

NECA: Neural Customizable Human Avatar

Junjin Xiao¹ Qing Zhang^{1,3*} Zhan Xu² Wei-Shi Zheng^{1,3}

¹School of Computer Science and Engineering, Sun Yat-sen University, China ²Adobe Research

³Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

xiaojj37@mail2.sysu.edu.cn zhangq93@mail.sysu.edu.cn

zhaxu@adobe.com wszheng@ieee.org




Figure 1. **Neural customizable human avatar.** Our method takes as input monocular or sparse multi-view videos and outputs disentangled human representations, including normal, albedo, shadow, and illumination. Such disentanglement enables control of the learned human avatar with arbitrary poses/viewpoints and various customization options such as adjusting shape, shadow, lighting and texture.

Abstract

Human avatar has become a novel type of 3D asset with various applications. Ideally, a human avatar should be fully customizable to accommodate different settings and environments. In this work, we introduce NECA, an approach capable of learning versatile human representation from monocular or sparse-view videos, enabling granular customization across aspects such as pose, shadow, shape, lighting and texture. The core of our approach is to represent humans in complementary dual spaces and predict disentangled neural fields of geometry, albedo, shadow, as well as an external lighting, from which we are able to derive realistic rendering with high-frequency details via volumetric rendering. Extensive experiments demonstrate the advantage of our method over the state-of-the-art methods in photorealistic rendering, as well as various editing tasks such as novel pose synthesis and relighting. Our code is available at <https://github.com/iSEE-Laboratory/NECA>.

*Corresponding author.

1. Introduction

There is widespread demand for human avatars in many emerging applications such as the metaverse, telepresence, and 3D games. Among them, a basic common requirement for human avatars is that they should be fully editable to allow easy customization across aspects like pose, shape, lighting, texture, and even shadow. While there are numerous works on neural human avatar modelling, as illustrated in Tab. 1, they are mostly tailored for either animation or relighting purposes, failing to offer full customization capabilities for avatars thus being limited in overall practicality.

In this work, we present NECA, a novel framework that allows learning fully customizable neural human avatars with photorealistic rendering under any novel pose, viewpoint and lighting, as well as the ability to edit shape, texture, and shadow. To this end, we propose to learn human representation in both Canonical space and the surface space consisting of UV features and local tangent coordinates, so as to capture high-frequency dynamic variations and the shared characteristics across poses with geometry priors for harvesting

Table 1. Comparison of customizable attributes enabled by existing neural human avatar reconstruction methods.

Method	Pose	Lighting	Shadow	Shape	Texture
Neural Body [42]	✗	✗	✗	✗	✗
Neural Actor [31]	✓	✗	✗	✓	✓
Ani-NeRF [41]	✓	✗	✗	✗	✗
DS-NeRF [75]	✓	✗	✗	✓	✓
SA-NeRF [63]	✓	✗	✗	✓	✓
HumanNeRF [59]	✓	✗	✗	✗	✗
ARAH [56]	✓	✗	✗	✗	✗
TAVA [29]	✓	✗	✗	✗	✓
Relighting4D [9]	✗	✓	✗	✗	✗
MonoHuman [68]	✓	✗	✗	✗	✗
CustomHumans [20]	✓	✗	✗	✓	✓
UV Volumes [8]	✓	✗	✗	✓	✓
PoseVocab [30]	✓	✗	✗	✗	✗
Sun <i>et al.</i> [53]	✓	✓	✗	✗	✗
Ours	✓	✓	✓	✓	✓

high-fidelity novel pose synthesis. Besides, to achieve full customization capabilities, we employ distinct MLPs to predict geometry, albedo and shadow separately, and optimize a spherical environmental lighting. The entire framework is trained in a self-supervised manner, constrained only by photometric losses and normal regularization.

In summary, the main contributions of this work are:

- We present NECA, which to our knowledge, is the first framework that is able to learn fully customizable neural human avatars allowing photorealistic rendering.
- We propose to learn human representation in dual spaces for capturing high-frequency motion details and geometry-aware characteristics, and represent avatars as disentangled neural fields with distinct geometry and appearance attributes for flexible control.
- Extensive experiments demonstrate the broad editing capabilities of our approach, and our significant improvements over prior state-of-the-arts in various tasks such as novel pose synthesis and relighting.

2. Related Work

Human Avatar Reconstruction. Early methods explicitly represent human avatars with deformable templates such as SMPL [32] and its derivatives [37, 40]. Such template-based representation has a fixed mesh topology, and always fails to capture high-fidelity details. In recent years, the advent of neural implicit fields has witnessed the great success in 3D reconstruction [10, 34–36, 39, 52, 54, 55]. Among them, SDF-based methods [18, 46, 47, 61, 62, 74] demonstrate promising results in creating lifelike human models. However, they are unable to represent animatable humans. Various later methods [19, 22, 31, 41, 75] are designed to address this limitation by deforming the human body from

observation space to canonical space [59, 68]. Particularly, Animatable NeRF [41] proposes to learn per-pose related latent codes for skinning weights in the observation space for deformation, but it requires extra fine-tuning in the inference stage. ARAH [56] uses a root-finding algorithm [6] for transformations, thereby enabling novel pose rendering, but at the cost of low rendering speed. Note, as lighting and texture information are usually integrated into a single MLP in these methods, they do not support any appearance editing. There also exist some works that allow appearance editing [8, 20, 29, 63], but they mainly focus on shape, pose, and texture adjustments.

Human Relighting. A widely adopted technique for relighting human is inverse rendering [2, 14, 27, 33, 38, 48, 67], which aims to disentangle geometry, material, and lighting from observed images. However, they are either restricted to novel view and novel pose synthesis because of lacking underlying 3D representation [27, 38], or depend heavily on the hard-to-obtain one-light-at-a-time (OLAT) images [14]. Recent implicit-field based methods [1, 11, 24, 51, 58, 69, 72, 73] address these limitations by directly learning the underlying 3D structure and unknown illumination from the input images, relying on photometric or geometry loss. While these methods demonstrate promising novel view synthesis results and fine-grained geometry reconstruction, they are mostly tailored for static scenes and not applicable to dynamic humans with complex non-rigid motion and shadow. By extending NeRF [35], some concurrent works [9, 53, 64] gain applicability to dynamic humans. Nevertheless, they struggle to render high-fidelity humans under arbitrary novel poses [9], and none of them provide capabilities for texture and shadow editing.

3. Method

Our method aims to generate a fully customizable human avatar from as few as a monocular video. We represent the human avatar in complementary dual spaces (Sec. 3.1), for capturing both high-frequency pose-aware features and geometry-aware subject-level characteristics. From such representation, we derive the disentangled neural fields, each dedicated to a specific attribute such as SDF, shadow, albedo, as well as environmental lighting (Sec. 3.2). The disentanglement is learned in a self-supervised manner, with only photometric losses and normal regularization (Sec. 3.3). Once the avatar has been generated, we can adjust each decomposed attributes separately for diverse applications, ranging from reposing and relighting to reshaping, shadow editing and texture swapping, as demonstrated in our experiments. An overview of our method is shown in Fig. 2.

3.1. Dual-space Dynamic Human Representation

Motivation. Most existing works define neural representation in either Canonical space or the surface space of the de-




Figure 2. **Overview of NECA.** We first sample points along the camera ray and transform the query points from observation space to canonical space. Next, we query the pose-aware feature by projecting points to factorized tri-plane that per-pose optimized. Then we construct tangent space of the nearest surface point to the query points on SMPL, and obtain the subject-level feature by concatenating the tangent space local coordinate and the learned latent code in the surface space. Finally, to enable flexible customization, we disentangle the neural fields into attributes including SDF, albedo and shadow, as well as a learnable environmental lighting, by decoding the extracted features with distinct MLPs. The entire network is trained in a self-supervised manner, with only photometric losses and normal regularization.

formable human template [32], each with its own advantages and limitations. Concretely, Canonical space is more generic and offers greater feature capacity to capture high-frequency details, but may lead to disturbing misalignment across poses because of ignoring the shape prior of the human body. In contrast, the surface space leverages the shape prior, focusing on the proximate surface volume. Nevertheless, it relies much on pose prior, which limits the expressiveness of the learned representation. Therefore, we propose to incorporate the strengths of both spaces. We define the pose-aware representation in Canonical space to mitigate the influence of strong pose prior and enhance expressiveness for diverse motion dynamics, and the subject-level representation in the surface space due to its consistency across poses. In a recent work [31], Canonical coordinates and UV features are both utilized in NeRF. They however do not explicitly learn Canonical space feature representation, resulting in relatively limited capability to represent motion variations. Moreover, the surface space feature is only employed for color prediction, leaving the shape prior unexploited.

Canonical space for pose-aware feature. Pose changes lead to high-frequency appearance variations, such as dynamics of clothing and shadows resulting from self-occlusion. To better capture such intricate pose-dependent details, We adopt tri-plane representation [21, 44, 49] to learn pose-aware features. Concretely, we construct tri-plane $\mathbf{T}_\theta^{\{XY, XZ, YZ\}} \in \mathbb{R}^{L \times L \times D}$ for each training pose θ in the canonical space. Given a query point \mathbf{x}_o in the observation space, we transform it to the canonical space using inverse

linear blend skinning [32]:

$$\mathbf{x}_c = \sum_b^{N_j} W_b \mathbf{B}_b^{-1} \mathbf{x}_o. \quad (1)$$

Here, b is index of joints, N_j denotes the total number of joints. W_b denotes the skinning weight of point \mathbf{x}_o associated with joint b , and \mathbf{B}_b^{-1} is the transformation of joint b from the observation space to the canonical space. We follow the approach outlined in [31] to calculate skinning weight W_b . We can then fetch the pose-aware feature \mathbf{p}_o for \mathbf{x}_o by projecting \mathbf{x}_c onto three feature planes, and sample corresponding features:

$$\mathbf{p}_o = \oplus(\mathbf{t}_\theta^{XY}, \mathbf{t}_\theta^{XZ}, \mathbf{t}_\theta^{YZ}, \theta), \quad (2)$$

where \oplus denotes concatenation, $\mathbf{t}_\theta^{\{XY, XZ, YZ\}}$ are the sampled features from the feature planes by bilinear interpolation. We append pose θ to be more robust.

However, this straightforward implementation incurs significant memory costs due to the per-pose unique tri-plane feature. To reduce memory cost, we employ CP decomposition [4, 5, 17] to factorize the feature planes into orthogonal feature vectors. This is achieved by factorizing the axis-aligned planes into a sum of R outer products of 1D vectors:

$$\mathbf{T}_\theta^{mn} = \sum_{r=1}^R \mathbf{v}_r^m \circ \mathbf{v}_r^n \circ \mathbf{v}_r^{mn}, \quad (3)$$

where m, n are different axes from $\{X, Y, Z\}$. $\mathbf{v}_r^{\{x, y, z\}} \in \mathbb{R}^L$ represent learned axis-aligned vectors, and $\mathbf{v}_r^{\{xy, xz, yz\}} \in$




Figure 3. **Local coordinate used by [20, 31, 70] and the one used by us.** (a) Previous works defined their local coordinate in UV space. The three pink points will be mapped to the same red point on the surface, thus share the same feature. Such many-to-one problem is also mentioned in [63]. (b) We complement UV space with local tangent space. Although pink points are still mapped to the same surface point, now they have different coordinates in the local tangent space. Their features are therefore different. As previous works, our local coordinate is also rotation invariant.

\mathbb{R}^D represent learned vectors in the feature dimension. We dramatically compress the tri-plane, reducing the memory complexity from $\mathcal{O}(L^2)$ to $\mathcal{O}(RL)$ where $R \ll L$.

The learned factorized tri-planes from all training poses form a pose-aware feature dictionary $\{\mathbf{T}_\theta\}$. During test time, given an unseen pose, we select the top 5 most similar training poses, and average the features obtained from corresponding tri-planes, weighted by the pose similarity. A recent work, PoseVocab [30], also uses a pose-aware embedding dictionary, but it learns global features for each local joint and its feature representativeness is degraded for achieving better memory efficiency. In our case, pose-aware features are learned with a global pose based on the factorized tri-plane, resulting in features with strong representation power while keeping a low memory footprint.

Surface space for subject-level feature. Subject-level feature describes pose-agnostic characteristics. We employ the SMPL model [32] as prior to align poses, and define our subject-level feature on its UV map. While earlier methods [20, 31, 70] have adopted the UV volume as pose-agnostic space, their approaches incur the many-to-one problem [63] illustrated in Fig. 3(a): multiple 3D points are mapped to the same point on edges and vertices, due to the finite area of each triangle. Inspired by [3], we address such limitation by constructing tangent space for any point on the SMPL surface, and encode the query 3D point based on its local coordinate in the tangent space of the nearest surface point, as a supplement to the UV feature map.

Specifically, for a surface point \mathbf{x}_s , we construct its associated tangent space by creating the local coordinates from its tangent vector \mathbf{t}_s , bitangent vector \mathbf{b}_s , and normal vector \mathbf{n}_s :

$$\mathcal{M}(\mathbf{x}_s) = [\mathbf{t}_s, \mathbf{b}_s, \mathbf{n}_s]^T. \quad (4)$$

More details on calculating TBN matrix \mathcal{M}_s can be found in our supplementary material. For each 3D query point \mathbf{x}_o in the observation space, we first find its nearest surface

point \mathbf{x}_s^* on the SMPL mesh, and then transform the global coordinate of \mathbf{x}_o to the local tangent space.

$$\mathbf{x}_l = \mathcal{M}(\mathbf{x}_s^*)(\mathbf{x}_o - \mathbf{x}_s^*). \quad (5)$$

Eq. (5) can differentiate all 3D points given a particular tangent space without the many-to-one problem. The derived \mathbf{x}_l is invariant to human pose change. To indicate the position of \mathbf{x}_s and provide geometry awareness, we calculate the associated UV feature \mathbf{g}_s of the surface point \mathbf{x}_s^* .

$$\mathbf{g}_s = \mathcal{B}_{u_s^*, v_s^*}(\mathbf{C}_{[\mathcal{V}_f]}). \quad (6)$$

Here, $\{\mathbf{C}_v\}$ is a set of learnable latent codes associated with SMPL vertices $v = [1, \dots, N_v]$. \mathcal{V}_f are the three vertices of the triangle \mathbf{x}_s falls in, and $\mathbf{C}_{[\mathcal{V}_f]}$ denotes their latent codes. $\mathcal{B}_{u_s^*, v_s^*}(\cdot)$ is the interpolation operation based on the barycentric coordinates (u_s^*, v_s^*) .

With both \mathbf{x}_l and \mathbf{g}_s , our subject-level feature \mathbf{s}_o of the query point \mathbf{x}_o is formulated as the concatenation of them. We in practice also encode \mathbf{x}_l with the positional encoding function introduced in [35].

3.2. Disentangled Neural Fields

The proposed NECA is designed to facilitate diverse customization tasks, including shape, pose, texture, lighting and shadow. Reposing and reshaping are achieved by the adoption of SMPL model. We enable other customization functionalities by decoding the extracted features into distinct neural fields. Unlike [35] which produces density and radiance, our method consumes the canonical position \mathbf{x}_c , the viewing direction \mathbf{e} , the pose-aware feature \mathbf{p}_o and the subject-level feature \mathbf{s}_o into SDF d , albedo \mathbf{a} , shadow v , while simultaneously optimizing the representation \mathbf{L} for lighting. Such an approach expands our editing capability by exposing more controllable attributes for rendering.

The overall process $(\mathbf{x}_c, \mathbf{e}, \mathbf{s}_o, \mathbf{p}_o) \rightarrow (d, \mathbf{a}, v)$ is achieved by three separate MLPs. First, the geometry MLP takes all the inputs except for the viewing direction and outputs SDF d and a latent feature \mathbf{h} , similar to [30, 56]:

$$d, \mathbf{h} = \text{MLP}_{geo}(\mathbf{x}_c, \mathbf{s}_o, \mathbf{p}_o; \Theta_{geo}). \quad (7)$$

The output SDF can be converted to density σ with methods such as [66] for volumetric rendering. Moreover, we can calculate normal vector by normalizing the gradient of SDF d with respect to \mathbf{x}_c [43, 60]:

$$\mathbf{n}_c = \frac{\partial d}{\partial \mathbf{x}_c} / \left\| \frac{\partial d}{\partial \mathbf{x}_c} \right\|_2. \quad (8)$$

The normal vector provides critical information about geometry details and orientation, which also benefits our prediction of shadow by another shadow MLP:

$$v = \text{MLP}_{shadow}(\mathbf{e}, \mathbf{n}_c, \mathbf{h}; \Theta_{shadow}). \quad (9)$$




Figure 4. Qualitative comparison of novel pose synthesis on ZJU-MoCap dataset.

Finally, the albedo MLP predicts the albedo given latent feature \mathbf{h} and \mathbf{x}_c :

$$\mathbf{a} = \text{MLP}_{\text{albedo}}(\mathbf{x}_c, \mathbf{h}; \Theta_{\text{albedo}}). \quad (10)$$

Besides the intrinsic characteristics, we also optimize a particular representation of the environmental lighting. We assume the lighting during capturing the input videos is unchanged and in gray-scale, and utilize the learnable light probes [9, 13] to represent it in latitude-longitude format as $\mathbf{I} \in \mathbb{R}^{16 \times 32 \times 1}$. Note that although we assume the light is gray-scale during training, we can replace it with arbitrary light with RGB color during testing time for relighting.

With all the above disentangled neural fields and the lighting, the color of a query point \mathbf{x}_o can be obtained using the following rendering equation as in [9, 25, 45]:

$$\mathbf{c} = \mathbf{a} \odot v \odot \sum_i I_i(\mathbf{n}_o \cdot \boldsymbol{\omega}_i) \Delta \boldsymbol{\omega}_i, \quad (11)$$

where \odot denotes the Hadamard Product, \mathbf{n}_o represents the normal vector transformed to the observation space, $\boldsymbol{\omega}_i$ denotes the direction from query points \mathbf{x}_o in the observation space to the light position, and (\cdot) indicates the dot product.

To render the final color of a pixel, we follow the volumetric rendering [26], by integrating the colors $\{\mathbf{c}_i, i \in [1, N]\}$ of N sampled query points along the ray r from that pixel position based on their density $\{\sigma_i\}$.

$$C(r) = \sum_i^N \exp(-\sum_{j=1}^{i-1} \sigma_j \delta_j) (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i, \quad (12)$$

where δ_i indicates the distance between two sampled points.

Our disentangled neural fields can produce higher-fidelity rendering while maximizing controllability for customization during rendering. Unlike traditional physically-based rendering approaches which rely on BRDFs for color and explicit visibility for shadows, our method leverages a data-driven approach for inverse rendering. Such choice is more robust to noise, and can better represent complex shadows, such as those caused by clothes wrinkles.

3.3. Loss Function

We supervise the training of NECA with photometric losses and normal regularization described below.

The photometric losses measure the consistency of the rendering of NECA to the ground-truth targets. First, we supervise our rendered pixel colors $C(r)$ by minimizing the squared Euclidean distance to the target pixel colors $C_{gt}(r)$:

$$\mathcal{L}_c = \|C_{gt}(r) - C(r)\|_2^2. \quad (13)$$

Next, we employ a perceptual loss LPIPS [71] with VGG network [50] as backbone to enhance the quality of the rendered images:

$$\mathcal{L}_p = \|VGG_{gt}(r) - VGG(r)\|_2^2. \quad (14)$$

Note, we apply patch-based ray sampling [59] to keep memory efficiency.

The foreground masks serve as supervisory signal for geometry learning. We employ the binary cross-entropy loss to reduce discrepancies in determining whether each ray intersects the subject or not [43, 65].

$$\mathcal{L}_m = BCE(\text{sigmoid}(-\rho d_r), M(r)). \quad (15)$$

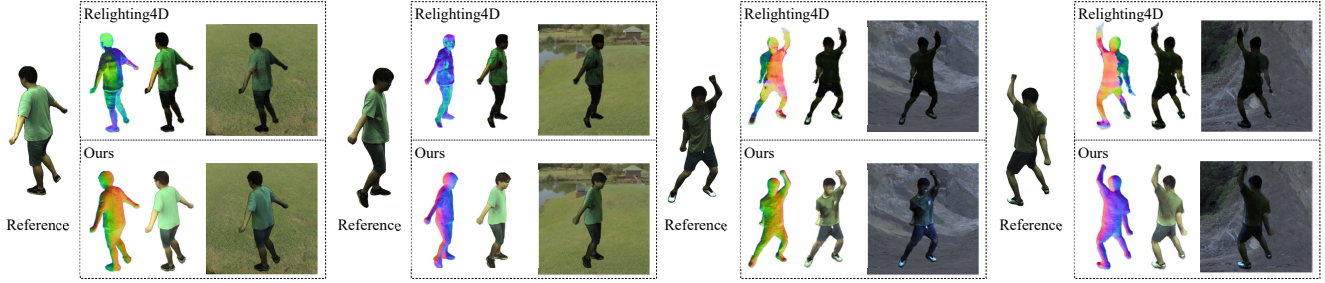


Figure 5. **Qualitative comparison of relighting under novel pose and view on ZJU-MoCap dataset.** Given the original frame as reference, we compare with Relighting4D on the estimated normal and albedo, as well as the generated relighting results. As shown, our method outperforms Relighting4D in both appearance disentanglement and relighting. Please see the supplementary material for more results.

Here, ρ is a scale factor, d_r is the minimal SDF value of ray r , and $M(r)$ is the ground truth binary mask.

We incorporate normal regularization to better capture geometry details, which is also beneficial to the prediction of albedo and shadow. One regularization is the eikonal loss [12] that encourages the SDF field to be smooth:

$$\mathcal{L}_e = \left(\left\| \frac{\partial d}{\partial \mathbf{x}_c} \right\|_2 - 1 \right)^2. \quad (16)$$

We also notice that normals occasionally face backward due to the inherent symmetry of SMPL. To address this, we introduce another normal regularization \mathcal{L}_n :

$$\mathcal{L}_n = \max(-\mathbf{n}_c \cdot \mathbf{n}_s, 0), \quad (17)$$

where \mathbf{n}_c represents the normal vector as defined in Eq. (8), and \mathbf{n}_s denotes the normal vector of the nearest SMPL vertex. $\mathbf{n}_c \cdot \mathbf{n}_s$ indicates the dot product of the two vectors.

The overall training objective function for our network is formulated as follows:

$$\mathcal{L} = \mathcal{L}_c + \lambda_p \mathcal{L}_p + \lambda_m \mathcal{L}_m + \lambda_e \mathcal{L}_e + \lambda_n \mathcal{L}_n, \quad (18)$$

where we set $\lambda_p = 0.1$, $\lambda_m = 1$, $\lambda_e = 0.1$, $\lambda_n = 0.1$.

3.4. Implementation Details

We implement our network in Pytorch and train it for 200K iterations with a mini-batch size of 1 on an NVidia RTX 4080 GPU. The entire network is optimized using the Adam optimizer [28] with a learning rate decaying exponentially from an initial value of 5×10^{-4} to 1×10^{-4} . For each training batch, we sample 1024 rays and 64 points per ray. The patch size for the perceptual loss is set as 32, with a single patch in use. We set $|\mathbf{C}_v| \in \mathbb{R}^{16}$ for the latent code attached to SMPL vertex. For tri-plane decomposition, we set the grid resolution $L_x \times L_y \times L_z$ as $512 \times 512 \times 128$, and a lower resolution along z-axis is used for memory efficiency due to the flatness of human body [7]. We set $\mathbf{v}_r^{\{xy,xz,yz\}} \in \mathbb{R}^{32}$, and $R = 48$.

	Novel View			Novel Pose			Param.
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	
NB [42]	28.2	0.944	0.096	23.7	0.893	0.145	4.37M
AN [41]	25.8	0.913	0.146	23.2	0.887	0.165	1.96M
DS [75]	27.8	0.937	0.125	23.8	0.892	0.157	0.50M
ARAH [56]	28.4	0.947	0.082	24.5	0.909	0.109	87.09M
PV [30]	27.5	0.939	0.068	24.3	0.905	0.095	24.08M
Ours	28.5	<u>0.946</u>	0.067	25.0	0.914	0.091	18.37M

Table 2. **Quantitative comparison of novel pose/view synthesis on all 9 subjects on ZJU-MoCap dataset.** As shown, our method achieves the best results with comparable model size.

4. Experiments

4.1. Evaluation Datasets and Metrics

Datasets. We evaluate our method on an indoor dataset ZJU-MoCap [42], an outdoor dataset NeuMan [23]. ZJU-MoCap consists of 9 sequences with dynamic humans with complex motion captured by multiple cameras in laboratory environment. We train and evaluate our method on all 9 sequences, and follow the training and testing view and pose split introduced in [42]. NeuMan consists of videos of approximately 10 to 20 seconds in length, where the “Bike” and “Seattle” sequences are employed for training and evaluation. Note, we also evaluate our method on DeepCap [15] and DynaCap [16] datasets, and quantitatively validate our superiority in relighting using a synthetic dataset. Please see the supplementary material for details.

Metrics. We follow previous work to evaluate our method on novel pose synthesis and relighting using three standard metrics, *i.e.*, Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM) [57], and Learned Perceptual Image Patch Similarity (LPIPS) [71]. Akin to [42], we employ a bounding box surrounding the human as mask for metric calculation. Note, for other edits allowed by our method (*e.g.*, reshaping, retexturing, and reshadowing), we only provide visual results as there are no ground-truths.

4.2. Comparison with State-of-the-art Methods

Evaluation on novel pose synthesis. We compare our method against the following five state-of-the-art methods

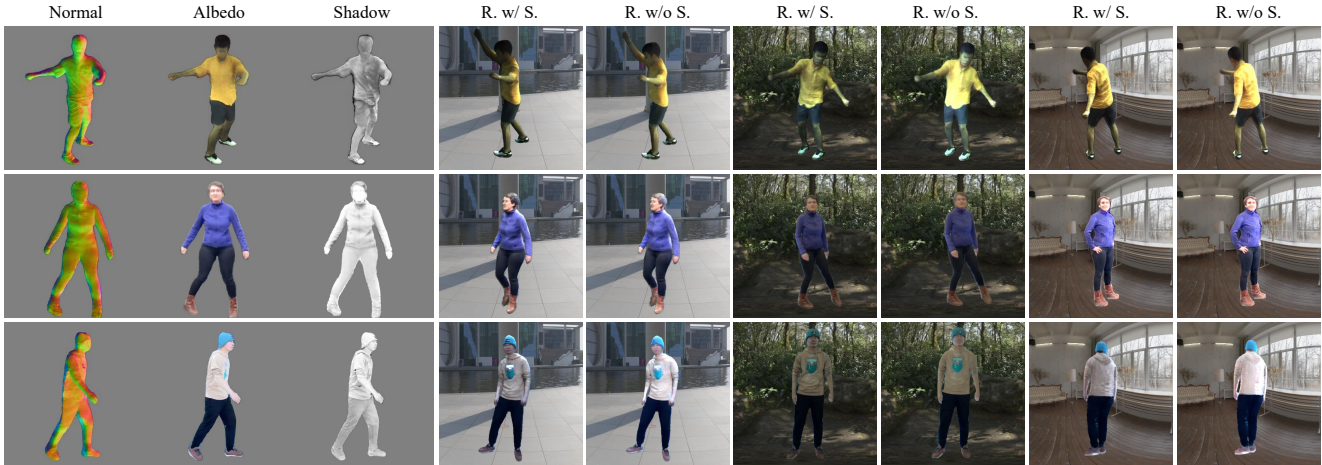


Figure 6. **More relighting results on the indoor ZJU-MoCap (top row) and the outdoor NeuMan (bottom two rows) datasets.** “R. w/ S.” and “R. w/o S.” refer to results with and without shadows, respectively. Our method produces high-fidelity dynamic human renderings that respond faithfully to novel lighting. Please check the supplementary videos for more results.



Figure 7. **Results of shape editing on subjects ‘315’ and ‘377’ of ZJU-MoCap.**

on novel pose synthesis: Neural Body (NB) [42], Animatable NeRF (AN) [41], Dual-Space NeRF (DS) [75], ARAH [56], and PoseVocab (PV) [30]. For fair comparison, we produce their results using publicly-available implementation provided by the authors with recommended parameter setting. Tab. 2 reports the quantitative results, where we can see that our method outperforms the compared baselines in all three metrics. It is worth mentioning that our method also produces good novel view synthesis results. In addition, we show visual comparison in Fig. 4. Comparing the results, it is clear that our method can render photorealistic images with much more high-frequency details.

Evaluation on relighting. For relighting, we compare our method with Relighting4D [9], by naively placing the relighted human in a new background. We did not compare with [53, 64], since they have no publicly available code for now. Fig. 5 gives the comparison results on relighting, as well as normal and albedo. As shown, our method generates



Figure 8. **Results of retexturing.** Our method allows to swap textures between body parts from different subjects.




Figure 9. **Results of shadow editing on the DynaCap dataset [16].** Our method allows to remove shadows (left two columns) and also transfer shadows between humans (right three columns).

physically realistic relighting results under novel pose, while Relighting4D fails to generate natural-looking results when the human in videos performs dramatic movement. Fig. 6 shows more relighting results produced by our method, in both indoor and outdoor scenes. As can be seen, for all these cases, our method produces high-quality results.

4.3. More Customization Results

In addition to novel pose synthesis and relighting, our method also allows editing of shape, texture, and shadow.




Figure 10. **Effect of normal regularization \mathcal{L}_n .** As shown, adopting \mathcal{L}_n benefits obtaining more accurate normal map.

Model	Num of views	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
w/o normal reg.	4	26.6	0.928	0.084
w/o pose-aware	4	26.6	0.929	0.078
w/o subject-level	4	26.6	0.928	0.079
full method	1	24.2	0.906	0.117
full method	4	26.8	0.930	0.079

Table 3. **Ablation studies on ZJU-MoCap.** “w/o normal reg.” refers to the exclusion of the proposed normal regularization \mathcal{L}_n , while “w/o pose-aware” indicates removal of the pose-aware feature, and “w/o subject-level” means that the subject-level feature is omitted. Note, here we report the average results of novel view and novel pose synthesis.

Fig. 7 shows our shape editing results, which are produced by adopting our local coordinate defined in tangent space to transform query points between different shape spaces. Following [63], we achieve retexturing by mapping the 3D query points from the target subject to the source subject. For each point in the canonical space of the target subject, we first adjust its position according to the bounding boxes of the two subjects, and then query the color from the canonical space of the source subject. By composing multiple learned neural human fields, we can achieve texture swapping, as shown in Fig. 8. Moreover, Fig. 9 shows that our method is applicable to shadow editing, and allows to eliminate self-cast shadows caused by clothing wrinkles and locally transfer shadows from one person to another. More customization results can be found in the supplementary material.

4.4. More Analysis

Ablation studies. Here we conduct ablation studies to validate the effectiveness of our subject-level feature, pose-aware feature, and normal regularization. Besides, we examine the performance of our method on single-view video. Tab. 3 summarizes the quantitative results of ablation studies. Fig. 10 presents visual comparison to verify the necessity of our normal regularization, while Fig. 11 demonstrates the benefit of subject-level and pose-aware features to our approach. Please see also the supplementary material for more results and analysis on different numbers of R in Eq. (3) and other possible alternatives (*i.e.*, relative position, direction, and

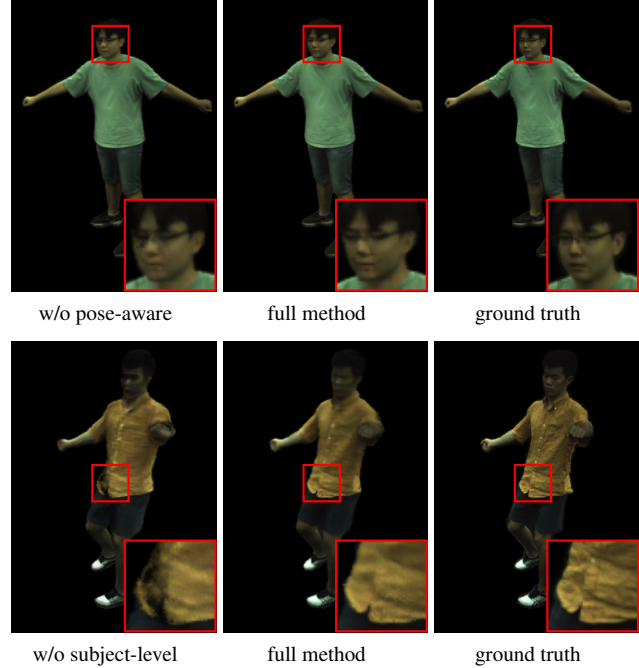


Figure 11. **Effect of pose-aware and subject-level features.** As shown, these features help achieve better alignment effect.

UVH) for our local tangent coordinate.

Limitations. Since obtaining the subject-level feature requires to project the query point to the nearest point on SMPL faces, our performance can be sensitive to the estimated SMPL parameters. In addition, as we do not explicitly model visibility, shadows under complex novel poses may be erroneous. Moreover, we may generate blurred retexturing results due to large pose variance, and unrealistic reshaping results for extreme shape. Finally, similar to the vanilla NeRF, our method has to be optimized for each human separately.

5. Conclusion

We have presented a novel framework called NECA for learning fully customizable human avatars from sparse-view or even monocular videos. In contrast to previous methods which offer limited editing capabilities, we provide neural avatar that allows high-fidelity editing on pose, viewpoint, lighting, shape, texture, and shadow. Extensive experiments validate the versatility and practicality of our approach, and our improvements over prior state-of-the-arts in novel pose synthesis and relighting. We hope that our work can shed light on the creation of customizable human avatar and its related applications.

Acknowledgement. This work was supported by the National Natural Science Foundation of China (U21A20471, 62072191), Guangdong Basic and Applied Basic Research Foundation (2023A1515030002, 2023B1515040025).

References

- [1] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *CVPR*, 2022. 2
- [2] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In *CVPR*, 2020. 2
- [3] James F. Blinn. Simulation of wrinkled surfaces. In *SIGGRAPH*, 1978. 4
- [4] J. Douglas Carroll and Jih Jie Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of “eckart-young” decomposition. *Psychometrika*, 35: 283–319, 1970. 3
- [5] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, 2022. 3
- [6] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 2
- [7] Xu Chen, Tianjian Jiang, Jie Song, Max Rietmann, Andreas Geiger, Michael J. Black, and Otmar Hilliges. Fast-SNARF: A fast deformer for articulated neural fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–15, 2023. 6
- [8] Yue Chen, Xuan Wang, Xingyu Chen, Qi Zhang, Xiaoyu Li, Yu Guo, Jue Wang, and Fei Wang. Uv volumes for real-time rendering of editable free-view human performance. In *CVPR*, 2023. 2
- [9] Zhaoxi Chen and Ziwei Liu. Relighting4d: Neural relightable human from videos. In *ECCV*, 2022. 2, 5, 7
- [10] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 2
- [11] Enric Corona, Mihai Zanfir, Thiemo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. Structured 3d features for reconstructing controllable avatars. In *CVPR*, 2023. 2
- [12] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*, 2020. 6
- [13] Jie Guo, Zijing Zong, Yadong Song, Xihao Fu, Chengzhi Tao, Yanwen Guo, and Ling-Qi Yan. Efficient light probes for real-time global illumination. *ACM Trans. Graphics*, 41(6), 2022. 5
- [14] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, Danhang Tang, Anastasia Tkach, Adarsh Kowdle, Emily Cooper, Mingsong Dou, Sean Fanello, Graham Fyffe, Christoph Rhemann, Jonathan Taylor, Paul Debevec, and Shahram Izadi. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Trans. Graphics*, 38(6), 2019. 2
- [15] Marc Habermann, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *CVPR*, 2020. 6
- [16] Marc Habermann, Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Gerard Pons-Moll, and Christian Theobalt. Real-time deep dynamic characters. *ACM Trans. Graphics*, 40(4), 2021. 6, 7
- [17] Richard A. Harshman. Foundations of the parafac procedure: Models and conditions for an “explanatory” multi-model factor analysis. 1970. 3
- [18] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. In *NeurIPS*, 2020. 2
- [19] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021. 2
- [20] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *CVPR*, 2023. 2, 4
- [21] Wenbo Hu, Yuling Wang, Lin Ma, Bangbang Yang, Lin Gao, Xiao Liu, and Yuewen Ma. Tri-miprf: Tri-mip representation for efficient anti-aliasing neural radiance fields. In *ICCV*, 2023. 3
- [22] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 2
- [23] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 6
- [24] Haian Jin, Isabella Liu, Peijia Xu, Xiaoshuai Zhang, Songfang Han, Sai Bi, Xiaowei Zhou, Zexiang Xu, and Hao Su. Tensorf: Tensorial inverse rendering. In *CVPR*, 2023. 2
- [25] James T. Kajiya. The rendering equation. In *SIGGRAPH*, 1986. 5
- [26] James T. Kajiya and Brian P Von Herzen. Ray tracing volume densities. In *SIGGRAPH*, 1984. 5
- [27] Yoshihiro Kanamori and Yuki Endo. Relighting humans: Occlusion-aware inverse rendering for full-body human images. *ACM Trans. Graphics*, 37(6), 2018. 2
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [29] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhofer, Jurgen Gall, Angjoo Kanazawa, and Christoph Lassner. Tava: Template-free animatable volumetric actors. In *ECCV*, 2022. 2
- [30] Zhe Li, Zerong Zheng, Yuxiao Liu, Boyao Zhou, and Yebin Liu. Posevocab: Learning joint-structured pose embeddings for human avatar modeling. In *SIGGRAPH*, 2023. 2, 4, 5, 6, 7
- [31] 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. Graphics*, 40(6), 2021. 2, 3, 4
- [32] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics*, 34(6):248:1–248:16, 2015. 2, 3, 4
- [33] Abhimitra Meka, Rohit Pandey, Christian Häne, Sergio Orts-Escolano, Peter Barnum, Philip David-Son, Daniel Erickson, Yinda Zhang, Jonathan Taylor, Sofien Bouaziz, Chloe Legendre, Wan-Chun Ma, Ryan Overbeck, Thabo Beeler, Paul Debevec, Shahram Izadi, Christian Theobalt, Christoph Rhemann, and Sean Fanello. Deep relightable textures: Volumet-

- ric performance capture with neural rendering. *ACM Trans. Graphics*, 39(6), 2020. 2
- [34] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 2
- [35] 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 *ECCV*, 2020. 2, 4
- [36] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *CVPR*, 2022. 2
- [37] Ahmed A A Osman, Timo Bolkart, and Michael J. Black. STAR: A sparse trained articulated human body regressor. In *ECCV*, 2020. 2
- [38] Rohit Pandey, Sergio Orts Escolano, Chloe Legendre, Christian Häne, Sofien Bouaziz, Christoph Rhemann, Paul Debevec, and Sean Fanello. Total relighting: Learning to relight portraits for background replacement. *ACM Trans. Graphics*, 40(4), 2021. 2
- [39] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 2
- [40] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 2
- [41] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 2, 5, 6, 7
- [42] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2, 5, 6, 7
- [43] Sida Peng, Shangzhan Zhang, Zhen Xu, Chen Geng, Boyi Jiang, Hujun Bao, and Xiaowei Zhou. Animatable neural implicit surfaces for creating avatars from videos. *arXiv preprint arXiv:2203.08133*, 2022. 4, 5
- [44] Sida Peng, Yunzhi Yan, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Representing volumetric videos as dynamic MLP maps. In *CVPR*, 2023. 3
- [45] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *ECCV*, 2022. 5
- [46] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2
- [47] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 2
- [48] Carolin Schmitt, Simon Donné, Gernot Riegler, Vladlen Koltun, and Andreas Geiger. On joint estimation of pose, geometry and svbrdf from a handheld scanner. In *CVPR*, 2020. 2
- [49] Ruizhi Shao, Zerong Zheng, Hanzhang Tu, Boning Liu, Hongwen Zhang, and Yebin Liu. Tensor4d: Efficient neural 4d decomposition for high-fidelity dynamic reconstruction and rendering. In *CVPR*, 2023. 3
- [50] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. 5
- [51] Pratul P. Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T. Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021. 2
- [52] J.-M. Sun, T. Wu, and L. Gao. Recent advances in implicit representation based 3d shape generation. *Visual Intelligence*, 2024. 2
- [53] Wenzhang Sun, Yunlong Che, Han Huang, and Yandong Guo. Neural reconstruction of relightable human model from monocular video. In *ICCV*, 2023. 2, 7
- [54] Jianwei Tang, Jieming Wang, and Jian-Fang Hu. Predicting human poses via recurrent attention network. *Visual Intelligence*, 2023. 2
- [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. In *NeurIPS*, 2021. 2
- [56] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdf. In *ECCV*, 2022. 2, 4, 5, 6, 7
- [57] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. 6
- [58] Zian Wang, Tianchang Shen, Jun Gao, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojcic, Wenzheng Chen, and Sanja Fidler. Neural fields meet explicit geometric representations for inverse rendering of urban scenes. In *CVPR*, 2023. 2
- [59] Chung-Yi Weng, Brian Curless, Pratul P. Srinivasan, Jonathan T. Barron, and Ira Kemelmacher-Shlizerman. HumanNeRF: Free-viewpoint rendering of moving people from monocular video. In *CVPR*, 2022. 2, 5
- [60] Tong Wu, Jia-Mu Sun, Yu-Kun Lai, and Lin Gao. De-nerf: Decoupled neural radiance fields for view-consistent appearance editing and high-frequency environmental relighting. In *SIGGRAPH*, 2023. 4
- [61] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: Implicit Clothed humans Obtained from Normals. In *CVPR*, 2022. 2
- [62] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. ECON: Explicit Clothed humans Optimized via Normal integration. In *CVPR*, 2023. 2
- [63] Tianhan Xu, Yasuhiro Fujita, and Eiichi Matsumoto. Surface-aligned neural radiance fields for controllable 3d human synthesis. In *CVPR*, 2022. 2, 4, 8

- [64] Zhen Xu, Sida Peng, Chen Geng, Linzhan Mou, Zihan Yan, Jiaming Sun, Hujun Bao, and Xiaowei Zhou. Relightable and animatable neural avatar from sparse-view video. In *arXiv preprint arXiv:2308.07903*, 2023. 2, 7
- [65] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020. 5
- [66] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *NeurIPS*, 2021. 4
- [67] Ye Yu and William A. P. Smith. Inverserendernet: Learning single image inverse rendering. In *CVPR*, 2019. 2
- [68] Zhengming Yu, Wei Cheng, xian Liu, Wayne Wu, and Kwan-Yee Lin. MonoHuman: Animatable human neural field from monocular video. In *CVPR*, 2023. 2
- [69] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, 2021. 2
- [70] Ruiqi Zhang and Jie Chen. Ndf: Neural deformable fields for dynamic human modelling. In *ECCV*, 2022. 4
- [71] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5, 6
- [72] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graphics*, 40(6), 2021. 2
- [73] Ruichen Zheng, Peng Li, Haoqian Wang, and Tao Yu. Learning visibility field for detailed 3d human reconstruction and relighting. In *CVPR*, 2023. 2
- [74] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(6):3170–3184, 2022. 2
- [75] Yihao Zhi, Shenhan Qian, Xinhao Yan, and Shenghua Gao. Dual-space nerf: Learning animatable avatars and scene lighting in separate spaces. In *3DV*, 2022. 2, 5, 6, 7