MeshAvatar: Learning High-quality Triangular Human Avatars from Multi-view Videos

Yushuo Chen¹^o, Zerong Zheng²^o, Zhe Li¹^o, Chao Xu²^o, and Yebin Liu¹^o

¹ Tsinghua University, Beijing, China
² NNKosmos Technology, Hangzhou, China

Abstract. We present a novel pipeline for learning high-quality triangular human avatars from multi-view videos. Recent methods for avatar learning are typically based on neural radiance fields (NeRF), which is not compatible with traditional graphics pipeline and poses great challenges for operations like editing or synthesizing under different environments. To overcome these limitations, our method represents the avatar with an explicit triangular mesh extracted from an implicit SDF field, complemented by an implicit material field conditioned on given poses. Leveraging this triangular avatar representation, we incorporate physics-based rendering to accurately decompose geometry and texture. To enhance both the geometric and appearance details, we further employ a 2D UNet as the network backbone and introduce pseudo normal ground-truth as additional supervision. Experiments show that our method can learn triangular avatars with high-quality geometry reconstruction and plausible material decomposition, inherently supporting editing, manipulation or relighting operations. The code is available at https://github.com/shadOwta9/meshavatar.

Keywords: full-body avatars \cdot relighting \cdot physics-based rendering

1 Introduction

Human avatar creation plays a crucial role in the animation and film industries, enabling the efficient creation of character animations by artists. Unfortunately, constructing a high-fidelity avatar for a specific character is a costly and laborious process [14,23]. Conventional graphics pipelines typically involve scanning, texturing, rigging, and skinning, each requiring substantial human efforts.

Benefiting from the huge progress of learning-based 3D modeling, researchers have devoted great efforts in automatically learning 3D human body avatars from images or videos of real humans. With the rise of neural radiance fields (NeRF) [63], recent works tend to represent the 3D character as a pose-conditioned NeRF, and enhance the avatar representation with forward/backward skinning [10, 11,15,49,51,62], splitting and structuring [103,104] or UV mapping [24,57]. In addition, concurrent works [30,43,50,66,105] introduce 3D Gaussian splatting [41], a discrete radiance field representation, into human avatar modeling to realize better fidelity and fast rendering. Although these methods can produce realistic rendering of human avatars, they typically entangle geometry, appearance



Fig. 1: Example results by our method. Given the multi-view videos of a specific subject, our method learns his triangular avatar with the geometry reconstruction and intrinsic material decomposition. After training, the avatar not only supports novel pose synthesis and relighting, but also enables texture editing and material manipulation. In this example, we make the arm metallic and edit the texture of the T-shirt.

and lighting into neural networks, neglecting the recovery of explicit geoemtric surface and limiting many applications in traditional graphics pipelines such as editing and relighting. Although several works attempted to realize inverse rendering for humans in NeRF, they either rely on an inaccurate, physics-agnostic network prediction [13,55], or require computationally intensive sphere ray tracing [95].

To overcome these limitations, in this paper, we deviate from the current trend and aim to learn human avatars in the form of triangular meshes, which is more compatible with traditional graphics engines and inherently enables physically-based rendering, relighting and editing. Although early works has explored mesh-based avatar modeling [3, 25, 44, 92, 93], they require per-frame mesh reconstruction [3, 92, 93] or pre-scanned templates [25, 44]. As a result, these avatars cannot be learned in an end-to-end manner and none of them support material decomposition.

Therefore, we propose a novel hybrid representation for triangular human avatar modeling, enabling complete end-to-end training from multi-view videos. Specifically, the learned model consists of an *explicit* triangular mesh geometry (with vertex skinning weights) extracted from an *implicit* SDF field, and an *implicit* dynamic texture and material field conditioned on given poses. The explicit mesh and implicit fields are bridged by differentiable marching tetrahedra (DMTet) [76]. On one hand, mesh-based human geometry can exploit the compatibility with conventional graphics pipelines for editing operations. Moreover, hardware-accelerated ray tracing algorithm could be applied for more realistic rendering and better geometry recovery and intrinsic material decomposition [26]. On the other hand, the implicit representation delivers higher representation flexibility while inherently regularizing spatial smoothness on both the geometry and the material, which enables effective, end-to-end learning of triangular avatars without any initial guess for the 3D geometry. After training, the learned avatar is capable of both novel pose synthesis and relighting/editing, as demonstrated in Fig. 1.

3

Nevertheless, acquiring such a representation directly from image observations poses challenges. Dynamic surface reconstruction from sparse views is inherently hindered by the shape-radiance ambiguity, and ensuring model representativeness for dynamic geometric details further complicates the learning process. Here we present several technical designs to address these challenges. Specifically, we incorporate shadow-aware Physics-based Rendering (PBR) into the differentiable rasterization process [26]. PBR accurately models the lighting process and shadow effects by analytically connecting the outgoing radiance with the underlying surface position. Compared to implicit radiance fields, PBR can prevent shadow baking, facilitating accurate surface reconstruction. To fully harness the benefits of PBR, a geometry reconstruction module capable of recovering high-quality surface details is essential. Drawing inspiration from [50], we employ powerful 2D neural networks to encode pose features and model high-frequency geometry information. To further enhance geometric details and reduce ambiguities, we leverage normal map prior estimated from image stereo as a weak supervision. Note that this is feasible owing to our triangular mesh representation that allows direct supervision on surface details, which is always difficult in implicit surface representation. Benefiting from these technical designs, our method achieves high-quality reconstruction of the shape of dynamic humans, without reliance on explicit surface tracking.

As shown in the experiments, our method can learn a triangular avatar with high-quality geometry and plausible materials in an end-to-end manner. Moreover, we achieve state-of-the-art dynamic human reconstruction in terms of both geometry and appearance. In summary, our technical contributions are:

- We propose an end-to-end pipeline for learning triangular human avatars from multi-view videos. The learned avatars are represented as triangular meshes, which is compatible with traditional graphics pipeline for editing or physics-based rendering.
- We propose a hybrid human avatar representation for triangular human avatar modeling. It bridges explicit triangular mesh geometry with implicit SDF/material fields, and recovers high-quality details by integrating 2D neural networks.
- We achieve high-quality geometry recovery by PBR differentiable rendering, which can prevent shadow baking and facilitate both surface reconstruction and material decomposition.
- We propose a stereo algorithm based on human prior for normal estimation. The estimated normal maps can be used as a weak supervision to further improve the geometry quality.

2 Related Work

2.1 Human Avatars using Explicit Representations

Over the last decade, efforts have been made to achieve drivable avatar modeling. Traditional pipelines typically employ explicit representations to model the

4 Y. Chen et al.

shapes of clothed human bodies. Among them, polygon meshes is one of the most popular representations. For example, early approaches first reconstruct a character-specific mesh template and animate it through simulation [21, 78] or retrieval [94]. Some works directly deform a generic human body template (e.g. SMPL [58]) to approximate the shape of clothed humans [2, 9, 42, 101]. Unfortunately, the quality of avatars is also constrained by the limitations of their deformed SMPL representation. In the deep learning era, recent works, leveraging convolutional networks as a powerful learning tool, propose to learn non-rigid deformation and dynamic appearance in the UV space of subject-specific mesh templates [3, 25, 44, 92, 93]. These methods usually require pre-scanning a subject-specific template or sophisticated non-rigid tracking throughout the sequence. In contrast, our hybrid representation eliminates these requirements and autonomously learns an triangular mesh template in an end-to-end fashion.

Apart from mesh-based representation, point clouds have gained popularity for their flexibility. These approaches mainly focus on modeling the human geometry [54,59,60,97,102]. Recently, 3D Gaussian splatting (3DGS) [41], an efficient differentiable point-based rendering method, has demonstrated real-time photo-realistic rendering. By incorporating 3DGS into avatar modeling, concurrent works realize high-fidelity human animation results [30, 43, 50, 53, 66, 105]. However, the unstructured nature of points poses significant challenges for recovering the underlying geometry and/or its intrinsic material properties of human shapes. In contrast, our method is able to recover the geometric shapes as well as intrinsic properties, enabling physics-based rendering and avatar relighting.

2.2 Human Avatars using Implicit Representations

Implicit representations model scene with coordinate-based MLPs that map continuous spatial coordinates to an implicit field, such as signed distance function (SDF) [67], occupancy [61], and radiance fields (NeRF) [63]. Since their debut, it becomes a popular trend to adopt implicit representations in avatar modeling framework. In the realm of geometric avatar modeling, numerous approaches leverages pose-conditioned SDF field [18, 28, 73, 85, 87] or occupancy fields [10,11,15,51,62] learned from human scans or depth sequences. To capture the avatar appearance, recent works leverage NeRF as the underlying representation [16, 19, 22, 36, 37, 48, 49, 68, 70, 75, 80, 84, 88, 91, 103, 104] for its impressive learning capability. Some research explores hybrid representations that combine the strengths of both explicit and implicit representations [12, 57, 71, 74, 79]. Despite the rapid progress, these implicit function-based methods necessitates dense sampling and extensive network evaluation along pixel rays, suffering from high computational complexity and making physics-based ray tracing extremely difficult. In contrast, our representation directly generates a polygon mesh for image rendering, which allows us to seamlessly integrate physics-based ray tracing methods. This enables high-fidelity material decomposition and reconstruction while mitigating the computational burden associated with ray tracing in implicit representations.

2.3 Neural Inverse Rendering

Inverse rendering aims to disentangle intrinsic scene properties including geometry, material and lighting conditions from image observations. This is a fundamental yet extremely ill-posed task. Therefore, previous methods [4,5,17,20, 33,46,47,52,65] often resorted to assumptions like known geometry, controlled lighting or other physics priors. Recently, advances in neural implicit representations aid in inverse rendering tasks by simultaneously recovering scene properties without explicit geometry guidance [6,7,38,77,98,100]. Despite achieving impressive results, the ambiguity in volume rendering lead to low geometric quality. Surface-based inverse rendering methods, on the contrary, directly optimize the underlying geometry with differentiable rasterization, showing higher accuracy in surface reconstruction [26, 64].

This progress have also inspired follow-up research on human performance relighting from sparse multi-view videos. A typical pipeline decomposes human performance into canonical reflectance fields and incorporate human templates for motion modeling [13, 32, 81]. However, these methods often struggle with inaccurate surface reconstruction due to erroneous motion estimation and lack of geometry supervision. The compromised quality of dynamic geometry hampers the modeling of material properties, resulting in unsatisfactory relighting effects. In contrast, our approach excels in high-fidelity surface reconstruction and material decomposition for complex human avatars, overcoming the limitations of existing methods and delivering superior results. Moreover, our triangular representation are more suitable in conventional graphics engines and applications. Concurrent to us, Wang et al. [86], Xu et al. [95] and Lin et al. [55] improve the previous pipelines of relightable avatar learning by introducing explicit ray tracing [86], hierarchical distance query [95] or part-wise light visibility estimation [55]. These techniques allow for better modeling of the lighting process. However, both of them are still built upon implicit representation, making them computationally expensive when synthesizing novel poses and views.

Some methods have been proposed to enable relighting of human portraits from low-cost inputs such as single images [34,35,40,45,82] or videos [89]. Despite their convenience, these methods cannot change the viewpoints or the body poses, so a deep discussion of these works is out of the scope of this paper.

3 Method

3.1 Overview

Given a collection of multiview videos capturing a clothed human' motions and the corresponding SMPL-X registrations, our task is to build a photo-realistic avatar for this subject in an end-to-end manner, *in the form of an explicit triangular mesh*. Such a goal deviates notably from the recently popular trend of NeRF-based avatar modeling, and this departure is particularly significant as it addresses the challenges associated with intrinsic decomposition in human avatar modeling. However, this task is by no means trivial and poses two critical

5



Fig. 2: Illustration of our method. Our pipeline learns a hybrid human avatar represented in the form of (a) an explicit skinned mesh and (b) implicit pose-dependent material fields. Such a representation inherently supports (c) physics-based ray tracing and can be trained in an end-to-end manner using (d) normal estimation as an additional supervision signal.

difficulties. One is to construct an accurate mesh template with pose-dependent fine-grained details for clothed human animation, and the other is to address the shape-radiance ambiguity in surface reconstruction. To tackle these difficulties, we propose a novel pipeline utilizing a hybrid representation to model triangular clothed human avatars.

As shown in Fig. 2, our pipeline learns a hybrid human avatar representation, which consists of an explicit skinned mesh extracted from an implicit SDF field using deep marching tetrahedra (Sec. 3.2), and a pose-dependent detail generation network that produces dynamic vertex offsets (Sec. 3.3). The intrinsic material properties including albedo and roughness are predicted in the form of volumetric fields, with which our avatar can be rendered using physics-based rendering given the learnable environmental lighting (Sec. 3.4). All networks are trained in an end-to-end manner supervised by the input images and normal estimations (Sec. 3.5).

3.2 Canonical Human Skinned Mesh

In this paper, we propose a hybrid representation to model the canonical implicit SDF field of the clothed human. Specifically, we represent the SDF field by a coordinate-based MLP

$$\mathcal{S}(x): \mathbb{R}^3 \mapsto \mathbb{R}.\tag{1}$$

in order to regularize the topology smoothness [67]. Instead of performing volume rendering from the SDF field, we leverage marching tetrahedra (MT) [76] algorithm to extract the mesh in the canonical space for the following processing. This algorithm differentiably connects the explicit triangular mesh and the implicit SDF field in the canonical space, making it possible to optimize for a variety of shapes in an end-to-end manner.

To associate the canonical space and the observation (posed) space, a skinning weight field is required for skinning the canonical mesh extracted by DMTet. Previous works project the spatial points onto the fitted SMPL surface to obtain an initial guess, and learn some residuals for refinements [69]. However, simultaneously solving the geometry and its skinning weights is highly ill-posed. Therefore, we follow FITE [54] to spread the skinning weights $\mathcal{W}^{\mathcal{T}}$ on SMPL surface \mathcal{T} into 3D space by minimizing the energy:

$$\mathcal{E} = \lambda_p \int_{p \in \mathcal{T}} \|\nabla \mathcal{W}(p) - \nabla_{\mathcal{T}} \mathcal{W}^{\mathcal{T}}(p)\|^2 + \lambda_{\mathcal{T}} \int_{p \in \mathcal{T}} \|\mathcal{W}(p) - \mathcal{W}^{\mathcal{T}}(p)\|^2 + \lambda_r \int_{p \in \mathbb{R}^3} \|\nabla^2 \mathcal{W}(p)\|^2,$$
(2)

where $\mathcal{W} : \mathbb{R}^3 \to \mathbb{R}^J$ denotes the diffused skinning weights. After minimizing the above energy, we cache the diffused skinning weights on a $256 \times 256 \times 256$ tetrahedra grid for subsequent queries. In this way, an arbitrary mesh extracted from the grid could be instantly transformed into any posed space via the linear blend skinning (LBS) algorithm. Specifically, let's denote the vertices on the extracted canonical mesh as $\{\mathbf{c}_n\}_{n=1}^{N_c}$. Given any pose $\theta \in \mathbb{R}^{3J}$, the corresponding vertex in the posed space

$$\mathbf{v}_{n} = \sum_{i=1}^{J} \mathcal{W}(\mathbf{c}_{n})_{i} \mathbf{M}_{i}(\theta) \mathbf{c}_{n}, \qquad (3)$$

where $\mathbf{M}_{i}(\theta)$ is the rigid transformation for the *i*-th joint.

3.3 Pose-dependent Geometric Detail Generation

Although the representation in Sec. 3.2 can roughly approximate the shape under different body poses, high-frequency details like the cloth wrinkles are missing due to the low-frequency bias in MLP [83]. Without high-quality surface geometry, material decomposition and physics-based rendering are impossible. Although previous works like LaplacianFusion [42] have represented the highfrequency details with Laplacian coordinates, they require expensive computations at test time. Thanks to our triangular mesh representation, we can represent the geometric details by pure vertex offsets on the skinned mesh, and the problem boils down to accurate vertex offset estimation.

To this end, we draw inspiration from [50] and employ techniques on 2D images to generate high-frequency features. Specifically, consider the vertices on the canonical SMPL template, we assign the posed coordinates to the vertex colors, then acquire the front/back position map $\mathcal{P}_f, \mathcal{P}_b \in \mathbb{R}^{512 \times 512 \times 3}$ by orthogonally project the colored mesh onto the front/back plane. The position maps are then fed into a UNet \mathcal{U} to obtain the vertex offset maps and an auxiliary feature map, as

$$(\mathcal{V}_f, \mathcal{V}_b, \mathcal{F}) = \mathcal{U}\left(\operatorname{concat}(\mathcal{P}_f(\theta), \mathcal{P}_b(\theta))\right),\tag{4}$$

8 Y. Chen et al.

where \mathcal{V}_f and \mathcal{V}_b denotes the front and back vertex offset maps, respectively, and \mathcal{F} is the auxiliary pose feature map. Then the high-frequency vertex offsets are directly extracted from the offset map via bilinear interpolation according to the canonical coordinates:

$$\Delta \mathbf{v}_n = \mathbf{1}_{\{\mathbf{c}_{n,z}>0\}} \mathcal{V}_f(\mathbf{c}_{n,x},\mathbf{c}_{n,y}) + \mathbf{1}_{\{\mathbf{c}_{n,z}\leq0\}} \mathcal{V}_b(\mathbf{c}_{n,x},\mathbf{c}_{n,y}).$$
(5)

To enhance generalizability of this pose encoding for novel poses, we also apply position map projection onto training space by PCA [50].

3.4 Material Maps and Physics-based Rendering

Following the conventions in previous works, we use Disney material model [8] and assume that the illumination comes from infinity, thus it could be represented as an environmental map. These material parameters and lighting parameters can be optimized through differentiable rendering. Considering inverse rendering itself is a highly ill-posed problem, we manually set the metallic coefficient to zero in order to to regularize the solution complexity for clothed human bodies. Although the network in Sec. 3.3 is powerful enough to generate high-frequency geometric details, tiny errors are still inevitable due to the limited mesh resolution. To compensate for this inaccuracy, we make the material map pose-dependent. Specifically, for any point \mathbf{v} on the posed human body, we decode its material properties (albedo color, roughness) from the feature map (Eq. 4) by a shallow MLP \mathcal{M} :

$$(k_d(\mathbf{v}), k_s(\mathbf{v})) = \mathcal{M}\left(\operatorname{concat}(\gamma(\mathbf{v}), \mathcal{F}(\mathbf{c}_x, \mathbf{c}_y))\right), \tag{6}$$

where **c** is the corresponding canonical coordinate calculated via barycentric interpolation, and $\gamma(\cdot)$ is the Fourier positional encoding function [63]. To guarantee material consistency across poses, the environmental lighting is restricted to be identical for different poses.

Since the human body is diffusive-dominant and consequently low-frequency environment map is sufficient to recover the materials [26,64,98], we apply a differentiable Monte-Carlo (MC) renderer [26] for image rendering. Following the rendering equation [39], the outgoing radiance $L(\omega_o)$ is computed by the integration of the product of incident radiance $L_i(\omega_i)$ and BRDF function $f_{\text{BRDF}}(\cdot, \cdot)$ over the upper hemisphere Ω around the surface normal \mathbf{n} , which can be approximated by Monte-Carlo sampling on reflection rays:

$$L(\omega_o) = \int_{\Omega} L_i(\omega_i) f_{\text{BRDF}}(\omega_i, \omega_o) \langle \omega_i, \mathbf{n} \rangle \, \mathrm{d}\omega_i$$
$$\approx \frac{1}{N} \sum_{i=1}^N \frac{L_i(\omega_i) f_{\text{BRDF}}(\omega_i, \omega_o) \langle \omega_i, \mathbf{n} \rangle \, \mathrm{d}\omega_i}{p(\omega_i)}.$$
(7)

Note that this sampling process is particularly effective under our assumption of low-frequency lighting. We also apply Multiple Importance Sampling (MIS)



Fig. 3: Animation and relighting results. The actors on the same row are driven with the same pose, while the actors on the same column are relighted under the same global illumination.

and image denoising techniques to further improve the sample efficiency, so that the image can be rendered with only 4 to 8 samples per pixel during training [26]. Moreover, this MC-based renderer is able to account for the occlusions and generate shadows, thereby helping us reconstruct accurate and high-quality materials and create photo-realistic animations. Please refer to [26] for more implementation details about the MC-based renderer.

3.5 Prior-based Normal Supervision

To further reduce ambiguities and enhance geometric detail reconstruction, we leverage normal maps estimated from the multi-view inputs as an additional supervision signal. For this purpose, we utilize 3D human scans collected from [96] to train a normal estimation network. The network consists of two modules designed to estimate normals in a coarse-to-fine fashion. The coarse module, adapted from RAFT-Stereo [56], takes a pair of images from two neighboring views as input and produces a disparity map, as is also done in StereoPIFu [29]. The disparity map is then converted into a depth map, from which coarse normals can be analytically calculated. After that, a refinement network is employed





Fig. 4: More reconstruction results, including reconstructed geometry, albedo color, material roughness, and light probes.

to further enhance the normal details. This is accomplished by an image-toimage translation module, which refines the coarse normal map by utilizing the RGB image as a condition. Both network modules are trained with ground-truth supervision, which is rendered from the 3D scans in a similar way to PIFuHD [72]. Once trained, the normal estimation network is applicable to different datasets.

To use the normal prior in our avatar learning pipeline, we simply apply it to the multi-view image inputs to obtain the normal estimations. These estimations serve as pseudo ground-truth to construct an additional loss for supervising the avatar geometry. The detailed learning objective will be presented in the supplemental material.

Experiments 4

4.1**Experimental Settings**

Datasets. Our method is evaluated on the datasets from ActorsHQ [31] and AvatarReX [104]. There are totally 8 + 2 character sequences. For setting consistency, we manually select 16 surrounding, full-body views from 160 views in ActorsHQ. The SMPL-X registrations for these sequences can be obtained by



Fig. 5: Qualitative comparisons on training frame reconstructions. Our method could reconstruct fine-grained dynamic human geometry.



Fig. 6: Qualitative comparisons on novel pose synthesis. Our method could not only synthesize high-quality geometry, but also more realistic appearance.

existing tools [1]. For quantitative comparison, we make a train-test split for each sequence and the training set contains 1000 frames.

Main Results. The animation and relighting results are presented in Fig. 3. We also present our reconstruction results, including reconstructed geometry, albedo color, material roughness and avatar relighting, in Fig. 4. Our approach demonstrates the capability not only in synthesizing human images in novel poses but also in adapting to diverse lighting conditions. These results prove the robustness and adaptability of our method in achieving realistic and dynamic human representations with high-quality geometry and material decomposition.

4.2 Comparison

We compare our methods with recent SOTA works on neural avatars, including ARAH [88], PoseVocab [49] and AvatarReX [104]. The results of ARAH and

12 Y. Chen et al.

Table 1: Quantitative comparisons on reconstructions and novel pose synthesis. The best results and the second best are highlighted in **bold** and <u>underlined</u> fonts, respectively.

	Training Frame Reconstruction				Novel Pose Synthesis			
	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	LPIPS \downarrow	FID \downarrow	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	LPIPS \downarrow	$\mathrm{FID}\downarrow$
ARAH	23.1071	0.9390	0.1232	98.1130	22.9917	0.9368	0.1254	104.3341
PoseVocab	27.7608	0.9580	0.0668	41.5524	26.2010	0.9490	0.0752	41.3314
AvatarReX	28.4015	0.9607	0.0579	30.6214	26.5853	0.9497	0.0684	36.1911
ours (radiance)	28.9939	0.9646	0.0497	24.1061	26.6194	0.9498	0.0657	31.5325
ours (PBR)	$\underline{28.9496}$	0.9630	0.0559	$\underline{27.0561}$	26.8526	0.9505	0.0693	35.4671



Fig. 7: Qualitative comparison against Relighting4D [13] in terms of material decomposition and relighting fidelity.

PoseVocab are reproduced by the public codes and the results of AvatarReX are requested from the authors. In order to have a fair and complete evaluation, we implemented a NeRF-version for our method, denoted as *ours (radiance)*, along-side the primary version referred to as *ours (PBR)*. Specifically, they are different only in the rendering process. In *Ours(radiance)*, the surface color is directly predicted by neural networks like NeRF, instead of integrating the reflective rays on the upper hemisphere in *ours (PBR)*. For quantitative comparisons on synthesized images, we use the common Peak Signal-to-Noise Ratio (PSNR), Structure Similarity Index Measure (SSIM) [90], Learned Perceptual Image Patch Similarity (LPIPS) [99] and Fréchet Inception Distance (FID) [27] as our metrics. We additionally evaluated the reconstructed geometry on the rendered normal map. The normal error with pseudo ground truth is 12.17 (degree).

We assess the effectiveness of our approach through the evaluations on both reconstruction and novel pose synthesis. To evaluate the reconstruction quality, we present the qualitative comparisons on both the AvatarReX and ActorsHQ datasets in Fig. 5. The 'geometry' column depicts the rendered normal map pro-



Fig. 8: Ablation Studies. (a) w/o PBR and normal; (b) w/o PBR; (c) w/o normal; (d) ours; (e) reference image and estimated normal.

duced by each method, with the final column corresponding to the estimated normals derived from the raw image data. Notably, our method excels in reconstructing dynamic geometric details, such as cloth wrinkles, and capturing high-fidelity appearances. In contrast, existing approaches often rely on baking these appearance details onto a smooth human body surface. In terms of novel pose synthesis, Fig. 6 showcases the synthesized results under novel poses. Our learned human geometry demonstrates the capability of generalizing to unseen poses, and generating plausible geometric details. This capability leads to better realism in the synthesis quality achieved by our method.

Tab. 1 provides the quantitative comparisons conducted on the AvatarReX dataset. Notably, the NeRF version of our method exhibits superior performance, proving that our hybrid representation surpasses all three recent approaches in human avatar modeling. Despite a minor performance decline due to rendering constraints, the PBR version still outperforms previous methods, underscoring the efficacy of our accurate geometry modeling.

To evaluate our performance on material decomposition and avatar relighting, we conduct a comparative analysis with Relighting4D [13], one of SOTA inverse rendering pipelines designed for human motion sequences. Fig. 7 visually compares the reconstruction and relighting outcomes on one of the training frame. As shown in the figure, our method excels in learning more accurate geometry and albedo color, leading to superior results in synthesizing realistic outputs under novel lighting conditions. In contrast, Relighting4D struggles with both geometry reconstruction and material decomposition. This comparison shows the enhanced capabilities of our approach in capturing and rendering intricate details in both geometry and appearance. Due to the absence of public code, we could not compare with other previous works like [32,81]. More experiments on material decomposition and relighting are presented in the supplemental material.



Fig. 9: Limitations of our method. (a) It's challenging to factorize all the shadows due to the limited resolution of the mesh geometry, which leads to pose-dependent shadow baking. In practice, the material field is conditioned on poses. (b) For loose garments, our method failed to learn the correct geometric deformation across different poses.

4.3 Ablation Study

In this section, we evaluate the effectiveness of PBR rendering and pseudonormal supervision in the context of learning human geometry. Fig. 8 visually presents rendered images and normal maps for a representative training frame, illustrating the impact of each component. (a) Similar to other NeRF-based neural avatars, the exclusive prediction of radiance suffers from shape-radiance ambiguity. While it may achieve fair rendering quality, it often falls short in reconstructing complex geometric details. (b) Pseudo normal ground-truth supervision encourages the generation of similar geometric structures, albeit with limitations imposed by multi-view consistency. (c) The introduction of shadowaware PBR proves beneficial in capturing geometric details; see the wrinkles on the arm. However, its efficacy is constrained when tasked with recovering concave structures, particularly when provided with sparse views. (d) By leveraging the strengths of both approaches, we successfully combine the merits of the two worlds, resulting in the reconstruction of high-fidelity human geometry.

5 Conclusion

We introduce a novel framework for high-quality triangular human avatar modeling from multi-view videos. Our approach represents the avatar with an explicit triangular mesh, which is extracted from an implicit SDF field, complemented by a pose-dependent material field learned through the utilization of differentiable physics-based ray tracing and rendering. Moreover, we integrate 2D feature maps as well as normal supervision to further facilitate geometric detail learning. **Limitations.** A limitation of our work is that we use pose-dependent material to compensate the geometric errors caused by the limited resolution of the triangular mesh. This design, though practical, is not physically plausible. (Fig. 9a). In addition, more challenging garment types like loose dresses exhibit complex

non-rigid deformations, which may poses difficulties in our framework (Fig. 9b). **Acknowledgements.** This paper is supported by National Science and Technology Major Project (2021ZD0113503) and the NSFC project No.62125107.

References

- 1. https://github.com/zju3dv/EasyMocap
- Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: CVPR. pp. 8387–8397 (2018)
- Bagautdinov, T., Wu, C., Simon, T., Prada, F., Shiratori, T., Wei, S.E., Xu, W., Sheikh, Y., Saragih, J.: Driving-signal aware full-body avatars. TOG 40(4), 1–17 (2021)
- Barron, J.T., Malik, J.: Shape, albedo, and illumination from a single image of an unknown object. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 334–341. IEEE (2012)
- Bi, S., Xu, Z., Sunkavalli, K., Kriegman, D., Ramamoorthi, R.: Deep 3d capture: Geometry and reflectance from sparse multi-view images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5960–5969 (2020)
- Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12684–12694 (2021)
- Boss, M., Jampani, V., Braun, R., Liu, C., Barron, J., Lensch, H.: Neural-pil: Neural pre-integrated lighting for reflectance decomposition. Advances in Neural Information Processing Systems 34, 10691–10704 (2021)
- 8. Burley, B.: Physically based shading at disney (2012)
- Burov, A., Nießner, M., Thies, J.: Dynamic surface function networks for clothed human bodies. In: ICCV. pp. 10754–10764 (2021)
- Chen, X., Jiang, T., Song, J., Rietmann, M., Geiger, A., Black, M.J., Hilliges, O.: Fast-snarf: A fast deformer for articulated neural fields. IEEE T-PAMI (2023)
- Chen, X., Zheng, Y., Black, M.J., Hilliges, O., Geiger, A.: Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In: ICCV. pp. 11594–11604 (2021)
- Chen, Y., Wang, X., Chen, X., Zhang, Q., Li, X., Guo, Y., Wang, J., Wang, F.: Uv volumes for real-time rendering of editable free-view human performance. In: CVPR. pp. 16621–16631 (2023)
- Chen, Z., Liu, Z.: Relighting4d: Neural relightable human from videos. In: European Conference on Computer Vision. pp. 606–623. Springer (2022)
- 14. Debevec, P.: The light stages and their applications to photoreal digital actors. SIGGRAPH Asia 2(4), 1–6 (2012)
- Deng, B., Lewis, J.P., Jeruzalski, T., Pons-Moll, G., Hinton, G., Norouzi, M., Tagliasacchi, A.: Nasa neural articulated shape approximation. In: ECCV. pp. 612–628. Springer (2020)
- Dong, J., Fang, Q., Guo, Y., Peng, S., Shuai, Q., Zhou, X., Bao, H.: Totalselfscan: Learning full-body avatars from self-portrait videos of faces, hands, and bodies. In: NeurIPS (2022)
- Dong, Y., Chen, G., Peers, P., Zhang, J., Tong, X.: Appearance-from-motion: Recovering spatially varying surface reflectance under unknown lighting. ACM Transactions on Graphics (TOG) 33(6), 1–12 (2014)
- Dong, Z., Guo, C., Song, J., Chen, X., Geiger, A., Hilliges, O.: Pina: Learning a personalized implicit neural avatar from a single rgb-d video sequence. In: CVPR (2022)

- 16 Y. Chen et al.
- Feng, Y., Yang, J., Pollefeys, M., Black, M.J., Bolkart, T.: Capturing and animation of body and clothing from monocular video. In: SIGGRAPH Asia 2022 Conference Proceedings. SA '22 (2022)
- Goldman, D.B., Curless, B., Hertzmann, A., Seitz, S.M.: Shape and spatiallyvarying brdfs from photometric stereo. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(6), 1060–1071 (2009)
- Guan, P., Reiss, L., Hirshberg, D.A., Weiss, A., Black, M.J.: Drape: Dressing any person. TOG 31(4), 1–10 (2012)
- Guo, C., Jiang, T., Chen, X., Song, J., Hilliges, O.: Vid2avatar: 3d avatar reconstruction from videos in the wild via self-supervised scene decomposition. In: CVPR. pp. 12858–12868 (2023)
- Guo, K., Lincoln, P., Davidson, P., Busch, J., Yu, X., Whalen, M., Harvey, G., Orts-Escolano, S., Pandey, R., Dourgarian, J., et al.: The relightables: Volumetric performance capture of humans with realistic relighting. ACM Transactions on Graphics (ToG) 38(6), 1–19 (2019)
- Habermann, M., Liu, L., Xu, W., Pons-Moll, G., Zollhoefer, M., Theobalt, C.: Hdhumans: A hybrid approach for high-fidelity digital humans. ACM SCA 6(3), 1–23 (2023)
- Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Real-time deep dynamic characters. TOG 40(4), 1–16 (2021)
- Hasselgren, J., Hofmann, N., Munkberg, J.: Shape, light, and material decomposition from images using monte carlo rendering and denoising. Advances in Neural Information Processing Systems 35, 22856–22869 (2022)
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. NeurIPS 30 (2017)
- Ho, H.I., Xue, L., Song, J., Hilliges, O.: Learning locally editable virtual humans. In: CVPR. pp. 21024–21035 (2023)
- Hong, Y., Zhang, J., Jiang, B., Guo, Y., Liu, L., Bao, H.: Stereopifu: Depth aware clothed human digitization via stereo vision. In: CVPR. pp. 535–545 (2021)
- 30. Hu, S., Liu, Z.: Gauhuman: Articulated gaussian splatting from monocular human videos. arXiv preprint arXiv: (2023)
- Işık, M., Rünz, M., Georgopoulos, M., Khakhulin, T., Starck, J., Agapito, L., Nießner, M.: Humanrf: High-fidelity neural radiance fields for humans in motion. TOG 42(4), 1–12 (2023)
- Iqbal, U., Caliskan, A., Nagano, K., Khamis, S., Molchanov, P., Kautz, J.: Rana: Relightable articulated neural avatars. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 23142–23153 (2023)
- Iwase, S., Saito, S., Simon, T., Lombardi, S., Timur, B., Joshi, R., Prada, F., Shiratori, T., Sheikh, Y., Saragih, J.: Relightablehands: Efficient neural relighting of articulated hand models. In: CVPR (2023)
- Ji, C., Yu, T., Guo, K., Liu, J., Liu, Y.: Geometry-aware single-image full-body human relighting (October 2022)
- Jiang, K., Chen, S.Y., Fu, H., Gao, L.: Nerffacelighting: Implicit and disentangled face lighting representation leveraging generative prior in neural radiance fields. ACM Transactions on Graphics (TOG) (2023)
- Jiang, T., Chen, X., Song, J., Hilliges, O.: Instantavatar: Learning avatars from monocular video in 60 seconds. In: CVPR. pp. 16922–16932 (2023)
- 37. Jiang, W., Yi, K.M., Samei, G., Tuzel, O., Ranjan, A.: Neuman: Neural human radiance field from a single video. In: ECCV. pp. 402–418. Springer (2022)

- Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: Tensoir: Tensorial inverse rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 165–174 (2023)
- 39. Kajiya, J.T.: The rendering equation. In: Proceedings of the 13th annual conference on Computer graphics and interactive techniques. pp. 143–150 (1986)
- 40. Kanamori, Y., Endo, Y.: Relighting humans: occlusion-aware inverse rendering for full-body human images. ACM Trans. Graph. **37**(6) (dec 2018)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. TOG 42(4), 1–14 (2023)
- Kim, H., Nam, H., Kim, J., Park, J., Lee, S.: Laplacianfusion: Detailed 3d clothedhuman body reconstruction. ACM Transactions on Graphics (TOG) 41(6), 1–14 (2022)
- 43. Kocabas, M., Chang, J.H.R., Gabriel, J., Tuzel, O., Ranjan, A.: Hugs: Human gaussian splats (2023)
- 44. Kwon, Y., Liu, L., Fuchs, H., Habermann, M., Theobalt, C.: Deliffas: Deformable light fields for fast avatar synthesis. In: NeurIPS (2023)
- 45. Lagunas, M., Sun, X., Yang, J., Villegas, R., Zhang, J., Shu, Z., Masia, B., Gutierrez, D.: Single-image full-body human relighting. In: Eurographics Symposium on Rendering (EGSR). The Eurographics Association (2021). https: //doi.org/10.2312/sr.20211301
- Lawrence, J., Rusinkiewicz, S., Ramamoorthi, R.: Efficient brdf importance sampling using a factored representation. ACM Transactions on Graphics (ToG) 23(3), 496–505 (2004)
- 47. LeGendre, C., Ma, W.C., Fyffe, G., Flynn, J., Charbonnel, L., Busch, J., Debevec, P.: Deeplight: Learning illumination for unconstrained mobile mixed reality. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5918–5928 (2019)
- Li, R., Tanke, J., Vo, M., Zollhöfer, M., Gall, J., Kanazawa, A., Lassner, C.: Tava: Template-free animatable volumetric actors. In: ECCV. pp. 419–436. Springer (2022)
- Li, Z., Zheng, Z., Liu, Y., Zhou, B., Liu, Y.: Posevocab: Learning joint-structured pose embeddings for human avatar modeling. In: ACM SIGGRAPH Conference Proceedings (2023)
- Li, Z., Zheng, Z., Wang, L., Liu, Y.: Animatable gaussians: Learning posedependent gaussian maps for high-fidelity human avatar modeling. arXiv preprint arXiv:2311.16096 (2023)
- Li, Z., Zheng, Z., Zhang, H., Ji, C., Liu, Y.: Avatarcap: Animatable avatar conditioned monocular human volumetric capture. In: ECCV. pp. 322–341. Springer (2022)
- 52. Li, Z., Xu, Z., Ramamoorthi, R., Sunkavalli, K., Chandraker, M.: Learning to reconstruct shape and spatially-varying reflectance from a single image. ACM Transactions on Graphics (TOG) 37(6), 1–11 (2018)
- Lin, S., Li, Z., Su, Z., Zheng, Z., Zhang, H., Liu, Y.: Layga: Layered gaussian avatars for animatable clothing transfer. In: SIGGRAPH Conference Papers (2024)
- 54. Lin, S., Zhang, H., Zheng, Z., Shao, R., Liu, Y.: Learning implicit templates for point-based clothed human modeling. In: ECCV. pp. 210–228. Springer (2022)
- 55. Lin, W., Zheng, C., Yong, J.H., Xu, F.: Relightable and animatable neural avatars from videos. AAAI (2024)
- 56. Lipson, L., Teed, Z., Deng, J.: Raft-stereo: Multilevel recurrent field transforms for stereo matching. In: International Conference on 3D Vision (3DV) (2021)

- 18 Y. Chen et al.
- 57. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. TOG 40(6), 1–16 (2021)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. TOG 34(6), 1–16 (2015)
- Ma, Q., Saito, S., Yang, J., Tang, S., Black, M.J.: Scale: Modeling clothed humans with a surface codec of articulated local elements. In: CVPR. pp. 16082–16093 (2021)
- Ma, Q., Yang, J., Tang, S., Black, M.J.: The power of points for modeling humans in clothing. In: ICCV. pp. 10974–10984 (2021)
- Mescheder, L., Oechsle, M., Niemeyer, M., Nowozin, S., Geiger, A.: Occupancy networks: Learning 3d reconstruction in function space. In: CVPR. pp. 4460–4470 (2019)
- Mihajlovic, M., Zhang, Y., Black, M.J., Tang, S.: Leap: Learning articulated occupancy of people. In: CVPR. pp. 10461–10471 (2021)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV. pp. 405–421. Springer (2020)
- 64. Munkberg, J., Hasselgren, J., Shen, T., Gao, J., Chen, W., Evans, A., Müller, T., Fidler, S.: Extracting Triangular 3D Models, Materials, and Lighting From Images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8280–8290 (June 2022)
- 65. Nam, G., Lee, J.H., Gutierrez, D., Kim, M.H.: Practical svbrdf acquisition of 3d objects with unstructured flash photography. ACM Transactions on Graphics (TOG) 37(6), 1–12 (2018)
- 66. Pang, H., Zhu, H., Kortylewski, A., Theobalt, C., Habermann, M.: Ash: Animatable gaussian splats for efficient and photoreal human rendering (2023)
- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: Deepsdf: Learning continuous signed distance functions for shape representation. In: CVPR. pp. 165–174 (2019)
- Peng, B., Hu, J., Zhou, J., Zhang, J.: Selfnerf: Fast training nerf for human from monocular self-rotating video. arXiv preprint arXiv:2210.01651 (2022)
- Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: ICCV. pp. 14314– 14323 (2021)
- Peng, S., Zhang, S., Xu, Z., Geng, C., Jiang, B., Bao, H., Zhou, X.: Animatable neural implicit surfaces for creating avatars from videos. arXiv preprint arXiv:2203.08133 (2022)
- Remelli, E., Bagautdinov, T., Saito, S., Wu, C., Simon, T., Wei, S.E., Guo, K., Cao, Z., Prada, F., Saragih, J., et al.: Drivable volumetric avatars using texelaligned features. In: ACM SIGGRAPH 2022 Conference Proceedings. pp. 1–9 (2022)
- 72. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: CVPR (June 2020)
- Saito, S., Yang, J., Ma, Q., Black, M.J.: Scanimate: Weakly supervised learning of skinned clothed avatar networks. In: CVPR. pp. 2886–2897 (2021)
- 74. Shao, R., Zhang, H., Zhang, H., Chen, M., Cao, Y., Yu, T., Liu, Y.: Doublefield: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In: CVPR (2022)
- Shen, K., Guo, C., Kaufmann, M., Zarate, J.J., Valentin, J., Song, J., Hilliges, O.: X-avatar: Expressive human avatars. In: CVPR. pp. 16911–16921 (2023)

- 76. Shen, T., Gao, J., Yin, K., Liu, M.Y., Fidler, S.: Deep marching tetrahedra: a hybrid representation for high-resolution 3d shape synthesis. In: Advances in Neural Information Processing Systems (NeurIPS) (2021)
- 77. Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7495–7504 (2021)
- Stoll, C., Gall, J., De Aguiar, E., Thrun, S., Theobalt, C.: Video-based reconstruction of animatable human characters. TOG 29(6), 1–10 (2010)
- Su, S.Y., Bagautdinov, T., Rhodin, H.: Npc: Neural point characters from video. In: ICCV (2023)
- Su, S.Y., Yu, F., Zollhöfer, M., Rhodin, H.: A-nerf: Articulated neural radiance fields for learning human shape, appearance, and pose. NeurIPS 34, 12278–12291 (2021)
- Sun, W., Che, Y., Huang, H., Guo, Y.: Neural reconstruction of relightable human model from monocular video. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 397–407 (2023)
- Tajima, D., Kanamori, Y., Endo, Y.: Relighting humans in the wild: Monocular full-body human relighting with domain adaptation. Computer Graphics Forum (Proc. of Pacific Graphics 2021) 40(7), 205–216 (2021)
- Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. NeurIPS **33**, 7537–7547 (2020)
- Te, G., Li, X., Li, X., Wang, J., Hu, W., Lu, Y.: Neural capture of animatable 3d human from monocular video. In: ECCV. pp. 275–291. Springer (2022)
- Tiwari, G., Sarafianos, N., Tung, T., Pons-Moll, G.: Neural-gif: Neural generalized implicit functions for animating people in clothing. In: ICCV. pp. 11708–11718 (2021)
- Wang, S., Antić, B., Geiger, A., Tang, S.: Intrinsicavatar: Physically based inverse rendering of dynamic humans from monocular videos via explicit ray tracing. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2024)
- 87. Wang, S., Mihajlovic, M., Ma, Q., Geiger, A., Tang, S.: Metaavatar: Learning animatable clothed human models from few depth images. NeurIPS **34** (2021)
- Wang, S., Schwarz, K., Geiger, A., Tang, S.: Arah: Animatable volume rendering of articulated human sdfs. In: ECCV. pp. 1–19. Springer (2022)
- Wang, Y., Holynski, A., Zhang, X., Zhang, X.: Sunstage: Portrait reconstruction and relighting using the sun as a light stage. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20792–20802 (2023)
- 90. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE T-IP **13**(4), 600–612 (2004)
- Weng, C.Y., Curless, B., Srinivasan, P.P., Barron, J.T., Kemelmacher-Shlizerman, I.: Humannerf: Free-viewpoint rendering of moving people from monocular video. In: CVPR. pp. 16210–16220 (2022)
- 92. Xiang, D., Bagautdinov, T., Stuyck, T., Prada, F., Romero, J., Xu, W., Saito, S., Guo, J., Smith, B., Shiratori, T., et al.: Dressing avatars: Deep photorealistic appearance for physically simulated clothing. TOG 41(6), 1–15 (2022)
- Xiang, D., Prada, F., Bagautdinov, T., Xu, W., Dong, Y., Wen, H., Hodgins, J., Wu, C.: Modeling clothing as a separate layer for an animatable human avatar. TOG 40(6), 1–15 (2021)

- 20 Y. Chen et al.
- Xu, F., Liu, Y., Stoll, C., Tompkin, J., Bharaj, G., Dai, Q., Seidel, H.P., Kautz, J., Theobalt, C.: Video-based characters: creating new human performances from a multi-view video database. TOG 30(4), 1–10 (2011)
- 95. Xu, Z., Peng, S., Geng, C., Mou, L., Yan, Z., Sun, J., Bao, H., Zhou, X.: Relightable and animatable neural avatar from sparse-view video. In: Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR) (2024)
- 96. Yu, T., Zheng, Z., Guo, K., Liu, P., Dai, Q., Liu, Y.: Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In: CVPR. pp. 5746–5756 (2021)
- 97. Zhang, H., Lin, S., Shao, R., Zhang, Y., Zheng, Z., Huang, H., Guo, Y., Liu, Y.: Closet: Modeling clothed humans on continuous surface with explicit template decomposition. In: CVPR (2023)
- 98. Zhang, K., Luan, F., Wang, Q., Bala, K., Snavely, N.: Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5453–5462 (2021)
- Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
- 100. Zhang, X., Srinivasan, P.P., Deng, B., Debevec, P., Freeman, W.T., Barron, J.T.: Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. ACM Transactions on Graphics (ToG) 40(6), 1–18 (2021)
- 101. Zhao, H., Zhang, J., Lai, Y.K., Zheng, Z., Xie, Y., Liu, Y., Li, K.: High-fidelity human avatars from a single rgb camera. In: CVPR. pp. 15904–15913 (2022)
- Zheng, Y., Yifan, W., Wetzstein, G., Black, M.J., Hilliges, O.: Pointavatar: Deformable point-based head avatars from videos. In: CVPR. pp. 21057–21067 (2023)
- 103. Zheng, Z., Huang, H., Yu, T., Zhang, H., Guo, Y., Liu, Y.: Structured local radiance fields for human avatar modeling. In: CVPR. pp. 15893–15903 (2022)
- 104. Zheng, Z., Zhao, X., Zhang, H., Liu, B., Liu, Y.: Avatarrex: Real-time expressive full-body avatars. TOG 42(4) (2023). https://doi.org/10.1145/3592101
- 105. Zielonka, W., Bagautdinov, T., Saito, S., Zollhöfer, M., Thies, J., Romero, J.: Drivable 3d gaussian avatars (2023)