

# Structured 3D Features for Reconstructing Controllable Avatars

Enric Corona<sup>1,3†</sup> Mihai Zanfir<sup>2†</sup> Thiemo Alldieck<sup>3</sup>  
 Eduard Gabriel Bazavan<sup>3</sup> Andrei Zanfir<sup>3</sup> Cristian Sminchisescu<sup>3</sup>  
<sup>1</sup>UPC, Barcelona <sup>2</sup>Newton <sup>3</sup>Google Research




Figure 1. We introduce *Structured 3D Features* (S3F), a new feature representation for monocular 3D Human Reconstruction. Using S3Fs, we train a new implicit reconstruction model for creating rigged and relightable avatars – without the need for post-processing. S3Fs store features in a dense 3D point cloud around a human body model, allowing for various applications: our model can produce 3D reconstructions in novel poses, under new illumination, under novel views, or with altered clothing for virtual try-on.

## Abstract

We introduce *Structured 3D Features*, a model based on a novel implicit 3D representation that pools pixel-aligned image features onto dense 3D points sampled from a parametric, statistical human mesh surface. The 3D points have associated semantics and can move freely in 3D space. This allows for optimal coverage of the person of interest, beyond just the body shape, which in turn, additionally helps modeling accessories, hair, and loose clothing. Owing to this, we present a complete 3D transformer-based attention framework which, given a single image of a person in an unconstrained pose, generates an animatable 3D reconstruction with albedo and illumination decomposition, as a result of a single end-to-end model, trained semi-supervised, and with no additional postprocessing. We show that our S3F model surpasses the previous state-of-the-art on various tasks, including monocular 3D reconstruction, as well as albedo & shading estimation. Moreover, we show that the proposed methodology allows novel view synthesis, re-lighting, and re-posing the reconstruction, and can naturally be extended to handle multiple input images (e.g. different views of a person, or the same view, in different poses, in video). Finally, we demonstrate the editing capabilities of our model for 3D virtual try-on applications.

† Work was done while Enric and Mihai were with Google Research.

## 1. Introduction

Human digitization is playing a major role in several important applications, including AR/VR, video games, social telepresence, virtual try-on, or the movie industry. Traditionally, 3D virtual avatars have been created using multi-view stereo [33] or expensive equipment [48]. More flexible or low-cost solutions are often template-based [1, 4], but these lack expressiveness for representing details such as hair. Recently, research on implicit representations [6, 16, 22, 52, 53] and neural fields [24, 46, 63, 70] has made significant progress in improving the realism of avatars. The proposed models produce detailed results and are also capable, within limits, to represent loose hair or clothing.

Even though model-free methods yield high-fidelity avatars, they are not suited for downstream tasks such as animation. Furthermore, proficiency is often limited to a certain range of imaged body poses. Aiming at these problems, different efforts have been made to combine parametric models with more flexible implicit representations [21, 22, 62, 71]. These methods support animation [21, 22] or tackle challenging body poses [62, 71]. However, most work relies on 2D pixel-aligned features. This leads to two important problems: (1) First, errors in body pose or camera parameter estimation will result in misalignment between the projections of 3D points on the body surface and image features, which ultimately results in low

quality reconstructions. (2) image features are coupled with the input view and cannot be easily manipulated *e.g.* to edit the reconstructed body pose.

In this paper we introduce Structure 3D Features (S3F), a flexible extension to image features, specifically designed to tackle the previously discussed challenges and to provide more flexibility during and after digitization. S3F store local features on ordered sets of points around the body surface, taking advantage of the geometric body prior. As body models do not usually represent hair or loose clothing, it is difficult to recover accurate body parameters for images in-the-wild. To this end, instead of relying too much on the geometric body prior, our model freely moves 3D body points independently to cover areas that are not well represented by the prior. This process results in our novel S3Fs, and is trained without explicit supervision only using reconstruction losses as signals. Another limitation of prior work is its dependence on 3D scans. We alleviate this dependence by following a mixed strategy: we combine a small collection of 3D scans, typically the only training data considered in previous work, with large-scale monocular in-the-wild image collections. We show that by guiding the training process with a small set of 3D synthetic scans, the method can efficiently learn features that are only available for the scans (*e.g.* albedo), while self-supervision on real images allow the method to generalize better to diverse appearances, clothing types and challenging body poses.

In this paper, we show how the proposed S3Fs are substantially more flexible than current state-of-the-art representations. S3Fs enable us to train a single end-to-end model that, based on an input image and matching body pose parameters, can generate a 3D human reconstruction that is relightable and animatable. Furthermore, our model supports the 3D editing of *e.g.* clothing, without additional post-processing. See Fig.1 for an illustration and Table 1 for a summary of our model’s properties, in relation to prior work. We compare our method with prior work and demonstrate state-of-the-art performance for monocular 3D human reconstruction from challenging, real-world and unconstrained images, and for albedo & shading estimation. We also provide an extensive ablation study to validate our different architectural choices and training setup. Finally, we show how the proposed approach can also be naturally extended to integrate observations across different views or body poses, further increasing reconstruction quality.

## 2. Related work

**Monocular 3D Human Reconstruction** is an inherently ill-posed problem and thus greatly benefits from strong human body priors. The reconstructed 3D shape of a human is often a byproduct of 3D pose estimation [15, 27–29, 31, 41, 44, 51, 54, 69] represented by a statistical human body model [34, 44, 50, 64]. Body models, however, only pro-

	end-to-end trainable	returns albedo	returns shading	true surface normals	challenging poses	animatable	semantic editing	allows multi-view	
✓	✗	✓	✓	✗	✗	✗	✓	✓	PIFu [52]
✗	✗	✗	✓	✗	✗	✗	✗	✗	PIFuHD [53]
✓	✗	✓	✗	✗	✓	✗	✗	✗	ARCH++ [21]
✗	✗	✓	✓	✓	✓	✗	✗	✗	PaMIR [71]
✓	✓	✓	✓	✗	✗	✗	✗	✗	PHORHUM [6]
✗	✗	✗	✓	✓	✓	✗	✗	✗	ICON [62]
✓	✓	✓	✓	✓	✓	✓	✓	✓	S3F (Ours)

Table 1. **Key properties of S3F compared to recent work.** Our model includes a number of desirable novel features with respect to previous state-of-the-art, and recovers rigged and relightable human avatars even for challenging body poses in input images.

vide mid-resolution body meshes that do not capture important elements of a person’s detail, such as clothing or accessories. To this end, one line of work extends parametric bodies to represent clothing through offsets to the body mesh [1–4, 8, 42, 73]. However, this is prone to fail for loose garments or those with a different topology than the human body. Other representations of the clothed human body have been explored, including voxels [58, 72], geometry images [47], bi-planar depth maps [17] or visual hulls [40]. To date, the most powerful representations are implicit functions [6, 10, 21, 22, 52, 53, 62, 66, 67, 71] that define 3D geometry via a decision boundary [13, 36] or level set [43]. A popular choice to condition implicit functions on an input image are pixel-aligned features [52]. This approach has been used to obtain detailed 3D reconstruction methods without relying on a body template [6, 52, 52, 66], and in combination with parametric models [21, 22, 62, 66] to take advantage of body priors. ARCH and ARCH++ map pixel-aligned features to a canonical space of SMPL [34] to support animatable reconstructions. However, this requires almost perfect SMPL estimates, which are hard to obtain for images in-the-wild. To address this, ICON [62] proposes an iterative refinement of the SMPL parameters during 3D reconstruction calculations. ICON requires different modules for predicting front & back normal maps, to compute features, and for SMPL fitting. In contrast, we propose an end-to-end trainable model that can correct potential errors in body pose and shape estimation, without supervision, by freely allocating relevant features around the approximate body geometry.

Somewhat related are also methods for human relighting in monocular images [23, 26, 30, 56]. However, these methods typically do not reconstruct the 3D human and instead rely on normal maps to transform pixel colors. Our work bears similarity with PHORHUM [6], in our prediction of albedo, and a global scene illumination code. However, we combine this approach with mixed supervision from both synthetic and real data, thus obtaining detailed, photo-realistic reconstructions, for images in-the-wild.

**Human neural fields.** Our work is also related to neural radiance fields, or NeRFs [38], from which we take

inspiration for our losses on images in-the-wild. NeRFs have been recently explored for novel human view synthesis [11, 14, 25, 45, 46, 55, 60, 61, 63]. These methods are trained per subject by minimizing residuals between rendered and observed images and are typically based on a parametric body model. See [57] for an extensive literature review. In contrast, our method is trained on images of different subjects and thus generalizes to unseen identities.

### 3. Method

We seek to estimate the textured and animation-ready 3D geometry of a person as viewed in a single image. Further, we compute a per-image lighting model to explain shading effects on 3D geometry and to enable relighting for realistic scene placement.

Our system utilizes the statistical body model GHUM [64], which represents the human body  $M(\cdot)$  as a parametric function of pose  $\theta$  and shape  $\beta$

$$M(\beta, \theta) : \theta \times \beta \mapsto \mathbf{V} \in \mathbb{R}^{3N}. \quad (1)$$

GHUM returns a set of 3D body vertices  $\mathbf{V}$ . We also use imGHUM [5], GHUM’s implicit counterpart. imGHUM computes the signed distance to the body surface  $s_x^{body}$  for any given 3D point  $\mathbf{x}$  under a given pose and shape configuration:  $\text{imGHUM}(\mathbf{x}, \beta, \theta) : \beta, \theta \mapsto s_x^{body}$ . Please refer to the original papers for details.

Given a monocular RGB image  $\mathbf{I}$  (where the person is segmented), together with an approximate 3D geometry represented by GHUM/imGHUM parameters  $\theta$  and  $\beta$ , our method generates an animation-ready, textured 3D reconstruction of that person. We represent the 3D geometry  $\mathcal{S}$  as the zero-level-set of a signed distance field parameterized by the network  $f$ ,

$$\mathcal{S}_{\phi_f} = \{\mathbf{x} \in \mathbb{R}^3 \mid f(\mathbf{I}, \theta, \beta, \mathbf{x}; \phi_f) = (0, \mathbf{a})\}, \quad (2)$$

with learnable parameters  $\phi_f$ . In addition to the signed distance  $s$  w.r.t. the 3D surface,  $f$  predicts per-point albedo color  $\mathbf{a}$ . In the sequel, we denote the signed distance and albedo for point  $\mathbf{x}$  returned by  $f$  as  $s_x$  and  $\mathbf{a}_x$ , respectively.  $\mathcal{S}$  can be extracted from  $f$  by running Marching Cubes [35] in a densely sampled 3D bounding box. In the process of computing  $s$  and  $\mathbf{a}$ ,  $f$  extracts novel *Structured 3D Features*. In contrast to pixel-aligned features – a popular representation used by other state-of-the-art methods – our structured 3D features can be re-posed via Linear Blend Skinning (LBS), thus enabling reconstructions in novel poses as well as multi-image feature aggregation. We explain our novel Structured 3D Features in more detail below.

#### 3.1. S3F architecture

We divide our end-to-end trainable network into two different parts. First, we introduce our novel Structured 3D

Feature (S3F) representation. Second, we describe the procedure to obtain signed distance and color predictions. Finally, we generalize our formulation to multi-view or multi-frame settings.

**Structured 3D Features.** Recent methods have shown the suitability of local features for the 3D human reconstruction, due to their ability to capture high-frequency details. Pixel-aligned features [6, 22, 52, 53] are a popular choice, despite having certain disadvantages: (1) it is not straightforward to integrate pixel-aligned features from different time steps or body poses. (2) a ray-depth ambiguity. In order to address these problems, we propose a natural extension of pixel-alignment by lifting the features to the 3D domain: structured 3D features. We obtain the initial 3D feature locations  $\mathbf{v}_i$  by densely sampling the surface of the approximate geometry represented by GHUM. We initialize the corresponding features  $\mathbf{f}_i$  by projecting and pooling from a 2D feature map obtained from the input image  $\mathbf{I}$

$$g : (\mathbf{I}, \pi(\mathbf{v}_i)) \mapsto \mathbf{f}_i \in \mathbb{R}^k, \quad (3)$$

where  $\pi(\mathbf{v}_i)$  is a perspective projection of  $\mathbf{v}_i$ , and  $g$  is a trainable feature extractor network. The 3D features obtained this way approximate the underlying 3D body but not the actual geometry, especially for loose fitting clothing. To this end, we propose to non-rigidly displace the initial 3D feature locations to better cover the human in the image. Given the initial 3D feature locations  $\mathbf{v}_i$ , we predict per-point displacements in camera coordinates using the network  $d$

$$\mathbf{v}'_i = \mathbf{v}_i + d(\mathbf{f}_i, \mathbf{e}_i, \mathbf{v}_i), \quad (4)$$

where  $\mathbf{e}_i \in \mathbb{R}^k$  is a learnt semantic code for  $\mathbf{v}_i$ , initialized randomly and jointly optimized during training. Finally, we project the updated 3D feature locations again to the feature map obtained from  $\mathbf{I}$  (following the approach from Eq. 3) obtaining our final structured 3D features  $\mathbf{f}'_i$ . See Fig. 2 for an overview.

**Predicting geometry and texture.** We continue describing the process of computing per-point albedo  $\mathbf{a}_x$  and signed distances  $s_x$ . We denote the set of structured 3D features as  $(\mathbf{F}' \in \mathbb{R}^{N \times k}, \mathbf{V}' \in \mathbb{R}^{N \times 3})$ , where  $\mathbf{F}'$  refers to the feature vectors and  $\mathbf{V}'$  to their 3D locations, respectively. First, we use a transformer encoder layer [59] to compute a master feature  $\mathbf{f}_x^*$  for each query point  $\mathbf{x}$ , using  $\mathbf{V}'$  as keys and  $\mathbf{F}'$  values, respectively. From  $\mathbf{f}_x^*$  we compute per-point albedo and signed distances using a standard MLP

$$t : (\mathbf{x}, \mathbf{V}', \mathbf{F}') \mapsto \mathbf{f}_x^* \mapsto s_x, \mathbf{a}_x. \quad (5)$$

While previous work maps query points  $\mathbf{x}$  to a continuous feature map (e.g. pixel-aligned features), we instead aim to pool features from the discrete set  $\mathbf{F}'$ . Intuitively, the most informative features for  $\mathbf{x}$  should be located in its 3D

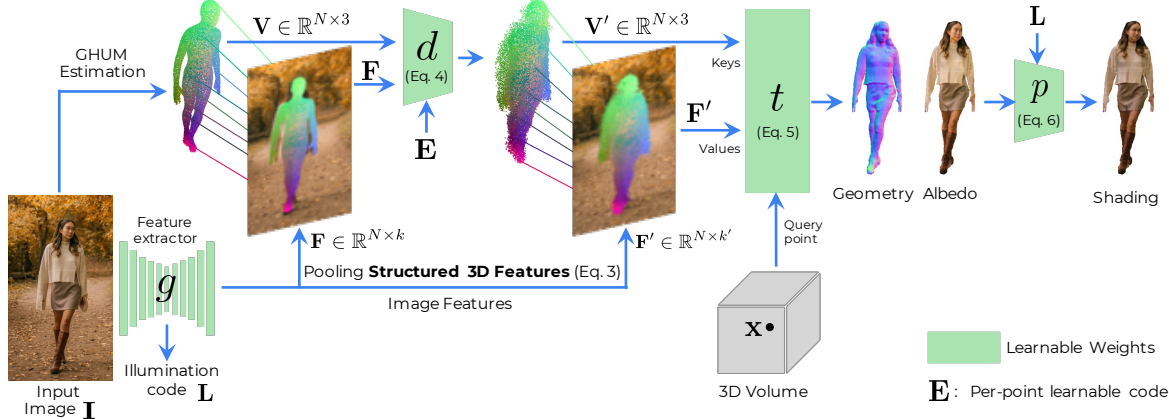


Figure 2. **Method overview.** We introduce a new implicit 3D representation S3F (Structured 3D Features) that utilizes  $N$  points sampled on the body surface  $\mathbf{V}$  and their 2D projection to pool features  $\mathbf{F}$  from a 2D feature map extracted from the input image  $\mathbf{I}$ . The initial body points are non-rigidly displaced by the network  $d$  to obtain  $\mathbf{V}'$  in order to sample new features  $\mathbf{F}'$ , on locations that are not covered by body vertices, such as loose clothing or hair. Given an input point  $x$ , we then aggregate representations from the set of points and their features using a transformer architecture  $t$ , to finally obtain per-point signed distance and albedo color. Finally, we can relight the reconstruction using the predicted albedo, an illumination representation  $L$  and a shading network  $p$ . We assume perspective projection to obtain more natural reconstructions, with correct proportions.

neighborhood. We therefore first explore a simple baseline, by collecting the features of the three closest points of  $x$  in  $\mathbf{V}'$  and interpolating them using barycentric coordinates. However, we found this approach not sufficient (see Tab. 2). With keys and queries based on 3D positions, we argue that a transformer encoder is a useful tool to integrate relevant structured 3D features based on learned distance-based attention. In practice, we use two different transformer heads and MLPs to predict albedo and distance, please see Sup. Mat. for details. Furthermore, we compute signed distances  $s_x$  by updating the initial estimate of imGHUM with a residual,  $s_x = s_x^{body} + \Delta s_x$ . This in practice makes the training more stable for challenging body poses.

To model scene illumination, we follow PHORHUM [6] and utilize the bottleneck of the feature extractor network  $g$  as scene illumination code  $\mathbf{L}$ . We then predict per-point shading coefficient  $\delta_x$  using a surface shading network  $p$

$$p(\mathbf{n}_x, \mathbf{L}; \phi_p) \mapsto \delta_x, \quad (6)$$

parameterized by weights  $\phi_p$ , where the normal of the input point  $\mathbf{n}_x = \nabla_x s_x$  is the estimated surface normal defined by the gradient of the estimated distance w.r.t.  $x$ . The final shaded color is then obtained as  $\mathbf{c}_x = \delta_x \odot \mathbf{a}_x$  with  $\odot$  denoting element-wise multiplication.

**3D Feature Manipulation & Aggregation.** Since the proposed approach relies on 3D feature locations originally sampled from GHUM’s body surface, we can utilize GHUM to re-pose the representation. This has two main benefits: (1) we can reconstruct in a pose different from the one in the image, *e.g.* by resolving self-contact, which is typically not possible after creating a mesh, and (2) we can aggregate information from several observations as follows. In a multi-view or multi-frame setting with  $O$  observations

of the same person under different views or poses, we define each input image as  $\mathbf{I}_t, \forall t \in \{1, \dots, O\}$ . Let us extend the previous notation to  $(\mathbf{V}'_t, \mathbf{F}'_t)$  as structured 3D features for a given frame  $t$ . To integrate features from all images, we use LBS to invert the posing transformation and map them to GHUM’s canonical pose denoted as  $\tilde{\mathbf{V}}'$ . For each observation, we compute point visibilities  $\mathbf{O}_t \in \mathbb{R}^{N \times 1}$  (based on the GHUM mesh) and weight all feature vectors and canonical feature positions by the overall normalized visibility, thus obtaining an aggregated set of structured features

$$\hat{\mathbf{V}}' = \sum_t \text{softmax}_t(\mathbf{O}_t) \odot \tilde{\mathbf{V}}'_t, \quad (7)$$

$$\hat{\mathbf{F}}' = \sum_t \text{softmax}_t(\mathbf{O}_t) \odot \tilde{\mathbf{F}}'_t \quad (8)$$

where  $\text{softmax}$  normalizes per-point contributions along all views. The aggregated structured 3D features  $(\hat{\mathbf{V}}', \hat{\mathbf{F}}')$  can be posed to the original body pose of each input image, or to any new target pose. Finally, we run the remaining part of the model in the posed space to predict  $s_x$  and  $\mathbf{a}_x$ .

### 3.2. Training S3F

The semi-supervised training pipeline leverages a small subset of 3D synthetic scans together with images in-the-wild. During training, we run a forward pass for both inputs and integrate the gradients to run a single backward step.

**Real data.** We use images in-the-wild paired with body pose and shape parameters obtained via fitting. We supervise on the input images through color and occupancy losses. Following recent work on neural rendering of combined signed distance and radiance fields [68], we convert

Method	Geometry			Color (PSNR)	
	Chamfer ↓	IoU ↑	NC ↑	Albedo ↑	Shading ↑
Architecture:					
Using pixel-aligned features	3.59	0.360	0.828	13.17	14.13
+ Predicting SDF residual	5.43	0.428	0.851	13.19	13.06
+ Lifting features to 3D	0.606	0.644	0.765	9.82	10.07
+ Transformer	0.532	0.720	0.911	<b>16.54</b>	15.85
<b>FULL: + Point displacement</b>	<b>0.339</b>	<b>0.734</b>	<b>0.924</b>	16.31	<b>16.67</b>
Supervision regime:					
<b>FULL: Only synthetic data</b>	0.381	0.719	0.916	16.34	15.84
<b>FULL: Only real data</b>	0.444	0.723	0.907	10.96	13.59

Table 2. **Ablation of several of our design choices.** Chamfer metrics are  $\times 10^{-3}$ .

predicted signed distance values  $s_{\mathbf{x}}$  into density such that

$$\sigma(\mathbf{x}) = \beta^{-1} \Psi_{\beta}(-s_{\mathbf{x}}), \quad (9)$$

where  $\Psi_{\beta}(\cdot)$  is the CDF of the Laplace distribution, and  $\beta$  is a learnable scalar parameter equivalent to a sharpness factor. After training, we no longer need  $\beta$  for reconstruction but can use it to render novel views (see Sec. 4.3).

The color of a specific image pixel is estimated via the volume rendering integral by accumulating shaded colors and volume densities along its corresponding camera ray (see [38] or Sup. Mat. for more details). After rendering, we minimize the residual between the ground-truth pixel color  $\mathbf{c}_r$  and the accumulation of shaded colors  $\hat{\mathbf{c}}_r$  along the ray of pixel  $r$  such that  $\mathcal{L}_{\text{rgb}} = \|\mathbf{c}_r - \hat{\mathbf{c}}_r\|_1$ . Additionally, we also define a VGG-loss [12]  $\mathcal{L}_{\text{vgg}}$  over randomly sampled front patches, enforcing the structure to be more realistic. In addition to color, we also supervise geometry by minimizing the difference between the integrated density  $\hat{\sigma}_r$  and the ground truth pixel mask  $\sigma_r$ :  $\mathcal{L}_{\text{mask}} = \|\sigma_r - \hat{\sigma}_r\|_1$ .

Finally, we regularize the geometry using the Eikonal loss [19] on a set of points  $\Omega$  sampled near the body surface such that

$$\mathcal{L}_{\text{eik}} = \sum_{\mathbf{x} \in \Omega} (\|\nabla_{\mathbf{x}} s_{\mathbf{x}}\|_2 - 1)^2. \quad (10)$$

The full loss for the real data is a linear combination of the previous components with weights  $\lambda_*$ :  $\mathcal{L}_{\text{real}} = \mathcal{L}_{\text{rgb}} + \lambda_{\text{vgg}} \mathcal{L}_{\text{vgg}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}} + \lambda_{\text{eik}} \mathcal{L}_{\text{eik}}$ .

**(Quasi-)Synthetic data.** Most previous work relies only on quasi-synthetic data in the form of re-rendered textured 3D human scans, for 3D human reconstruction. While we aim to alleviate the need for synthetic data, it is useful for supervising features not observable in images in-the-wild, such as albedo, color, or geometry of unseen body parts. To this end, we additionally supervise our model on a small set of pairs of 3D scans and synthetic images. For the quasi-synthetic subset, we use the same losses as for the real data and add additional supervision from the 3D scans. Given a 3D scan, we sample  $B$  points on the surface together with its ground truth albedo  $\mathbf{a}_i^{\text{GT}}$  and shaded color  $\mathbf{s}_i^{\text{GT}}$ . We minimize the difference in predicted albedo and shaded color

Method	Geometry			Color (PSNR)	
	Chamfer ↓	IoU ↑	NC ↑	Albedo ↑	Shading ↑
GHUM [64]	3.56	0.562	0.750	-	-
PIFu [52]	6.61	0.519	0.738	-	11.15
Geo-PIFu [20]	9.61	0.453	0.710	-	10.62
PIFuHD [53]	4.94	0.552	0.749	-	-
PaMIR [71]	5.35	0.597	0.763	-	-
ARCH [22]	7.52	0.549	0.712	-	10.66
ARCH++ [21]	6.53	0.549	0.722	-	10.37
ICON [62]	3.53	0.622	0.785	-	-
ICON [62] + [37]	4.47	0.599	0.764	-	-
PHORHUM [6]	2.92	0.594	0.814	<u>11.97</u>	11.20
Ours (No shading)	<u>2.27</u>	<u>0.675</u>	<u>0.827</u>	-	<b>16.62</b>
Ours	<b>1.88</b>	<b>0.694</b>	<b>0.847</b>	<b>15.06</b>	<u>14.81</u>
Ours (GT pose/shape)	0.339	0.734	0.924	16.31	16.67

Table 3. **Quantitative comparison against other monocular 3D human reconstruction methods.** Chamfer metrics are  $\times 10^{-3}$  and PSNR is obtained from all 3D scan vertices, including those not visible in the input image.

such that  $\mathcal{L}_{3\text{D rgb}} = \sum_{\mathbf{x} \in B} \|\mathbf{a}_{\mathbf{x}}^{\text{GT}} - \hat{\mathbf{a}}_{\mathbf{x}}\|_1 + \|\mathbf{c}_{\mathbf{x}}^{\text{GT}} - \hat{\mathbf{c}}_{\mathbf{x}}\|_1$ . We additionally sample a set of 3D points  $\Omega$  close to the scan surface and compute inside/outside labels  $l$  for our final loss,  $\mathcal{L}_{3\text{D label}} = \sum_{\mathbf{x} \in \Omega} \text{BCE}(l_{\mathbf{x}}, \sigma(\mathbf{x}))$ , where BCE is the Binary Cross Entropy function. The overall loss for synthetic data  $\mathcal{L}_{\text{synth}}$  is a linear combination of all losses described above. We train our model by minimizing both real and synthetic objectives  $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{real}} + \mathcal{L}_{\text{synth}}$ . Implementation details are available in the Sup. Mat.

## 4. Experiments

Our goal is to design a single 3D representation and an end-to-end semi-supervised training methodology that is flexible enough for a number of tasks. In this section, we evaluate our model’s performance on the tasks of monocular 3D human reconstruction, albedo, and shading estimation, as well as for novel view rendering and reconstruction from both single-images and video. Finally, we illustrate our model’s capabilities for 3D garment editing.

### 4.1. Data

We follow a mixed supervision approach by combining synthetic data with limited clothing or pose variability, with in-the-wild images of people in challenging poses/scenes.

**Synthetic data.** We use 35 rigged and 45 posed scans from RenderPeople [49] for training. We re-pose the rigged scans to 200 different poses and render them five times with different HDRI backgrounds and lighting using Blender [9]. With probability 0.4 we use a frontal view, otherwise we render the person from a random azimuth and a uniform random elevation in  $[-20, 20]^\circ$ . We obtain ground-truth albedo directly from the scan’s texture map, and shade the full scan (including occluded regions) to obtain ground-truth shaded colors before rendering.

**Images in-the-wild.** We use the HITI dataset [7], a collec-



Figure 3. **Qualitative 3D human reconstruction comparisons against the SOTA method ICON [62]** for complex body poses. Images on the right column feature challenging poses and viewpoints. Notice the additional detail provided by our method, especially for hair and faces. We additionally recover per-vertex albedo (not shown) and shaded color. See Sup. Mat. for more examples and failure cases.

tion of images in-the-wild with 2D key-point annotations and human segmentation masks, to which we fit GHUM by minimizing the 2D joint reprojection error. We automatically remove images with high fitting error, resulting in 40K images for training. We use the provided test set to validate results. More details are provided in the Sup. Mat.

**Evaluation data.** We compare our method for monocular reconstruction using a synthetically generated test set based on a test-split of 3D scans, featuring significant diversity of body poses (*e.g.* running), illumination, and viewpoints. For images in-the-wild, we obtain GHUM parameter initialisation using [18] and further optimize by minimizing 2D joint reprojection errors. For the task of novel-view rendering we use GHS3D [64], which consists of 14 multi-view videos of different people, between 30-60 frames each, with 4 views for training and 2 for testing. We evaluate on test views and compare results based on either a single training image or full video. Note that our method does not “train”, but only computes a forward pass based on a training image.

## 4.2. Monocular 3D Human Reconstruction

**Ablation Study.** We ablate our main methodological choices in Table 2. We use ground-truth pose and shape parameters for the test set, to factor out potential errors in body estimation that otherwise dominate the reconstruction metrics. We provide different baselines to demonstrate the contribution at each step. Albedo and shading evaluation is based on nearest neighbors of the scan, including the back or non-visible regions, which tend to dominate PSNR error rates. In the first two rows we show results of pixel-aligned only baselines. By lifting features to the 3D domain and placing them around the body, our method already achieves competitive performance. This is in line with other works [21, 22, 62] which showed that extending a statistical body model leads to better reconstruction for challenging poses and viewpoints. We improve on this by using a transformer architecture associating query points with 3D features, especially when using our S3F (full model). Finally, we compare our performance when training only on real or synthetic data and show that we get the best results when using a mix of both. Moreover, real data helps generalization



Figure 4. **Output of S3F.** Given an input image (first col.) and GHUM pose estimates (second col.), S3F displaces points efficiently to cover relevant features from the image (third col.). Displaced points cover areas that not necessarily belong to the foreground, but tend to represent loose clothing (*e.g.* hair and loose clothing in first and second examples respectively). This coverage and the efficient feature retrieval allows to correct potential errors in the body such as the feet in the first two examples or the overall body orientation in the third row. Zoom-in recommended.

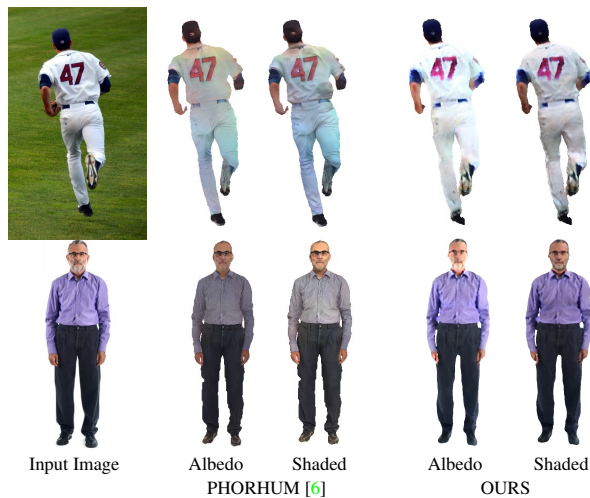


Figure 5. **Qualitative comparison against the state-of-the-art method PHORHUM [6]** for albedo and shading estimation from monocular images.

to in-the-wild images, not captured in our test set.

**SOTA Comparison.** We evaluate our method for monocular 3D human reconstruction and compare with previous methods quantitatively in Table 3. Here, we do not as-

	Chamfer↓	IoU↑	NC↑	PSNR↑	SSIM↑	LPIPS↓	Train time
NeuralBody [46]	0.790	0.887	0.810	24.70	0.829	0.236	hours
H-NeRF [63]	0.218	0.932	0.890	24.92	0.852	0.232	hours
Ours (finetuned)	0.244	0.841	0.910	27.54	0.871	0.148	30 min
One-shot novel view rendering							
Ours (1 image)	0.504	0.795	0.899	23.75	0.825	0.167	0
Ours (video)	0.473	0.807	0.905	24.82	0.836	0.148	0

Table 4. **Quantitative comparison against recent novel view rendering methods for the GHS3D Dataset [63].** We achieve competitive reconstruction, and view rendering metrics using just one forward pass, especially when taking multiple frames as input. This validates our proposed feature integration scheme.

sume ground-truth pose/shape for each image, and instead obtain them from the input images alone, for all methods. We observe that while metrics are dominated by errors in pose estimation, our method is more robust and can correct these errors by using S3Fs. Additionally, our model recovers albedo and shading significantly better than baselines.

Qualitatively, we compare against ICON [62] in Fig. 3, for different levels of body pose complexity. Our method is more robust to both challenging poses and loose clothing than previous SOTA while additionally predicting color, being animatable and relightable. We provide additional comparisons to other methods in Sup. Mat. We also compare our method qualitatively against PHORHUM [6] on the task of albedo and shading estimation from monocular images in Fig. 5. We observe that our method is more robust and we hypothesize this is due to PHORHUM’s trained on synthetic data only, which impacts its ability to generalize to images in-the-wild. We show more results of our method in Fig. 4, where we also visualize the initial and displaced 3D locations of our body points. While the initial GHUM fit might lead to noisy estimates, the displacement step is able to correct errors (feet in first & second rows) and better cover loose clothing (second row).

### 4.3. Applications

The main applications of human digitization (*e.g.* VR/AR or telepresence) require considerable control over virtual avatars, in order to animate, relight and edit them. Here, we evaluate the performance of our approach for novel view rendering to a full video, given as little as one input image. This process entails the ability to re-pose reconstructions and integrate features from different views or video frames. We then showcase the model’s potential for relighting and clothing editing.

**Novel view synthesis.** For novel-view synthesis, previous work often requires tens of GPU hours to train person-specific networks. Here, we compare against two such recent works, NeuralBody [46] and H-Nerf [63], on the GHS3D dataset [63]. We report reconstruction and novel view metrics in Table 4. These show that our method obtains comparable performance without retraining, just from a forward pass. For video, we have enough supervision to

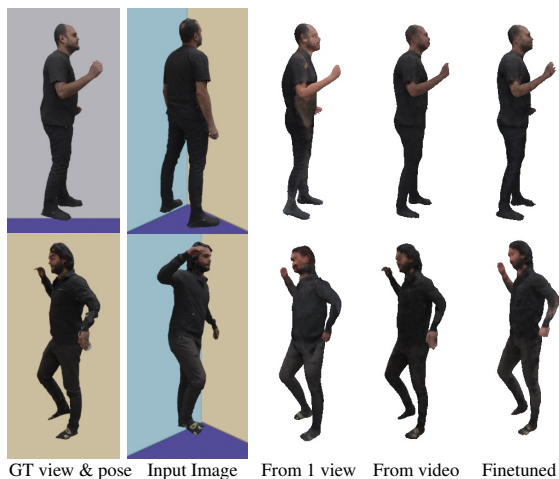


Figure 6. **3D Human reconstruction from one image, video and after finetuning** our network geometry and color heads for 30 minutes. The image used as input in the single-view case is shown in the second column. By using only this information, previously occluded body areas might lead to larger errors when rendering from novel viewpoints or in a different pose, *e.g.* first example, where face and frontal details need to be hallucinated. This can be corrected by integrating information from multiple frames.

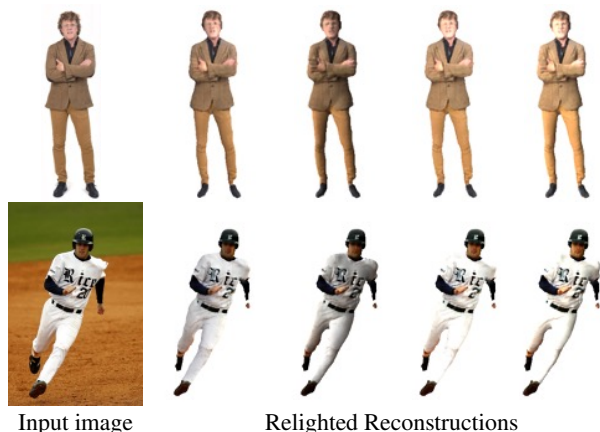


Figure 7. **Qualitative results on 3D Human Relighting.**

robustly fine-tune the geometric head of our network, in just 30 minutes, to obtain a significant performance boost in both geometry and rendering metrics. In Fig. 6 we show the difference between reconstructing from a single view or full video. Even though our method yields a reasonable 3D reconstruction from just a single image, it has to hallucinate color in occluded areas (*e.g.* face and front body parts in the first row). The resulting inconsistency can be largely resolved when novel views are considered.

**Relighting.** By predicting per-point albedo and a scene illumination, our model returns a relightable reconstruction of the person in the target image. We show a number of relighted reconstructions in Fig. 7 given a target input image (left column). The reconstructions are consistently reshaded in novel scenes, thus enabling 3D human compositing applications.

**Clothing editing.** We harvest the semantic properties of the proposed Structured 3D Features to explore potential appli-




Figure 8. **Clothing texture transfer.** Given an image of a target person, we identify S3Fs that project inside upper-body cloth segmentation [32], and replace their feature vectors for those obtained from other subjects. More examples, cloth types and reference images and segmentation masks are shown in Sup. Mat.

cations for 3D virtual try-on. Given an input image of a person and a segmentation mask [32] of a particular piece of clothing (*e.g.* upper-body), we first find body points that project inside the target mask. By exchanging their features with those obtained from other images, we can effectively transfer clothing texture. Fig. 8 showcases diverse clothing editing cases for the upper-body. See Sup. Mat. for other subjects and garments. While 3D clothing transfer is a complex problem by itself [39,65], we tackle a much more challenging scenario than previous methods that typically map color to a known 3D template; and we use a single, versatile end-to-end model for monocular 3D reconstruction, relighting, and 3D editing.

## 5. Discussion

**Limitations.** Failures in monocular human reconstruction occur mainly due to incorrect GHUM fits or very loose clothing. The lack of ambient occlusions in the proposed scheme might lead to incorrect albedo estimates for images in-the-wild. Misalignment in estimated GHUM pose & shape might lead to blurriness, especially for faces. We show failure cases in Supp. Mat.

**Ethical considerations.** We present a human digitization tool that allows relighting and re-posing avatars. However, neither is it intended, nor particularly useful for any form of deep fakes, since the quality of reconstructions is not at par with facial deep fakes. We aim to improve model coverage for diverse subject distributions by means of a weakly supervised approach and by taking advantage of extensive image collections, for which labeled 3D data may be difficult or impossible to collect.

**Conclusion.** We have presented a controllable transformer methodology, S3F, relying on versatile attention-based features, that enables the 3D reconstruction of rigged and relightable avatars by means of a single semi-supervised end-to-end trainable model. Our experimental results illustrate how S3F models obtain state-of-the-art performance for the task of 3D human reconstruction and for albedo and shading estimation. We also show the potential of our models for 3D virtual try-on applications.



## References

- [1] Thiemo Alldieck, Marcus Magnor, Bharat Lal Bhatnagar, Christian Theobalt, and Gerard Pons-Moll. Learning to reconstruct people in clothing from a single rgb camera. In *CVPR*, 2019. 1, 2
- [2] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *3DV*, 2018. 2
- [3] Thiemo Alldieck, Marcus Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, 2018. 2
- [4] Thiemo Alldieck, Gerard Pons-Moll, Christian Theobalt, and Marcus Magnor. Tex2shape: Detailed full human body geometry from a single image. In *ICCV*, 2019. 1, 2
- [5] Thiemo Alldieck, Hongyi Xu, and Cristian Sminchisescu. imGHUM: Implicit generative models of 3D human shape and articulated pose. In *ICCV*, 2021. 3
- [6] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *CVPR*, 2022. 1, 2, 3, 4, 5, 7
- [7] Eduard Gabriel Bazavan, Andrei Zanfir, Mihai Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Hspace: Synthetic parametric humans animated in complex environments. *arXiv*, 2021. 5
- [8] Bharat Lal Bhatnagar, Garvita Tiwari, Christian Theobalt, and Gerard Pons-Moll. Multi-garment net: Learning to dress 3d people from images. In *ICCV*, 2019. 2
- [9] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 2020. 5
- [10] Yukang Cao, Guanying Chen, Kai Han, Wenqi Yang, and Kwan-Yee K. Wong. Jiff: Jointly-aligned implicit face function for high quality single view clothed human reconstruction. In *CVPR*, 2022. 2
- [11] Jianchuan Chen, Ying Zhang, Di Kang, Xuefei Zhe, Linchao Bao, Xu Jia, and Huchuan Lu. Animatable neural radiance fields from monocular rgb videos, 2021. 3
- [12] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, 2017. 5
- [13] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 2
- [14] Enric Corona, Tomas Hodan, Minh Vo, Francesc Moreno-Noguer, Chris Sweeney, Richard Newcombe, and Lingni Ma. Lisa: Learning implicit shape and appearance of hands. *CVPR*, 2022. 3
- [15] Enric Corona, Gerard Pons-Moll, Guillem Alenyà, and Francesc Moreno-Noguer. Learned vertex descent: A new direction for 3d human model fitting. *ECCV*, 2022. 2
- [16] Enric Corona, Albert Pumarola, Guillem Alenyà, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *CVPR*, 2021. 1
- [17] Valentin Gabeur, Jean-Sébastien Franco, Xavier Martin, Cordelia Schmid, and Gregory Rogez. Moulding humans: Non-parametric 3d human shape estimation from single images. In *ICCV*, 2019. 2
- [18] Ivan Grishchenko, Valentin Bazarevsky, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Zanfir, Richard Yee, Karthik Raveendran, Matsvei Zhdanovich, Matthias Grundmann, and Cristian Sminchisescu. BlazePose ghum holistic: Real-time 3d human landmarks and pose estimation. *arXiv*, 2022. 6
- [19] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. *ICML*, 2020. 5
- [20] Tong He, John Collomosse, Hailin Jin, and Stefano Soatto. Geo-pifu: Geometry and pixel aligned implicit functions for single-view human reconstruction. *NeurIPS*, 2020. 5
- [21] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. Arch++: Animation-ready clothed human reconstruction revisited. In *CVPR*, 2021. 1, 2, 5, 6
- [22] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. Arch: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 1, 2, 3, 5, 6
- [23] Chaonan Ji, Tao Yu, Kaiwen Guo, Jingxin Liu, and Yebin Liu. Geometry-aware single-image full-body human relighting. In *ECCV*, 2022. 2
- [24] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. Selfrecon: Self reconstruction your digital avatar from monocular video. In *CVPR*, 2022. 1
- [25] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. *ECCV*, 2022. 3
- [26] Yoshihiro Kanamori and Yuki Endo. Relighting humans: occlusion-aware inverse rendering for full-body human images. *SIGGRAPH Asia*, 2019. 2
- [27] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 2
- [28] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. 2
- [29] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. 2
- [30] Manuel Lagunas, Xin Sun, Jimei Yang, Ruben Villegas, Jianming Zhang, Zhixin Shu, Belen Masia, and Diego Gutierrez. Single-image full-body human relighting. *ECCV*, 2022. 2
- [31] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *CVPR*, 2017. 2
- [32] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. Self-correction for human parsing. *PAMI*, 2020. 8
- [33] Tianye Li, Shichen Liu, Timo Bolkart, Jiayi Liu, Hao Li, and Yajie Zhao. Topologically consistent multi-view face inference using volumetric sampling. In *ICCV*, 2021. 1
- [34] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ToG*, 2015. 2
- [35] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *SIGGRAPH*, 1987. 3

- [36] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 2
- [37] Marko Mihajlovic, Aayush Bansal, Michael Zollhoefer, Siyu Tang, and Shunsuke Saito. KeypointNeRF: Generalizing image-based volumetric avatars using relative spatial encoding of keypoints. In *ECCV*, 2022. 5
- [38] 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. *ECCV*, 2020. 2, 5
- [39] Aymen Mir, Thimeo Alldieck, and Gerard Pons-Moll. Learning to transfer texture from clothing images to 3d humans. In *CVPR*, 2020. 8
- [40] Ryota Natsume, Shunsuke Saito, Zeng Huang, Weikai Chen, Chongyang Ma, Hao Li, and Shigeo Morishima. Siclope: Silhouette-based clothed people. In *CVPR*, 2019. 2
- [41] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*, 2018. 2
- [42] Hayato Onizuka, Zehra Hayirci, Diego Thomas, Akihiro Sugimoto, Hideaki Uchiyama, and Rin-ichiro Taniguchi. Tetratsdf: 3d human reconstruction from a single image with a tetrahedral outer shell. In *CVPR*, 2020. 2
- [43] 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
- [44] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 2
- [45] 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 *CVPR*, 2021. 3
- [46] 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. 1, 3, 7
- [47] Albert Pumarola, Jordi Sanchez-Riera, Gary Choi, Alberto Sanfeliu, and Francesc Moreno-Noguer. 3dpeople: Modeling the geometry of dressed humans. In *ICCV*, 2019. 2
- [48] Edoardo Remelli, Timur Bagautdinov, Shunsuke Saito, Chenglei Wu, Tomas Simon, Shih-En Wei, Kaiwen Guo, Zhe Cao, Fabian Prada, Jason Saragih, et al. Drivable volumetric avatars using texel-aligned features. In *SIGGRAPH*, 2022. 1
- [49] Renderpeople dataset. <https://renderpeople.com/>. 5
- [50] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ToG*, 2017. 2
- [51] Yu Rong, Takaaki Shiratori, and Hanbyul Joo. Frankmocap: Fast monocular 3d hand and body motion capture by regression and integration. *arXiv*, 2020. 2
- [52] 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. 1, 2, 3, 5
- [53] 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. 1, 2, 3, 5
- [54] David Smith, Matthew Loper, Xiaochen Hu, Paris Mavroidis, and Javier Romero. Facsimile: Fast and accurate scans from an image in less than a second. In *ICCV*, 2019. 2
- [55] Shih-Yang Su, Frank Yu, Michael Zollhoefer, and Helge Rhodin. A-nerf: Surface-free human 3d pose refinement via neural rendering. *NeurIPS*, 2021. 3
- [56] Daichi Tajima, Yoshihiro Kanamori, and Yuki Endo. Relighting humans in the wild: Monocular full-body human relighting with domain adaptation. In *CGF*, 2021. 2
- [57] Ayush Tewari, Justus Thies, Ben Mildenhall, Pratul Srinivasan, Edgar Tretschk, W Yifan, Christoph Lassner, Vincent Sitzmann, Ricardo Martin-Brualla, Stephen Lombardi, et al. Advances in neural rendering. In *CGF*, 2022. 3
- [58] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *ECCV*, 2018. 2
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3
- [60] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human sdf. In *ECCV*, 2022. 3
- [61] 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. 3
- [62] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J Black. Icon: Implicit clothed humans obtained from normals. In *CVPR*, 2022. 1, 2, 5, 6, 7
- [63] Hongyi Xu, Thimeo Alldieck, and Cristian Sminchisescu. H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *NeurIPS*, 2021. 1, 3, 7
- [64] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Ghum & ghuml: Generative 3d human shape and articulated pose models. In *CVPR*, 2020. 2, 3, 5, 6
- [65] Xiangyu Xu and Chen Change Loy. 3d human texture estimation from a single image with transformers. In *CVPR*, 2021. 8
- [66] Ze Yang, Shenlong Wang, Sivabalan Manivasagam, Zeng Huang, Wei-Chiu Ma, Xinchun Yan, Ersin Yumer, and Raquel Urtasun. S3: Neural shape, skeleton, and skinning fields for 3d human modeling. In *CVPR*, 2021. 2
- [67] Kennard Yanting Chan, Guosheng Lin, Haiyu Zhao, and Weisi Lin. Integratedpifu: Integrated pixel aligned implicit function for single-view human reconstruction. In *ECCV*, 2022. 2

- [68] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 2021. [4](#)
- [69] Mihai Zanfir, Andrei Zanfir, Eduard Gabriel Bazavan, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. Thundr: Transformer-based 3d human reconstruction with markers. In *CVPR*, 2021. [2](#)
- [70] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. Im avatar: Implicit morphable head avatars from videos. In *CVPR*, 2022. [1](#)
- [71] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. *PAMI*, 2021. [1](#), [2](#), [5](#)
- [72] Zerong Zheng, Tao Yu, Yixuan Wei, Qionghai Dai, and Yebin Liu. Deephuman: 3d human reconstruction from a single image. In *ICCV*, 2019. [2](#)
- [73] Hao Zhu, Xinxin Zuo, Sen Wang, Xun Cao, and Ruigang Yang. Detailed human shape estimation from a single image by hierarchical mesh deformation. In *CVPR*, 2019. [2](#)