickson Fong. 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*, page 165–172, USA, 2000. ACM Press/Addison-Wesley Publishing Co. **2, 4**
- [29] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. **4, 1**
- [30] Peizhuo Li, Kfir Aberman, Rana Hanocka, Libin Liu, Olga Sorkine-Hornung, and Baoquan Chen. Learning skeletal articulations with neural blend shapes. *ACM Transactions on Graphics (TOG)*, 40(4):1, 2021. **1, 3**
- [31] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. 2022. **3, 6, 1, 2**
- [32] Yue Li, Marc Habermann, Bernhard Thomaszewski, Stelian Coros, Thabo Beeler, and Christian Theobalt. Deep physics-aware inference of cloth deformation for monocular human performance capture. 2020. **2**
- [33] Zhouyingcheng Liao, Jimei Yang, Jun Saito, Gerard Pons-Moll, and Yang Zhou. Skeleton-free pose transfer for stylized 3d characters. In *European Conference on Computer Vision (ECCV)*. Springer, 2022. **1, 2**
- [34] 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. Graph.*, 38(4), 2019. **1, 2, 6**
- [35] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Trans. Graph.*, 40(6), 2021. **2**
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34(6):248:1–248:16, 2015. **2, 3**
- [37] William Lorensen and Harvey Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM SIGGRAPH Computer Graphics*, 21:163–, 1987. **4, 3**
- [38] N. Magnenat-Thalmann, A. Laperrière, and D. Thalmann. Joint-dependent local deformations for hand animation and object grasping. In *Proceedings of Graphics Interface '88*, pages 26–33. Canadian Man-Computer Communications Society, 1988. **2**
- [39] 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 *Computer Vision – ECCV 2020*, pages 405–421, Cham, 2020. Springer International Publishing. **1**
- [40] Alex Mohr and Michael Gleicher. Building efficient, accurate character skins from examples. *ACM Trans. Graph.*, 22(3):562–568, 2003. **1, 2**
- [41] A. Mosella-Montoro and J. Ruiz-Hidalgo. Skinningnet: Two-stream graph convolutional neural network for skinning prediction of synthetic characters. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18572–18581, Los Alamitos, CA, USA, 2022. IEEE Computer Society. **1, 2, 6**
- [42] Tomohiko Mukai and Shigeru Kuriyama. Efficient dynamic skinning with low-rank helper bone controllers. *ACM Trans. Graph.*, 35(4), 2016. **2**
- [43] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. **8**
- [44] Haokai Pang, Heming Zhu, Adam Kortylewski, Christian Theobalt, and Marc Habermann. Ash: Animatable gaussian splats for efficient and photoreal human rendering. 2023. **2**
- [45] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. **3**
- [46] Sang Il Park and Jessica K. Hodgins. Capturing and animating skin deformation in human motion. *ACM Trans. Graph.*, 25(3):881–889, 2006. **2**
- [47] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Animatable neural radiance fields for human body modeling. *ICCV*, 2021. **2**
- [48] 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. *CVPR*, 1(1): 9054–9063, 2021. **2**
- [49] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. SCANimate: Weakly supervised learning of skinned clothed avatar networks. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. **3, 6, 7**
- [50] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. *Advances in neural information processing systems*, 29, 2016. **1**
- [51] Soumyadip Sengupta, Vivek Jayaram, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Background matting: The world is your green screen. In *Computer Vision and Pattern Recognition (CVPR)*, 2020. **3**
- [52] Peter-Pike J. Sloan, Charles F. Rose, and Michael F. Cohen. Shape by example. In *Proceedings of the 2001 Symposium on Interactive 3D Graphics*, pages 135–143, New York, NY, USA, 2001. Association for Computing Machinery. **1, 3**
- [53] TheCapturey. The Capturey. <http://www.thecapturey.com/>, 2020. **1, 3, 4**
- [54] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algo-

- rithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. [3](#)
- [55] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. [1](#), [3](#), [4](#)
- [56] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdfs. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXII*, page 1–19, Berlin, Heidelberg, 2022. Springer-Verlag. [3](#), [6](#), [1](#), [2](#)
- [57] 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*, page 129–138, New York, NY, USA, 2002. Association for Computing Machinery. [1](#), [2](#)
- [58] Rich Wareham and Joan Lasenby. Bone glow: An improved method for the assignment of weights for mesh deformation. In *Articulated Motion and Deformable Objects*, pages 63–71, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg. [2](#)
- [59] Zhan Xu, Yang Zhou, Evangelos Kalogerakis, Chris Landreth, and Karan Singh. Rignet: Neural rigging for articulated characters. *ACM Trans. on Graphics*, 39, 2020. [1](#), [2](#), [6](#), [7](#), [3](#)
- [60] Ji Yang, Xinxin Zuo, Sen Wang, Zhenbo Yu, Xingyu Li, Bingbing Ni, Minglun Gong, and Li Cheng. Object wake-up: 3d object rigging from a single image. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 311–327. Springer, 2022. [2](#)
- [61] Xiaosong Yang, Arun Somasekharan, and Jian Jun Zhang. Curve skeleton skinning for human and creature characters. *Computer Animation and Virtual Worlds*, 17, 2006. [2](#)