






Relightable Neural Actor with Intrinsic Decomposition and Pose Control

Diogo Carbonera Luvizon^{1,2}, Vladislav Golyanik¹, Adam Kortylewski^{1,3},
Marc Habermann^{1,2}, and Christian Theobalt^{1,2}

¹ Max Planck Institute for Informatics, Saarland Informatics Campus

² Saarbrücken Research Center for Visual Computing, Interaction and AI

³ University of Freiburg

{dluvizon,golyanik,akortyle,mhaberma,theobalt}@mpi-inf.mpg.de
<https://vcai.mpi-inf.mpg.de/projects/RNA>

Abstract. Creating a controllable and relightable digital avatar from multi-view video with fixed illumination is a very challenging problem since humans are highly articulated, creating pose-dependent appearance effects, and skin as well as clothing require space-varying BRDF modeling. Existing works on creating animatable avatars either do not focus on relighting at all, require controlled illumination setups, or try to recover a relightable avatar from very low cost setups, i.e. a single RGB video, at the cost of severely limited result quality, e.g. shadows not even being modeled. To address this, we propose *Relightable Neural Actor*, a new video-based method for learning a pose-driven neural human model that can be relighted, allows appearance editing, and models pose-dependent effects such as wrinkles and self-shadows. Importantly, for training, our method solely requires a multi-view recording of the human under a known, but static lighting condition. To tackle this challenging problem, we leverage an implicit geometry representation of the actor with a drivable density field that models pose-dependent deformations and derive a dynamic mapping between 3D and UV spaces, where normal, visibility, and materials are effectively encoded. To evaluate our approach in real-world scenarios, we collect a new dataset with four identities recorded under different light conditions, indoors and outdoors, providing the first benchmark of its kind for human relighting, and demonstrating state-of-the-art relighting results for novel human poses.

1 Introduction

Creating a relightable and controllable virtual actor of a real person solely from multi-view videos under static illumination is an exciting and challenging topic that has numerous applications in virtual and augmented reality, gaming, human-machine interaction, and content creation. However, it requires modeling the outgoing radiance of a particular surface point on the human, which depends on the pose-dependent and highly non-rigid surface deformation of the human body, the complex and spatially varying materials, as well as the scene lighting. Intrinsically decomposing these quantities solely from multi-view video and



Fig. 1: Our method learns a neural actor that is driven only by a 3D skeletal pose and allows rendering and editing of the actor’s appearance under new lightning and poses not seen during training. Our approach models pose-dependent deformations, self-shadows, and performs intrinsic decomposition of dynamic humans through a pose-dependent mapping between 3D and UV spaces, which also enables editing the appearance and material properties at inference time. For training, our method only needs a multi-view video and an environment map.

enabling novel compositions that are of highly photorealistic quality defines an unsolved and challenging inverse problem in the literature.

Relighting *static* scenes from multi-view imagery is a well-studied problem [22, 33, 49, 62]. However, directly applying such methods to moving humans, which are dynamic, highly articulated, and have complex material composition, is not trivially possible, as these approaches do not account for dynamics. On the contrary, some works have focused on building drivable and dynamic full-body avatars from multi-view video [32, 42, 54] without considering relightability at all. Despite some works [4, 11, 18, 36, 66] assuming a *well-calibrated lightstage settings*, only few human-specific works [6, 17, 64] try to achieve human intrinsic decomposition and relighting from *videos with static lights*. While they are able to recover a relightable human representation from a single RGB video, they can solely replay the human motion under novel lighting conditions [6], or do not model self-shadows and space-varying BRDF [17], which results in low visual fidelity. In stark contrast, we assume a multi-view setting while accounting for high-frequency surface deformations as well as induced high-frequency lighting effects such as self-shadows and shading, which are indispensable for a truly immersive experience.

To address this, we propose *Relightable Neural Actor*, a new method that achieves intrinsic decomposition of dynamic human avatars, enabling pose-control, material editing, and realistic rendering with pose-dependent effects, while solely requiring a multi-view video of the real human under a single and static light condition during training (see Fig. 1). At inference time, we only need an arbitrary driving skeletal pose and an environment map.

We solve this challenging problem by leveraging an implicit density field from a pose-driven neural actor and by proposing a novel intrinsic decomposition guided by the implicit density field. This strategy allows our method to model pose-dependent deformations in the geometry as dynamic maps, while

representing the albedo and roughness properties of the actor as static maps. Our approach not only makes the problem computationally tractable for the long and dynamic human sequences, but also enables editing the neural actor’s appearance and materials at inference time. Our key contribution is to sample normal and visibility from the posed density field and aggregate them into sparse UV maps, which are then densified by the proposed NormalNet and VisibilityNet models. This allows our approach to effectively model wrinkles and self-shadows even for long sequences. The albedo and roughness maps are optimized during training and sampled by the proposed UVDeltaNet, which accounts for pose-dependent deformations in the UV map and is intrinsically learned in our pipeline. Finally, given an environment map and a virtual camera position, our neural renderer produces the final image of the neural actor.

To train and evaluate our method, we collected a real-world dataset for novel pose and novel lighting human synthesis, composed of four actors recorded under six different light conditions by a multi-view camera setup, including indoor and outdoor environments and a light probe for each sequence, providing the first real in-the-wild benchmark for controllable human relighting.

To summarize, our main technical contributions are as follows: 1) The first method to learn a relightable neural actor from multi-view video under static lights that allows for pose control, appearance editing, and free-viewpoint rendering under arbitrary illumination while modeling pose-dependent effects. 2) We propose an intrinsic decomposition of diffuse and specular reflection based on a new approach to compute pose-dependent normals and visibility by sparsely sampling a neural field and densifying the results with NormalNet and VisibilityNet, making the problem of pose-driven human relighting tractable. Our approach relies on a mapping between 3D and UV space implicitly learned by UVDeltaNet, allowing us to learn an explicit and static material representation that is easy to edit at inference time.

In our experimental evaluation, we compare our approach with two baselines in the proposed real data benchmark, showing the superiority of our approach in synthesizing relightable humans on challenging motion sequences. The effectiveness of our method is demonstrated by ablation studies and qualitative results.

2 Related Work

Since the goal of this work is to build a *full-body* relightable and animatable avatar, we will not discuss works [9, 10, 28, 44, 46, 47, 50, 56, 58, 59] focusing on face relighting as both settings have different requirements and challenges.

2.1 Relighting Static Scenes

Classic approaches for relighting and editing perform intrinsic decomposition through inverse rendering [43]. With recent advances in neural rendering [51], light transportation can be taken into account with fully differentiable rendering

pipelines. For instance, Zhang et al. [62] address the problem of recovering geometry and spatially-varying BRDF with an implicit scene representation. However, this requires solving the light transport by evaluating the NeRF model for each light ray, which is very inefficient. NeRV [49] and Zhang et al. [63] mitigate this problem with a neural visibility field that is jointly optimized with the NeRF model, and NeRD [5] adds further relaxation by ignoring cast shadows. However, this assumption cannot be made for humans, which are highly articulated and often experience prominent self-casting shadows. Jin et al. [22] propose TensorIR, an efficient tensor-based representation that solves the inverse problem by casting camera and secondary rays. IRON [60] proposes a two-stage process for obtaining a mesh with material properties, and Lyu et al. [35] and NeRO [33] focus on multi-bounce illumination and reflective objects, which is usually not the case for humans. The methods mentioned above share the assumption of a static scene and cannot be trivially adapted to dynamic and controllable humans.

2.2 Human Models with Pose Control

Modeling humans is a long-standing topic. Traditionally, meshes from an actor can be automatically rigged to a skeleton [7] using Linear Blending Skinning (LBS) [27]. Parametric human models offer pose and shape control [34, 41] and, more recently, implicit human bodies can be articulated [3, 37, 38]. Such models offer a generic and controllable representation, but are unable to represent individual details. Template-based methods synthesize a realistic level of details [14, 15, 21], but have to deal with non-linear pose-dependent deformations.

Several approaches model appearance and geometry with an implicit field anchored on a coarse human mesh. Neural Body [42] optimizes latent vectors anchored in the vertices of the human mesh and Neural Actor [32] incorporates texture maps as a conditioning to break down the mapping from pose to dynamic effects (one-to-many mapping). ARAH [54] models the human surface with an SDF, Dual-Space NeRF [65] learns the color and light in separate spaces, and Kwon et al. [24] propose an approach that generalizes to new identities and poses. Shyshey et al. [48] jointly learn textures and a neural renderer but fall short in reconstructing the 3D human body. A few works [20, 57] approach a restrictive scenario of reconstructing 3D humans from a monocular video but fail to model fine details under new poses. Common to all these type of methods [12, 13, 25, 40] is the entanglement of material properties and illumination, making them restricted to the same appearance as observed during training.

2.3 Relighting Human Models

Only a few methods in the literature tackle the problem of relighting humans from videos under static illumination. Li et al. [29] create a relightable human mesh by estimating a low-frequency environment using Spherical Harmonics. Some approaches target the problem of single-image [19] and static pose [64] relighting under unknown lights, which is an extremely ill-posed problem due to the high ambiguity between albedo and lighting. Relighting4D [6] extends

Neural Body [42] to estimate geometry and reflectance, but can only replay the same sequence observed during training under new lights and is not able to generate new poses. RANA [17] allows for pose control but assumes Lambertian reflectance, ignoring different material properties such as skin and clothes, and does not model cast shadows, resulting in unrealistic renderings for articulated humans. In contrast, our method performs intrinsic decomposition with space-varying BRDF, models illumination and self-shadows, allows for pose control and novel lights, and offers an easy manner to edit the material properties of the learned neural actor.

3 Method

Our goal is to learn a digital human avatar (neural actor) from RGB images that models pose-dependent effects and can be relighted, pose-controlled, and rendered from novel viewpoints at test time. To achieve relighting, we decompose the observed appearance into normals, visibility with respect to each light source, and BRDF represented as albedo and roughness. With these components and an environment map, we employ the microfacet model [52] to render the output image of our method, which is compared to the observed images. In addition—since we target a drivable actor—the geometry components need to be controllable by the skeletal pose.

Fig. 2 shows an overview of our method divided into three main components. The first is a pose-driven implicit geometry model, where the goal is to obtain a skeleton-controllable implicit 3D density field that models the geometry of the actor with pose-dependent deformations. This model is adapted from Neural Actor [32] and is depicted here since our approach builds upon its internal design. Note that Neural Actor does not perform intrinsic decomposition and is unable to perform relighting or appearance editing, which is our main goal. The second component is our proposed intrinsic decomposition approach whose goal is to disentangle geometry represented as normal and visibility information, and material properties as albedo and roughness. Our key insight is to represent these components in UV space rather than in 3D coordinates, using the density field as guidance in the mapping function between 3D and UV spaces. This allows our method to model complex deformations that are pose-dependent in a compact UV map that can be easily computed and queried. Finally, we integrate the microfacet model [52] into our rendering pipeline that takes as input the disentangled components (i.e., normal, visibility, albedo, and roughness) and an environment map, and renders the image of the neural actor in an arbitrary pose and light condition from a free-viewpoint.

3.1 Preliminaries

Human Pose and Environment Map. Our method takes as input a 3D articulated human skeleton pose $\mathbf{P}_t = \{\boldsymbol{\theta}_t, \mathbf{R}_t, \mathbf{T}_t\}$, defined by the relative joint rotation $\boldsymbol{\theta} \in \mathbb{R}^{J \times 3}$ of J body joints, the global rotation $\mathbf{R} \in \mathbb{R}^{3 \times 3}$, and global

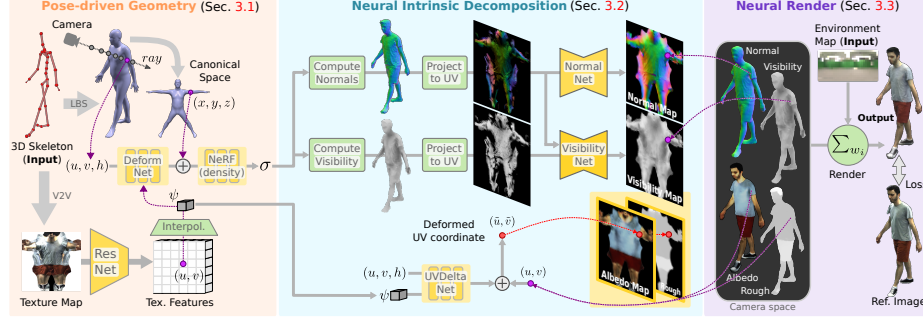


Fig. 2: Our method takes as input a 3D skeletal pose and a static environment map, and renders the neural actor from a virtual camera position. The pose-driven geometry (Sec. 3.1) learns an implicit density function, from which we compute the normal and visibility information. Our intrinsic decomposition disentangles normal, visibility, albedo and roughness maps in UV space (Sec. 3.2). The neural renderer (Sec. 3.3) outputs our prediction, which is supervised with the reference image. **Yellow** represents learnable components and **green** non-learnable components.

translation $\mathbf{T} \in \mathbb{R}^3$ of the person at frame t . The 3D pose can be used to drive a coarse human mesh with linear blending skinning (LBS) [34]. We represent the global lights as an environment map $\mathbf{L} \in \mathbb{R}^{16 \times 32 \times 3}$, which can be obtained from HDR images of a mirror sphere [8]—see the supplement for more details. The environment map can be static during training and changed during inference to render results under novel illumination.

Training Data. During training, our method is supervised with multi-view calibrated video streams $\mathbf{V}_c = \{\mathbf{I}_{c,t}\}$, where $c \in \{1, \dots, C\}$ is the camera index and $\mathbf{I}_{c,t}$ is the image from camera c at frame t . We obtain a foreground mask $\mathbf{M}_{c,t}$ for each video frame [30]. The coarse human mesh and 3D skeletons can be obtained with multi-view pose tracking [1]. The input 3D skeletons drive the LBS model to obtain posed human meshes used to extract texture maps from the multi-view images [2]. These texture maps are used during training to learn pose-dependent deformations on the implicit human model. At inference time, a V2V [55] network predicts texture maps from posed normal maps [32].

Implicit Human Model. For realistically relighting articulated avatars, we need to model human geometry with deformations caused by dynamic and varied poses. Therefore, we follow Neural Actor [32] to represent the pose-driven geometry as an implicit field guided by the 3D skeletal pose parameters and conditioned by texture maps, which provide the ability to model high-frequency deformations. Specifically, a standard NeRF model [39] is optimized through volumetric rendering in the canonical space of the coarse human mesh and the canonical space is deformed based on ResNet [16] features extracted from the

texture maps, which provide the ability to represent detailed wrinkles and deformations not modeled by the coarse mesh. During optimization, from each pixel in the reference images we cast a ray $\mathbf{r}_n = \mathbf{o} + \delta_n \mathbf{d}$, where $n = 1, \dots, N$ is the index of each point in the ray, $\mathbf{o} \in \mathbb{R}^3$ is the origin defined in the optical center of the camera, $\mathbf{d} \in \mathbb{R}^3$ is the ray direction, and δ_n is the depth (distance to the camera). The sampled points \mathbf{r}_n are projected to the mesh surface, resulting in the local coordinates (u_n, v_n, h_n) , where (u, v) is the surface coordinate in the UV space and h is the signed distance to the mesh.

3.2 Neural Intrinsic Decomposition

At the core of our method, we perform intrinsic decomposition to obtain the normal, visibility, albedo, and roughness components, required for rendering the neural actor under new light conditions. Specifically, for each point in each ray cast into the scene we want to obtain the normal direction $\mathbf{n} \in \mathbb{R}^3$, the visibility coefficients $\mathbf{v} \in \mathbb{R}^{512}$ respective to each light source in the environment map, the RGB albedo $\mathbf{a} \in \mathbb{R}^3$, and the scalar roughness ρ . The normal and visibility values can be computed from the geometry model. However, computing the visibility of each point in 3D with respect to the light sources is computationally intractable for a dynamic and long sequence, since for every point in 3D and for every frame, several rays have to be cast towards each pixel of the environment map. Therefore, we propose a different approach.

We notice that the intrinsic decomposition is only relevant for points close to the surface. Since the coarse human mesh provides a mapping from the mesh surface to UV space and the implicit human model provides a density field that models the geometry of the actor, we propose an aggregation strategy that samples normal and visibility from a given point of view (therefore, only a few points are sampled in 3D) and aggregates the obtained values in a per-frame UV map. Once we obtain the normal and visibility in UV space for a given pose, we can simply query the values by projecting the sampled 3D point \mathbf{r}_n into the UV space, preventing our method from the expensive computation of normal and visibility at every possible 3D coordinate over multiple frames. We next explain this approach in more detail.

Normal Maps. For each point in a ray, we want to obtain the normal value that will be required for solving the rendering equation during relighting (will be introduced in Sec. 3.3). This value could be obtained by computing the gradient from the neural field [62]. However, this process requires back-propagating gradients throughout the entire network and often results in noisy normals. Instead, we explicitly compute the normal vector close to the visible surface by deriving the unbiased depth $\hat{\mathbf{D}}$ [53]:

$$\hat{\mathbf{D}}(\mathbf{r}) = \frac{\sum_{n=1}^N w_n \delta_n}{\sum_{n=1}^N w_n}, \quad \text{with } w_n = T_n (1 - \exp(-\sigma_n \bar{\delta}_n)), \quad (1)$$

where $T_n = \exp\left(-\sum_{j=1}^{n-1} \sigma_j \bar{\delta}_j\right)$, w_n is the density weight from volumetric rendering [39], and σ_n is the density predicted by the NeRF model at point \mathbf{r}_n .

Then, as illustrated in Fig. 3, we convert the depth values relative to the camera to 3D coordinates and obtain the normal direction from the tangent surface computed at the intersection point, defined as the maximum value of the density weight w_n . This process can be performed for multiple viewpoints to cover the surface of the human.

Once the normal values are obtained in 3D, we project them to the UV map and aggregate them with the obtained normals from different views. In this process, the (u^*, v^*) coordinate where the obtained normal will be projected onto the UV map is defined by

$$(u^*, v^*) = \Psi(\mathbf{r}_n^*), \quad \text{with } \mathbf{r}_n^* = \arg \max_{\mathbf{r}_n} w_n, \quad (2)$$

where Ψ is a function that projects a 3D point to the closest surface point in the human mesh. Note that w_n depends on the sampled points \mathbf{r}_n (from T_n in Eq. 1) and (u^*, v^*) corresponds to the (u_n, v_n) value obtained from the point \mathbf{r}_n on the ray that has the largest density weight w_n . This process ensures that the normal map obtained aggregates the normals most likely close to the surface.

In this process, the resulting UV map can be sparse, which is undesirable for free-viewpoint rendering. Thus, the NormalNet is used to inpaint the missing normals in UV space. This lightweight CNN is implemented with 8 layers of partial convolutions [31], takes as input sparse normal maps with size $512 \times 512 \times 3$, and predicts a dense normal map (see Fig. 2). During training, NormalNet is supervised with densely sampled normal maps from the training data.

Visibility Maps. The visibility information is necessary to model pose-dependent shadows in the neural avatar, without which it is not possible to produce realistic relighting results. A naive approach for obtaining the visibility for each point in 3D would require querying the neural field thousands of times while casting rays from the intersecting point to all the light sources, which is prohibitive during training. In contrast, in our approach we efficiently encode visibility of each point in 3D close to the surface by applying an approximate mapping to UV space. Specifically, we pre-compute the visibility maps for each pose by sampling multiple camera views and aggregate the results in UV map following Eq. (2). The resulting visibility maps with size $256 \times 256 \times 512$ encode the visibility information w.r.t. each light source in the environment map (represented by 512 individual lights). Similarly to the computation of normal maps, this process results in sparse visibility maps, which are inpainted by the VisibilityNet (as




Fig. 3: The unbiased depth from Eq. (1) is projected into 3D coordinates, where we approximate the tangent surface at the intersecting point and obtain the normal vector. Normals are sampled from multiple viewpoints and aggregated in the UV map.

shown in Fig. 2). This shallow CNN takes as input sparse visibility and normal maps, as the visibility highly depends on normals (e.g. the visibility is always zero for the backlight source), and is trained with densely sampled visibility from the training sequence. See the supplement for architecture details.

Albedo and Roughness Maps. As discussed above, the normal and visibility components can be computed from the density field by using UV maps as proxies. However, the albedo and roughness properties cannot be directly obtained from the pre-trained Neural Actor model, since it predicts radiance which is the product of the rendering equation. Instead, we want to learn a decomposed representation. In intrinsic decomposition, the material properties are commonly optimized via differentiable rendering. Hence, we optimize the albedo and roughness of the neural actor by leveraging a differentiable renderer (Sec. 3.3) and by imposing priors through our design.

The material properties of an actor should not depend on the pose and should be efficiently encoded to enable fast computations and editing. Therefore, we model albedo and roughness as static UV maps, respectively defined as $\mathbf{A} \in \mathbb{R}^{1024 \times 1024 \times 3}$ and $\mathbf{B} \in \mathbb{R}^{1024 \times 1024}$. These tensors are jointly optimized for the whole training sequence and regularized with an L1 loss on the spatial gradients, which encourages smoothness in both maps:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{N} \sum |\nabla_x(\mathbf{A})| + |\nabla_y(\mathbf{A})| + |\nabla_x(\mathbf{B})| + |\nabla_y(\mathbf{B})|, \quad (3)$$

where ∇_x and ∇_y represent the spatial gradient operators.

Nonetheless, by assuming a material representation that does not depend on the skeletal pose, we still have to cope with intrinsic mesh registration errors. Such registration errors could cause blurry results or artifacts in the final renderings. To cope with this problem, we propose a corrective term to the UV mapping, which is predicted by the UVDeltaNet, as shown in Fig. 2. The UVDeltaNet is an MLP that takes as input the local coordinates (u, v, h) for each point on the ray and the sampled texture features ψ , and outputs a corrective term $\epsilon = (u', v')$. The texture features are necessary to provide additional information, since the local coordinates alone cannot account for deviations in the mesh registration.

Initially, UVDeltaNet is frozen and outputs zeros. After the initial convergence of \mathbf{A} and \mathbf{B} , UVDeltaNet is indirectly supervised by the rendering loss and regularized with:

$$\mathcal{L}_{\text{uv}} = \frac{1}{N} \sum \epsilon^2 + |\nabla_x(\epsilon)| + |\nabla_y(\epsilon)|, \quad (4)$$

which minimizes the magnitude of the deformation and encourages its spatial gradient to be sparse, i.e., it makes the UV deformation spatially coherent.

3.3 Neural Renderer

The third part of our method is our neural renderer, which takes the disentangled components as normal, visibility, albedo, and roughness; integrates the lights from the environment map, and outputs the rendered image.

Our neural renderer starts by recasting the ray \mathbf{r} with importance sampling around the estimated depth value $\hat{\mathbf{D}}$. Thus, all the points \mathbf{r}_n are most likely close to the surface. For each sampled point in the ray, we obtain the normal direction \mathbf{n}_n and the visibility values \mathbf{v}_n by interpolating the normal and visibility maps in UV around the point $(u_n, v_n) = \Psi(\mathbf{r}_n)$. The albedo \mathbf{a}_n and roughness ρ_n are obtained by interpolating the albedo and roughness maps in UV around the point $(u_n + u'_u, v_n + v'_v)$, i.e., following the deformation predicted by UVDeltaNet.

Before rendering the final image of the posed neural actor, we consider the light transport equation defined as

$$\mathbf{L}_0(\mathbf{r}_n, \boldsymbol{\omega}_0) = \sum_i \mathbf{R}(\mathbf{a}_n, \rho_n, \boldsymbol{\omega}_i, \boldsymbol{\omega}_0) \mathbf{v}_{n,i} \mathbf{L}(\boldsymbol{\omega}_i)(\boldsymbol{\omega}_i \mathbf{n}_n), \quad (5)$$

where $\mathbf{R}(\cdot)$ is the microfacet BRDF model [52], $\boldsymbol{\omega}_0$ and $\boldsymbol{\omega}_i$ are the camera and the light source directions, respectively, $\mathbf{v}_{n,i}$ is the visibility of point \mathbf{r}_n w.r.t. the i th light source, and $\mathbf{L}(\boldsymbol{\omega}_i)$ is the incoming light from the environment map. Eq. (5) is defined for a single point \mathbf{r}_n on the ray and is integrated into our volume rendering as

$$\tilde{\mathbf{I}}(\mathbf{r}) = \sum_{n=1}^N w_n \mathbf{L}_0(\mathbf{r}_n, \boldsymbol{\omega}_0), \quad (6)$$

where $\tilde{\mathbf{I}}$ is the output image of the relighted neural actor.

During training, our relightable model is supervised with the following loss:

$$\mathcal{L}_{\text{relight}} = \lambda_{\text{L2}} \mathcal{L}_{\text{L2}} + \lambda_{\text{vgg}} \mathcal{L}_{\text{vgg}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \mathcal{L}_{\text{uv}}, \quad (7)$$

where \mathcal{L}_{L2} and \mathcal{L}_{vgg} are the L2 and the perceptual losses [26] between $\tilde{\mathbf{I}}(\mathbf{r})$ and the reference image. The hyperparameters are defined as $\lambda_{\text{L2}} = 100$, $\lambda_{\text{vgg}} = 0.01$, and $\lambda_{\text{smooth}} = 0.1$ in all our experiments.

4 Relightable Dynamic Actors Dataset

Existing datasets for human reenactment and relighting are often monocular [2, 17], contain very short sequences with only a few pose variations, and do not provide environment maps [2, 17, 32, 42]. Hence, we cannot use them for evaluation under realistic scenarios with real ground truth. In addition, we observe that the only existing synthetic dataset for human relighting and pose control [17] suffers from artifacts in the geometry and salt-and-pepper noise in the renderings.

With these shortcomings in mind, we collect a new real dataset, which we call *Relightable Dynamic Actors*, with four identities performing a series of ten different actions (Fig. 4). Each subject is recorded six times wearing the same clothes but under different light conditions, including indoor and outdoor, where we used around 100 and eight cameras, respectively. For each sequence, we also capture HDR light probes, which can be used for training or testing. In total, our dataset contains around 90k multi-view video frames divided into 24 multi-view video sequences. We refer to the supplemental document for more details.



Fig. 4: *Relightable Dynamic Actors* dataset. For each actor, we recorded six sequences under different light conditions, including five indoor sequences and one outdoor. Environment maps were obtained from HDR pictures of a mirror sphere.

5 Experiments

In this section, we provide an experimental evaluation of our method on the Relightable Dynamic Actors dataset and show qualitative and quantitative results. Please refer to the supplementary video for dynamic results on long sequences.

5.1 Evaluation

Metrics. For quantitative evaluation, we report results on the PSNR, SSIM and LPIPS metrics [61], comparing the rendered relight results with the ground-truth images from our dataset. Importantly, our work is the first to be able to report such metrics on a real dataset since no previous dataset in the literature provided a real human recorded under different and known light conditions.

Baselines. Since the only method in the literature capable of reposing a human under new light conditions does not provide implementation [17], and Relighting4D [6] is not able to perform inference for new poses, we defined two baselines for a fair comparison. Our baselines are derived from Neural Actor [32] and adapted to perform intrinsic decomposition based on the obtained normal and visibility information, as discussed in Sec. 3.2. However, instead of using NormalNet and VisibilityNet to obtain these components, we densely sample the neural field and perform morphological inpainting to obtain dense UV maps. The first baseline, referred to as *Neural Actor+IR*, predicts the albedo and roughness components instead of RGB color directly from the deformed canonical space $\mathbf{x}_n = (x, y, z)$ (see Fig. 2). The second baseline, *Neural Actor+IR+Tex*, takes as input the texture features ψ as a conditioning for predicting the albedo and roughness values. In both cases, the predictor is a NeRF MLP with 8 layers.

Implementation Details. Our model is trained in a three-stage process. First, we train the geometry component (Sec. 3.1) with the Adam [23] optimizer with a learning rate of 0.001 until convergence, which takes about two days on four NVidia A40 GPUs. Then, we pre-train the NormalNet and VisibilityNet with pre-computed normal and visibility UV maps sampled from the geometry model



Fig. 5: Qualitative results of our method on *novel poses*. Our approach generalizes to both indoor and outdoor illumination, and generates realistic renderings for arbitrary environment maps with detailed shadows and pose-dependent deformations.

from 40 different views. Both networks are trained in parallel using Adam with a learning rate of 0.01 for about 12 hours. Finally, we train our intrinsic decomposition model with the same optimizer until convergence, which takes about one day. Since the obtained environment maps are up to factor of scale [8], we apply global color correction during training and evaluation. The predicted texture maps are obtained with a V2V model [55] trained (in parallel to our model) to translate a sequence of SMPL normal maps into a sequence of texture maps. Please refer to our supplementary materials for architecture details.

5.2 Results

Qualitative Results. In Fig. 5, we show qualitative results of our method. Our results are all from a trained model driven by a *new skeletal pose*, *new camera viewpoint*, and *new environment map* not seen during training. Note how our approach produces realistic appearance and prominent cloth wrinkles, which can be easily observed under different light conditions. Importantly, notice how our results look realistic under novel lights, even though our model is trained on a single sequence under a static light condition.

Table 1: Quantitative results from our method on real data, subject S1, considering indoor and outdoor scenes. Our method consistently improves the baselines, especially in outdoor scenarios, where the gap between training and test lights is more pronounced.

Method	T2			T3			T4			T5			T6 (outdoors)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Neural Actor+IR	17.522	0.750	0.181	17.473	0.768	0.178	17.745	0.767	0.171	18.244	0.768	0.175	18.355	0.710	0.203
Neural Actor+IR+Tex	18.175	0.778	0.164	17.799	0.790	0.167	18.132	0.785	0.167	18.531	0.789	0.166	18.704	0.737	0.186
Ours	18.602	0.792	0.163	18.240	0.800	0.169	18.388	0.800	0.164	18.820	0.800	0.165	19.461	0.753	0.173

In Fig. 6, we show results from our intrinsic decomposition. Note that our normal representation nicely captures the geometry details and the visibility. This is important to obtain an albedo that does not contain baked-in lighting effects, e.g. shadows, as can be observed from our results.



Fig. 6: Results from our intrinsic decomposition. Our method provides detailed normals and visibility (averaged for visualization), and plausible albedo and roughness predictions that enable realistic relighting.

Another important and unique characteristic of our method is the ability to model self-shadows and cloth wrinkles. In Fig. 7, we show results from two subjects, where a “sunrise” light probe was used to produce cast shadows. Note how pose-dependent effects are consistent with the mostly directional illumination.

Comparisons. We evaluated our method on the test sequences from our dataset. Specifically, we compare against two baselines in the indoor sequences (T2-T5) and outdoor sequence (T6) of subject S1 in every 10th frame, which sums up to 4000 evaluation frames. In Tab. 1 we show, for the first time, quantitative results of *real human relighting* under new poses and new light conditions. Our approach consistently improves upon the two baselines.

Ablation. In Tab. 2, we show an ablation with different components of our method. Compared to our baseline (Neural Actor+IR), learning the albedo and roughness as UV maps guided by the neural field significantly improves the performance on relighting and reposing. This is because, in the baselines, the albedo and roughness are under-constrained and can possibly embed shading effects, overfitting to the training data. The NormalNet and VisibilityNet provide additional capacity to the model to learn how to inpaint the sampled UV maps. Note that without these two models our results are computed using morphological region growth in the UV maps. Finally, the UVDeltaNet provides more flexibility to the albedo and roughness representation, accounting for fine details. This can be observed in the improvement in the LPIPS metric in the last



Fig. 7: Our method models pose-dependent effects and cast shadows, as can be observed from our results under the “sunrise” light.

row of Tab. 2. Nonetheless, this also means more capacity to overfit the training sequences, which translates into a small decrease in the PSNR metric.

5.3 Applications

An important characteristic of our approach is our static representation in UV for the albedo and roughness. This allows exciting and new applications such as color and material editing of an animatable and relightable neural actor, as shown in Fig. 8 and in Fig. 1. With a simple edit in the obtained albedo or roughness 2D map, we can insert new textures or new materials and render the edited neural actor under new poses, new viewpoints, and new light conditions.

Table 2: Ablation study with different components of our method. The proposed intrinsic decomposition with static albedo and roughness maps significantly improves over the baseline, and the proposed NormalNet and VisibilityNet further improve the results while enabling free-viewpoint rendering. UVDeltaNet provides more realistic results (LPIPS metric) with a small sacrifice in the PSNR.

Method	Indoors (S1,T2)			Outdoors (S1,T6)		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Baseline	17.522	0.750	0.181	18.355	0.710	0.203
+Albedo/Rough. Maps	18.334	0.789	0.168	19.341	0.747	0.181
+Normal/Vis.Net	18.634	0.792	0.173	19.704	0.760	0.180
+UVDeltaNet (Ours)	18.602	0.792	0.163	19.461	0.753	0.173

6 Limitations

This work focuses on visual fidelity, and representations faster than NeRF could be tested in our framework in future. Moreover, like several previous techniques for the free-viewpoint rendering of humans and avatar generation [32,45], the proposed method can currently only deal with tight clothing due to the coarse geometry registration. The coarse mesh registration can also negatively affect the fine details in the face and hand synthesis.



Fig. 8: Our method allows editing the material properties and simultaneously performing skeletal pose control, free-viewpoint rendering, and relighting of the edited neural actor.

7 Conclusion

We conclude that the proposed method for the free-viewpoint rendering of humans enables simultaneous relighting, material editing, and pose control at test time while requiring a single multi-view video sequence with static illumination during training. Our method demonstrates functionality that no competing method making the same assumptions at training time (i.e. from multi-view RGB video inputs only) can provide. We hope that our method and the proposed dataset will provoke further research in this exciting field, bringing the neural rendering of humans to a new qualitative level and paving the way for next-generation metaverse applications.

Acknowledgements

This work was funded by the ERC Consolidator Grant 4DRepLy (770784), the Saarbrücken Research Center for Visual Computing, Interaction and AI, and the German Research Foundation (DFG) under Grant No. 468670075.

References

1. Easymocap - make human motion capture easier. Github (2021), <https://github.com/zju3dv/EasyMocap>
2. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: Computer Vision and Pattern Recognition (CVPR). pp. 8387–8397 (2018)
3. Alldieck, T., Xu, H., Sminchisescu, C.: imghum: Implicit generative models of 3d human shape and articulated pose. In: International Conference on Computer Vision (ICCV). pp. 5461–5470 (2021)
4. Bi, S., Lombardi, S., Saito, S., Simon, T., Wei, S.E., Mcphail, K., Ramamoorthi, R., Sheikh, Y., Saragih, J.: Deep relightable appearance models for animatable faces. *ACM Transactions on Graphics (TOG)* **40**(4), 1–15 (2021)
5. Boss, M., Braun, R., Jampani, V., Barron, J.T., Liu, C., Lensch, H.: Nerd: Neural reflectance decomposition from image collections. In: International Conference on Computer Vision (ICCV). pp. 12684–12694 (2021)
6. Chen, Z., Liu, Z.: Relighting4d: Neural relightable human from videos. In: European Conference on Computer Vision (ECCV) (2022)
7. De Aguiar, E., Theobalt, C., Thrun, S., Seidel, H.P.: Automatic conversion of mesh animations into skeleton-based animations. In: Computer Graphics Forum (EUROGRAPHICS). vol. 27, pp. 389–397 (2008)
8. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: ACM SIGGRAPH 2008 classes, pp. 1–10 (2008)
9. Deng, B., Wang, Y., Wetzstein, G.: Lumigan: Unconditional generation of relightable 3d human faces. In: International Conference on 3D Vision (3DV). pp. 302–312. IEEE (2024)
10. d’Eon, E., Luebke, D., Enderton, E.: Efficient Rendering of Human Skin. In: Kautz, J., Pattanaik, S. (eds.) *Rendering Techniques*. The Eurographics Association (2007). <https://doi.org/10.2312/EGWR/EGSR07/147-157>
11. 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)
12. Habermann, M., Liu, L., Xu, W., Pons-Moll, G., Zollhoefer, M., Theobalt, C.: Hdhumans: A hybrid approach for high-fidelity digital humans. *Proc. ACM Comput. Graph. Interact. Tech.* **6**(3) (aug 2023). <https://doi.org/10.1145/3606927>, <https://doi.org/10.1145/3606927>
13. Habermann, M., Liu, L., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Real-time deep dynamic characters. *ACM Transactions on Graphics* **40**(4) (aug 2021)
14. Habermann, M., Xu, W., Zollhoefer, M., Pons-Moll, G., Theobalt, C.: Livecap: Real-time human performance capture from monocular video. *ACM Transactions On Graphics (TOG)* **38**(2), 1–17 (2019)

15. Habermann, M., Xu, W., Zollhofer, M., Pons-Moll, G., Theobalt, C.: Deepcap: Monocular human performance capture using weak supervision. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 5052–5063 (2020)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
17. Iqbal, U., Caliskan, A., Nagano, K., Khamis, S., Molchanov, P., Kautz, J.: Rana: Relightable articulated neural avatars. In: *International Conference on Computer Vision (ICCV)*. pp. 23142–23153 (2023)
18. Iwase, S., Saito, S., Simon, T., Lombardi, S., Bagautdinov, T., Joshi, R., Prada, F., Shiratori, T., Sheikh, Y., Saragih, J.: Relightablehands: Efficient neural relighting of articulated hand models. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 16663–16673 (2023)
19. Ji, C., Yu, T., Guo, K., Liu, J., Liu, Y.: Geometry-aware single-image full-body human relighting. In: *European Conference on Computer Vision (ECCV)*. pp. 388–405. Springer (2022)
20. Jiang, W., Yi, K.M., Samei, G., Tuzel, O., Ranjan, A.: Neuman: Neural human radiance field from a single video. In: *European Conference on Computer Vision (ECCV)*. pp. 402–418 (2022)
21. Jiang, Y., Habermann, M., Golyanik, V., Theobalt, C.: Hifecap: Monocular high-fidelity and expressive capture of human performances. In: *British Machine Vision Conference (BMVC)* (2022)
22. Jin, H., Liu, I., Xu, P., Zhang, X., Han, S., Bi, S., Zhou, X., Xu, Z., Su, H.: Tensorial inverse rendering. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 165–174 (2023)
23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *CoRR* **abs/1412.6980** (2014), <https://api.semanticscholar.org/CorpusID:6628106>
24. Kwon, Y., Kim, D., Ceylan, D., Fuchs, H.: Neural human performer: Learning generalizable radiance fields for human performance rendering. *Advances in Neural Information Processing Systems (NeurIPS)* **34**, 24741–24752 (2021)
25. Kwon, Y., Liu, L., Fuchs, H., Habermann, M., Theobalt, C.: Deliffas: Deformable light fields for fast avatar synthesis. *Advances in Neural Information Processing Systems* (2023)
26. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 4681–4690 (2017)
27. Lewis, J.P., Cordner, M., Fong, N.: Pose space deformation: A unified approach to shape interpolation and skeleton-driven deformation. In: *SIGGRAPH*. p. 165–172. ACM Press/Addison-Wesley Publishing Co., USA (2000). <https://doi.org/10.1145/344779.344862>, <https://doi.org/10.1145/344779.344862>
28. Li, G., Meka, A., Mueller, F., Buehler, M.C., Hilliges, O., Beeler, T.: Eyenerf: a hybrid representation for photorealistic synthesis, animation and relighting of human eyes. *ACM Transactions on Graphics (TOG)* **41**(4), 1–16 (2022)
29. Li, G., Wu, C., Stoll, C., Liu, Y., Varanasi, K., Dai, Q., Theobalt, C.: Capturing relightable human performances under general uncontrolled illumination. In: *Computer Graphics Forum*. vol. 32, pp. 275–284. Wiley Online Library (2013)
30. Lin, S., Ryabtsev, A., Sengupta, S., Curless, B.L., Seitz, S.M., Kemelmacher-Shlizerman, I.: Real-time high-resolution background matting. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 8762–8771 (2021)

31. Liu, G., Reda, F.A., Shih, K.J., Wang, T.C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: European Conference on Computer Vision (ECCV) (2018)
32. Liu, L., Habermann, M., Rudnev, V., Sarkar, K., Gu, J., Theobalt, C.: Neural actor: Neural free-view synthesis of human actors with pose control. *ACM Transactions on Graphics (TOG)* **40**(6), 1–16 (2021)
33. Liu, Y., Wang, P., Lin, C., Long, X., Wang, J., Liu, L., Komura, T., Wang, W.: Nero: Neural geometry and brdf reconstruction of reflective objects from multiview images. In: SIGGRAPH (2023)
34. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics (SIGGRAPH Asia)* **34**(6), 248:1–248:16 (2015)
35. Lyu, L., Tewari, A., Leimkuehler, T., Habermann, M., Theobalt, C.: Neural radiance transfer fields for relightable novel-view synthesis with global illumination. In: European Conference on Computer Vision (ECCV) (2022)
36. Meka, A., Pandey, R., Häne, C., Orts-Escolano, S., Barnum, P., David-Son, P., Erickson, D., Zhang, Y., Taylor, J., Bouaziz, S., Legendre, C., Ma, W.C., Overbeck, R., Beeler, T., Debevec, P., Izadi, S., Theobalt, C., Rhemann, C., Fanello, S.: Deep relightable textures: Volumetric performance capture with neural rendering. *ACM Transactions on Graphics (TOG)* **39**(6) (nov 2020). <https://doi.org/10.1145/3414685.3417814>, <https://doi.org/10.1145/3414685.3417814>
37. Mihajlovic, M., Saito, S., Bansal, A., Zollhoefer, M., Tang, S.: COAP: Compositional articulated occupancy of people. In: Computer Vision and Pattern Recognition (CVPR) (2022)
38. Mihajlovic, M., Zhang, Y., Black, M.J., Tang, S.: Leap: Learning articulated occupancy of people. In: Computer Vision and Pattern Recognition (CVPR). pp. 10461–10471 (2021)
39. 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: European Conference on Computer Vision (ECCV) (2020)
40. Pang, H., Zhu, H., Kortylewski, A., Theobalt, C., Habermann, M.: Ash: Animatable gaussian splats for efficient and photoreal human rendering. In: Computer Vision and Pattern Recognition (CVPR). pp. 1165–1175 (June 2024)
41. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: Computer Vision and Pattern Recognition (CVPR) (2019)
42. Peng, S., Zhang, Y., Xu, Y., Wang, Q., Shuai, Q., Bao, H., Zhou, X.: Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In: Computer Vision and Pattern Recognition (CVPR) (2021)
43. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: Conference on Computer Graphics and Interactive Techniques. p. 117–128. SIGGRAPH '01, Association for Computing Machinery, New York, NY, USA (2001). <https://doi.org/10.1145/383259.383271>, <https://doi.org/10.1145/383259.383271>
44. Ranjan, A., Yi, K.M., Chang, J.H.R., Tuzel, O.: Facelit: Neural 3d relightable faces. In: Computer Vision and Pattern Recognition (CVPR). pp. 8619–8628 (2023)
45. 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 texel-aligned features. In: ACM SIGGRAPH 2022 Conference Proceedings (2022)

46. Schwartz, G., Wei, S.E., Wang, T.L., Lombardi, S., Simon, T., Saragih, J., Sheikh, Y.: The eyes have it: An integrated eye and face model for photorealistic facial animation. *ACM Transactions on Graphics (TOG)* **39**(4), 91–1 (2020)
47. Shu, Z., Sahasrabudhe, M., Guler, R.A., Samaras, D., Paragios, N., Kokkinos, I.: Deforming autoencoders: Unsupervised disentangling of shape and appearance. In: *European Conference on Computer Vision (ECCV)*. pp. 650–665 (2018)
48. Shysheya, A., Zakharov, E., Aliev, K.A., Bashirov, R., Burkov, E., Iskakov, K., Ivakhnenko, A., Malkov, Y., Pasechnik, I., Ulyanov, D., et al.: Textured neural avatars. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 2387–2397 (2019)
49. 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: *Computer Vision and Pattern Recognition (CVPR)*. pp. 7491–7500 (2021). <https://doi.org/10.1109/CVPR46437.2021.00741>
50. Tan, F., Fanello, S., Meka, A., Orts-Escolano, S., Tang, D., Pandey, R., Taylor, J., Tan, P., Zhang, Y.: Volux-gan: A generative model for 3d face synthesis with hdri relighting. In: *ACM SIGGRAPH 2022 Conference Proceedings*. pp. 1–9 (2022)
51. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Yifan, W., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., Simon, T., Theobalt, C., Nießner, M., Barron, J.T., Wetzstein, G., Zollhöfer, M., Golyanik, V.: Advances in Neural Rendering. *Computer Graphics Forum (EG STAR 2022)* (2022)
52. Walter, B., Marschner, S.R., Li, H., Torrance, K.E.: Microfacet Models for Refraction through Rough Surfaces. In: Kautz, J., Pattanaik, S. (eds.) *Rendering Techniques. The Eurographics Association* (2007). <https://doi.org/10.2312/EGWR/EGSR07/195-206>
53. Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2021)
54. Wang, S., Schwarz, K., Geiger, A., Tang, S.: Arah: Animatable volume rendering of articulated human sdf. In: *European Conference on Computer Vision (ECCV)* (2022)
55. Wang, T.C., Liu, M.Y., Zhu, J.Y., Liu, G., Tao, A., Kautz, J., Catanzaro, B.: Video-to-video synthesis. In: *Conference on Neural Information Processing Systems (NeurIPS)* (2018)
56. Wang, Y., Holynski, A., Zhang, X., Zhang, X.: Sunstage: Portrait reconstruction and relighting using the sun as a light stage. In: *Computer Vision and Pattern Recognition (CVPR)*. pp. 20792–20802 (2023)
57. 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: *Computer Vision and Pattern Recognition (CVPR)*. pp. 16210–16220 (June 2022)
58. Wu, C.h., Zheng, N., Ardisson, S., Bali, R., Belko, D., Brockmeyer, E., Evans, L., Godisart, T., Ha, H., Huang, X., Hypes, A., Koska, T., Krenn, S., Lombardi, S., Luo, X., McPhail, K., Millerschoen, L., Perdoch, M., Pitts, M., Richard, A., Saragih, J., Saragih, J., Shiratori, T., Simon, T., Stewart, M., Trimble, A., Weng, X., Whitewolf, D., Wu, C., Yu, S.I., Sheikh, Y.: Multiface: A dataset for neural face rendering. In: *arXiv* (2022). <https://doi.org/10.48550/ARXIV.2207.11243>, <https://arxiv.org/abs/2207.11243>
59. Yang, H., Zheng, M., Feng, W., Huang, H., Lai, Y.K., Wan, P., Wang, Z., Ma, C.: Towards practical capture of high-fidelity relightable avatars. In: *SIGGRAPH Asia 2023 Conference Papers*. pp. 1–11 (2023)

60. Zhang, K., Luan, F., Li, Z., Snively, N.: IRON: Inverse rendering by optimizing neural sdfs and materials from photometric images. In: Computer Vision and Pattern Recognition (CVPR). pp. 5565–5574 (2022)
61. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Computer Vision and Pattern Recognition (CVPR) (2018)
62. 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**, 1–18 (2021)
63. Zhang, Y., Sun, J., He, X., Fu, H., Jia, R., Zhou, X.: Modeling indirect illumination for inverse rendering. In: Computer Vision and Pattern Recognition (CVPR). pp. 18643–18652 (2022)
64. Zheng, R., Li, P., Wang, H., Yu, T.: Learning visibility field for detailed 3d human reconstruction and relighting. In: Computer Vision and Pattern Recognition (CVPR). pp. 216–226 (2023)
65. Zhi, Y., Qian, S., Yan, X., Gao, S.: Dual-space nerf: Learning animatable avatars and scene lighting in separate spaces. In: International Conference on 3D Vision (3DV) (Sep 2022)
66. Zhou, T., He, K., Wu, D., Xu, T., Zhang, Q., Shao, K., Chen, W., Xu, L., Yu, J.: Relightable neural human assets from multi-view gradient illuminations. In: Computer Vision and Pattern Recognition (CVPR). pp. 4315–4327 (2023)