

KeypointNeRF: Generalizing Image-based Volumetric Avatars using Relative Spatial Encoding of Keypoints

Marko Mihajlovic^{1,2*}, Aayush Bansal², Michael Zollhoefer², Siyu Tang¹,
Shunsuke Saito²

¹ETH Zürich ²Reality Labs Research

markomih.github.io/KeypointNeRF

Abstract. Image-based volumetric humans using pixel-aligned features promise generalization to unseen poses and identities. Prior work leverages global spatial encodings and multi-view geometric consistency to reduce spatial ambiguity. However, global encodings often suffer from overfitting to the distribution of the training data, and it is difficult to learn multi-view consistent reconstruction from sparse views. In this work, we investigate common issues with existing spatial encodings and propose a simple yet highly effective approach to modeling high-fidelity volumetric humans from sparse views. One of the key ideas is to encode relative spatial 3D information via sparse 3D keypoints. This approach is robust to the sparsity of viewpoints and cross-dataset domain gap. Our approach outperforms state-of-the-art methods for head reconstruction. On human body reconstruction for unseen subjects, we also achieve performance comparable to prior work that uses a parametric human body model and temporal feature aggregation. Our experiments show that a majority of errors in prior work stem from an inappropriate choice of spatial encoding and thus we suggest a new direction for high-fidelity image-based human modeling.

Keywords: Neural Radiance Field, Pixel-Aligned Features

1 Introduction

3D renderable human representations are an important component for social telepresence, mixed reality, and a new generation of entertainment platforms. Classical mesh-based methods require dense multi-view stereo [34, 57, 58] or depth sensor fusion [68]. The fidelity of these methods is limited due to the difficulty of accurate geometry reconstruction. Recently, neural volumetric representations [29, 40] have enabled high-fidelity human reconstruction, especially where accurate geometry is difficult to obtain (e.g. hair). By injecting human-specific parametric shape models [7, 32], extensive multi-view data capture can

*The work was primarily done during an internship at Meta.

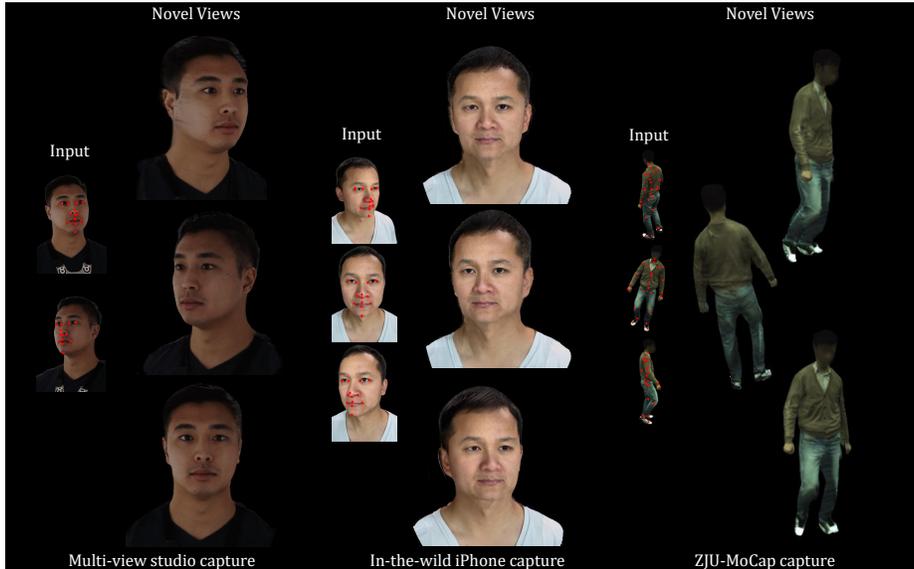


Fig. 1. (a) Our approach allows us to synthesize high-fidelity volumetric humans from two or three views of unseen subjects. (b) A model learned on studio captures can be used without modification on in-the-wild iPhone captures; and (c) finally, the same approach also allows us to synthesize novel views of unseen human subjects (faces are blurred). The figure is best viewed in electronic format.

be reduced to sparse camera setups [15, 43]. However, these learning-based approaches are subject-specific and require days of training for each individual subject, which severely limits their scalability. Democratizing digital volumetric humans requires an ability to instantly create a personalized reconstruction of a user from two or three snaps (from different views) taken from their phone. Towards this goal, we learn to generalize metrically accurate image-based volumetric humans from two or three views.

Fully convolutional pixel-aligned features utilizing multi-scale information have enabled better performance for various 2D computer vision tasks [6, 11, 31, 41], including the generalizable reconstruction of unseen subjects [27, 46, 49, 67]. Pixel-aligned neural fields infer field quantities given a pixel location and spatial encoding function (to avoid ray-depth ambiguity). Different spatial encoding functions [16, 19, 21, 46] have been proposed in the literature. However, the effect of spatial encoding is not fully understood. In this paper, we provide an extensive analysis of spatial encodings for modeling pixel-aligned neural radiance fields for human faces. Our experiments show that the choice of spatial encoding influences the reconstruction quality and generalization to novel identities and views. The models that use camera depth overfit to the training distribution, and multi-view stereo constraints are less robust to sparse views with large baselines.

We present a simple yet highly effective approach based on sparse 3D keypoints to address the limitations of existing approaches. 3D keypoints are easy to obtain using an off-the-shelf 2D keypoint detector [11] and a simple triangulation of multi-views [18]. We treat 3D keypoints as spatial anchors and encode relative 3D spatial information to those keypoints. Unlike global spatial encoding [49, 50, 67], the relative spatial information is agnostic to camera parameters. This property allows the proposed approach to be robust to changes in pose. 3D keypoints also allow us to use the same approach for both human faces and bodies. Our approach achieves state-of-the-art performance for generating volumetric humans for unseen subjects from sparse-and-wide two or three views, and we can also incorporate more views to further improve performance. We also achieve performance comparable to Neural Human Performer (NHP) [27] when it comes to full-body human reconstruction. NHP relies on accurate parametric body fitting and temporal feature aggregation, whereas our approach employs 3D keypoints alone. Our method is not biased [56] to the distribution of the training data. We can use the learned model (without modification) to never-before-seen iPhone captures. We attribute our ability to generalize image-based volumetric humans to an unseen data distribution to our choice of spatial encoding. Our key contributions include:

- A simple approach that leverages sparse 3D keypoints and allows us to create high-fidelity state-of-the-art volumetric humans for unseen subjects from widely spread out two or three views.
- Extensive analysis on the use of spatial encodings to understand their limitations and the efficacy of the proposed approach.
- We demonstrate generalization to never-before-seen iPhone captures by training with only a studio-captured dataset. To our knowledge, no prior work has shown these results.

2 Related Work

Our goal is to create high-fidelity volumetric humans for unseen subjects from as few as two views.

Classical Template-based Techniques: Early work on human reconstruction [22] required dense 3D reconstruction from a large number of images of the subject and non-rigid registration to align a template mesh to 3D point clouds. Cao et al. [10] employ coarse geometry along with face blendshapes and a morphable hair model to address restrictions posed by dense 3D reconstruction. Hu et al. [20] retrieve hair templates from a database and carefully compose facial and hair details. Video Avatar [1] obtains a full-body avatar based on a monocular video captured using silhouette-based modeling. The dependence on geometry and meshes restricts the applicability of these methods to faithfully reconstruct regions such as the hair, mouth, teeth, etc., where it is non-trivial to obtain accurate geometry.

Neural Rendering: Neural rendering [54, 55] has tackled some of the challenges classical template-based approaches struggle with by directly learning compo-

nents of the image formation process from raw sensor measurements. 2D neural rendering approaches [25, 33, 35, 37, 44, 45] employ surface rendering and a convolutional network to bridge the gap between rendered and real images. The downside of these 2D techniques is that they struggle to synthesize novel viewpoints in a temporally coherent manner. Deep Appearance Models [28] employ a coarse 3D proxy mesh in combination with view-dependent texture mapping to learn personalized face avatars from dense multi-view supervision. Using a 3D proxy mesh significantly helps with viewpoint generalization, but the approach faces challenges in synthesizing certain regions for which it is hard to obtain good 3D reconstruction, such as the hair and inside the mouth. Current state-of-the-art methods such as NeuralVolumes [29] and NeRF [40] employ differentiable volumetric rendering instead of relying on meshes. Due to their volumetric nature, these methods enable high-quality results even for regions where estimating 3D geometry is challenging. Various extensions [5, 30, 63] have further improved quality. These methods require dense multi-view supervision for person-specific training and take 3–4 days to train a single model.

Sparse View Reconstruction: Large scale deployment requires approaches that allow a user to take two or three pictures of themselves from multi-views and generate a digital human from this data. The use of pixel-aligned features [49, 50] further allows the use of large datasets for learning generalized models from sparse views. Different approaches [13, 14, 60, 67] have combined multi-view constraints and pixel-aligned features alongside NeRF to learn generalizable view-synthesis. In this work, we observe that these approaches struggle to generate fine details given sparse views for unseen human faces and bodies.

Learning Face and Body Reconstruction: Generalizable parametric mesh [4, 32, 66] and implicit [2, 36, 38, 61, 65] body models can provide additional constraints for learning details from sparse views. Recent approaches have incorporated priors specific to human faces [8, 9, 15, 17, 48, 59, 70] and human bodies [42, 43, 62, 64, 65, 69, 71, 72] to reduce the dependence on multi-view captures. Approaches such as H3DNet [47] and SIDER [12] use signed-distance functions (SDFs) for learning geometry priors from large datasets and perform test-time fine-tuning on the test subject. PaMIR [72] uses volumetric features guided by a human body model for better generalization. Neural Human Performer [27] employs SMPL with pixel-aligned features and temporal feature aggregation. In this work, we observe that the use of human 3D-keypoints provides necessary and sufficient constraints for learning from sparse-view inputs. Our approach has high flexibility since it only relies on 3D keypoints alone and thus enables us to work both on human faces and bodies. Prior methods have also employed various forms of spatial encoding for better learning. For example, PVA [46] and PortraitNeRF [16] use face-centric coordinates. ARCH/ARCH++ [19, 21] use canonical body coordinates. In this work, we extensively study the role of spatial encoding, and found that the use of a relative depth encoding using 3D keypoints leads to the best results. Our findings enable us to learn a representation that generalizes to never-before-seen iPhone camera captures for unseen human faces. In addition to achieving state-of-the-art results on volumetric face

reconstruction from as few as two images, our approach can also be used for synthesizing novel views of unseen human bodies and achieves competitive performance to prior work [27] in this setting.

3 Preliminaries: Neural Radiance Fields

NeRF [40] is a continuous function that represents the scene as a volumetric radiance field of color and density. Given a 3D point and a viewing direction $d \in \mathbb{R}^3$, NeRF estimates RGB color values and density (c, σ) that are then accumulated via quadrature to calculate the expected color of each camera ray $r(t) = o + td$:

$$C(r) = \int_{t_n}^{t_f} \exp\left(-\int_{t_n}^t \sigma(s) ds\right) \sigma(t)c(t, d) dt, \quad (1)$$

where t_n and t_f define near and far bounds.

Pixel-aligned NeRF. One of the core limitations of NeRF is that the approach requires per-scene optimization and does not work well for extremely sparse input views (e.g., two images). To address these challenges, several recent methods [46, 60, 67] propose to condition NeRF on pixel-aligned image features and generalize to novel scenes without retraining.

Spatial Encoding. To avoid the ray-depth ambiguity, pixel-aligned neural fields [46, 49, 50, 67] attach spatial encoding to the pixel-aligned feature. PIFu [49] and related methods [50, 67] use depth value in the camera coordinate space as spatial encoding, while PVA [46] uses coordinates relative to the head position. However, we argue that such spatial encodings are global and sub-optimal for learning generalizable volumetric humans. In contrast, our proposed relative spatial encoding provides a localized context that enables better learning and is more robust to changes in human pose.

4 KeypointNeRF

Our method is based on a radiance field function:

$$f(X, d | \{(I_n, P_n)\}_{n=1}^N) = (c, \sigma) \quad (2)$$

that infers a color $c \in \mathbb{R}^3$ and a density $\sigma \in \mathbb{R}$ value for any point in 3D space given its position $X \in \mathbb{R}^3$ and its viewing direction $d \in \mathbb{R}^3$ as input. In addition, the function has access to the N input images I_n and their camera calibrations P_n . We model the function f as a neural network that consists of four main parts; a spatial keypoint encoder (Sec. 4.1), two convolutional image encoders that extract pixel-aligned features (Sec. 4.2), an MLP fusion network that aggregates multiple pixel-aligned features (Sec. 4.3), and two MLPs that predict density σ and color values c (Sec. 4.4). The high-level overview of our method is illustrated in Fig. 2 and in the following we further describe its components.

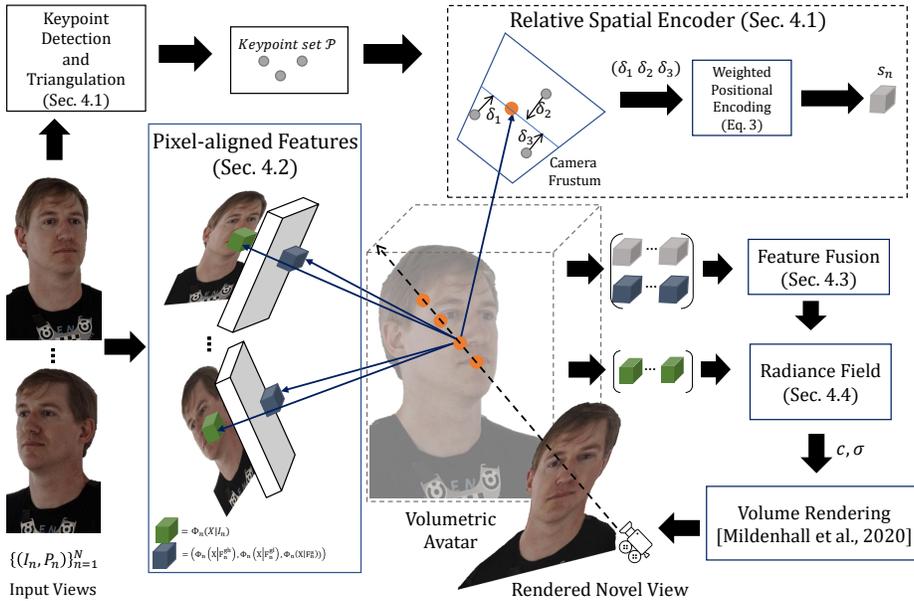


Fig. 2. Overview. We learn a generalizable image-based neural radiance representation for volumetric humans. Given a sparse set of input images $\{I_n\}_{n=1}^N$ and their respective camera parameters P_n , we first detect keypoints and lift them to 3D \mathcal{P} . The keypoints are used to provide the relative spatial encoding (Sec. 4.1). The input images are simultaneously encoded via convolutional encoders and provide image-guided pixel-aligned features (Sec. 4.2). These two types of encoding are fused (Sec. 4.3) and condition the radiance field (Sec. 4.4). The radiance field is decoupled by two MLPs, one that directly predicts view-independent density value σ , and the other one which predicts blending weights that are used to output the final color value c by blending image pixel values $\{\Phi(X|I_n)\}_{n=1}^N$. The predicted color and density values are accumulated along the ray via volume rendering [39] to render the volumetric representation from novel views. The rendered example in the figure is an actual output of our method for the displayed two input images of an unseen subject.

4.1 Relative Spatial Keypoint Encoding

Our method first leverages an off-the-shelf keypoint regressor [11] to extract K 2D keypoints from at least two input views. Then these points are triangulated and lifted to 3D $\mathcal{P} = \{p_k \in \mathbb{R}^3\}_{k=1}^K$ by using the direct linear transformation algorithm [18]. To spatially encode the query point X , we first compute the depth values of the query point and all keypoints w.r.t each input view $z(p_k|P_n)$ by the 2D projection and then estimate their relative depth difference $\delta_n(p_k, X) = z(p_k|P_n) - z(X|P_n)$. We further employ positional encoding $\gamma(\cdot)$ [40] and the Gaussian exponential kernel to create the final relative spatial encoding for the

query point X as:

$$s_n(X|\mathcal{P}) = \left[\exp\left(\frac{-l_2(p_k, X)^2}{2\alpha^2}\right) \gamma(\delta_n(p_k, X)) \right]_{k=1}^K, \quad (3)$$

where α is a fixed hyper-parameter that controls the impact of each keypoint. We set this value to 5cm for facial keypoints and to 10cm for the human skeleton.

4.2 Convolutional Pixel-aligned Features

In addition to the spatial encoding $s_n(X)$, we extract pixel-aligned features for the query point X by encoding the input images $I_n \in \mathbb{R}^{H \times W \times 3}$ using two convolutional encoders.

Image Encoders. The first image encoder uses a single HourGlass [41] network that generates both deep low-resolution $F_n^{gl} \in \mathbb{R}^{H/8 \times W/8 \times 64}$ and shallow high-resolution $F_n^{gh} \in \mathbb{R}^{H/2 \times W/2 \times 8}$ feature maps. This network learns a geometric prior of humans and its output is used to condition the density estimation network. The second encoder is a convolutional network with residual connections [23] that encodes input images $F_n^a \in \mathbb{R}^{H/4 \times W/4 \times 8}$ and provides an alternative pathway for the appearance information which is independent of the density prediction branch in the spirit of DoubleField [52]. Please see the supplemental material for further architectural details.

Pixel-aligned Features. To compute the pixel-aligned features, we project the query point on each feature plane $x = \pi(X|P_n) \in \mathbb{R}^2$ and bi-linearly interpolate the grid values. We define this operation of computing the pixel-aligned features (2D projection and interpolation) by the operator $\Phi_n(X|F)$, where F can represent any grid of vectors for the n th camera: $F_n^{gl}, F_n^{gh}, F_n^a, I_n$.

4.3 Multi-view Feature Fusion

To model a multi-view consistent radiance field, we need to fuse the per-view spatial encodings s_n (Eq. 3) and the pixel-aligned features Φ_n .

The spatial encoding s_n is first translated into a feature vector via a single-layer perceptron. This feature is then jointly blended with the deep low-resolution pixel-aligned feature $\Phi_n(X|F_n^{gl})$ by a two-layer perceptron. The output is then concatenated with the shallow high-resolution feature $\Phi_n(X|F_n^{gh})$ and processed by an additional two-layer perceptron that outputs per-view 64-dimensional feature vector that jointly encodes the blended spatial encoding and the pixel-aligned information. These multi-view features are then fused into a single feature vector $G_X \in \mathbb{R}^{128}$ by the mean and variance pooling operators as in [60].

4.4 Modeling Radiance Fields

The radiance field is modeled via decoupled MLPs for density σ and color c prediction.

Density Fields. The density network is implemented as a four-layer MLP that takes as input the geometry feature vector G_X and predicts the density value σ .

View-dependent Color Fields. We implement an additional MLP to output the consistent color value c for a given query point X and its viewing direction d by blending image pixel values $\{\Phi(X|P_n)\}_{n=1}^N$ similarly to IBRNet [60]. The input to this MLP is 1) the extracted geometry feature vector G_X that ensures geometrically consistent renderings, 2) the additional pixel-aligned features $\Phi_n(X|F_n^a)$, 3) the corresponding pixel values $\Phi_n(X|I_n)$, and 4) the view direction that is encoded as the difference between the view direction d and the camera views along with their dot product.

These inputs are concatenated and augmented with the mean and variance vectors computed over the multi-view pixel-aligned features, and jointly propagated through a nine-layer perceptron with residual connections which predicts blending weights for each input view $\{w_n\}_{n=1}^N$. These blending weights form the final color prediction by fusing the corresponding pixel-aligned color values:

$$c = \sum_{n=1}^N \frac{\exp(w_n)\Phi_n(X|I_n)}{\sum_{i=1}^N \exp(w_i)}. \quad (4)$$

5 Novel View Synthesis

Given our radiance field function $f(X, d) = (c, \sigma)$, we render novel views via the volume rendering equation (1), in which we define the near and far bound by analytically computing the intersection of the pixel ray and a geometric proxy that over-approximates the volumetric human and use the entrance and exit points as near and far bounds respectively. For the experiments on human heads, we use a sphere with a radius of 30 centimeters centered around the keypoints, while for the human bodies we follow the prior work [27, 43] and use a 3D bounding box. Similar to NeRF [40], we employ a coarse-to-fine rendering strategy, but we employ the same network weights for both levels.

5.1 Training and Implementation Details

To train our network, we render $H' \times W'$ patches (as in [51]) by accumulating color and density values for 64 sampled points along the ray for the coarse and 64 more for the fine rendering. Our method is trained end-to-end by minimizing the mean ℓ_1 -distance between the rendered and the ground truth pixel values and the VGG [53]-loss applied over rendered and ground truth image patches:

$$\mathcal{L} = \mathcal{L}_{\text{RGB}} + \mathcal{L}_{\text{VGG}}. \quad (5)$$

The use of the VGG loss for NeRF training was also leveraged by the concurrent methods [3, 64] to better capture high-frequency details. The final loss \mathcal{L} is minimized by the Adam optimizer [26] with a learning rate of $1e^{-4}$ and a batch size of one. For the other parameters, we use their defaults. The background from all training and test input images is removed via an off-the-shelf matting network [24]. Additionally for more temporally coherent novel-view synthesis at inference time, we clip the maximum of the dot product (introduced in Sec. 4.4) to 0.8 when the number of input images is two in the supplementary video.

6 Experiments

In this section, we validate our method on three different reconstruction tasks and datasets: 1) reconstruction of human heads from images captured in a multi-camera studio, 2) reconstruction of human heads from in-the-wild images taken with the iPhone’s camera, and 3) reconstruction of human bodies on the public ZJU-MoCap dataset [43]. As evaluation metrics, we follow prior work [27, 46, 60] and report the standard SSIM and PSNR metrics.

6.1 Reconstruction of Human Heads from Studio Data

Dataset and Experimental Setup. Our captured data consists of 29 1280×768 -resolution cameras positioned in front of subjects. We use a total of 351 identities and 26 cameras for training and 38 novel identities for evaluation. At inference time, we reconstruct humans only from 2–3 input views.

Baselines. As baseline, we employ the current state-of-the-art model IBRNet [60]. In addition, we add several other baselines by varying different types of encoding for the query points in our proposed reconstruction pipeline. Specifically, 1) our pipeline without any encoding, 2) with the camera z encoding used in [49, 50], 3) with the encoding of xyz coordinates in the canonical space of a human head that is used in [46], 4) relative encoding of xyz as the distance between the query point and estimated keypoints, 5) our relative spatial encoding without distance weighing ($\alpha \rightarrow \infty$ in Eq. 3), and 6) the proposed weighted relative encoding as described in the method section (Sec. 4.3). The last three models use a total of 13 facial keypoints that are visualized in Fig. 1. All methods are trained with a batch size of one for 150k training steps, except IBRNet [60] which was trained for 200k iterations. For more comparisons and baselines we refer the reader to the supplementary material and video.

Table 1. Studio Capture Results. Quantitative comparison with IBRNet [60] and different types of spatial encoding. Visual results are provided in Fig. 3

	SSIM↑	PSNR↑
PVA [46]	81.95	25.87
IBRNet [60]	82.39	27.14
1) no encoding	84.38	27.16
2) camera z encoding [49]	77.86	22.66
3) canonical xyz encoding [46]	83.11	26.33
4) relative xyz encoding	83.66	27.05
5) relative z encoding	84.72	27.66
6) KeypointNeRF (distance weighted relative z encoding, Eq. 3)	85.19	27.64

Results. We provide novel view synthesis (Fig. 3) results for unseen identities that have been reconstructed from only two input images. The results clearly demonstrate that the rendered images of our method are significantly sharper

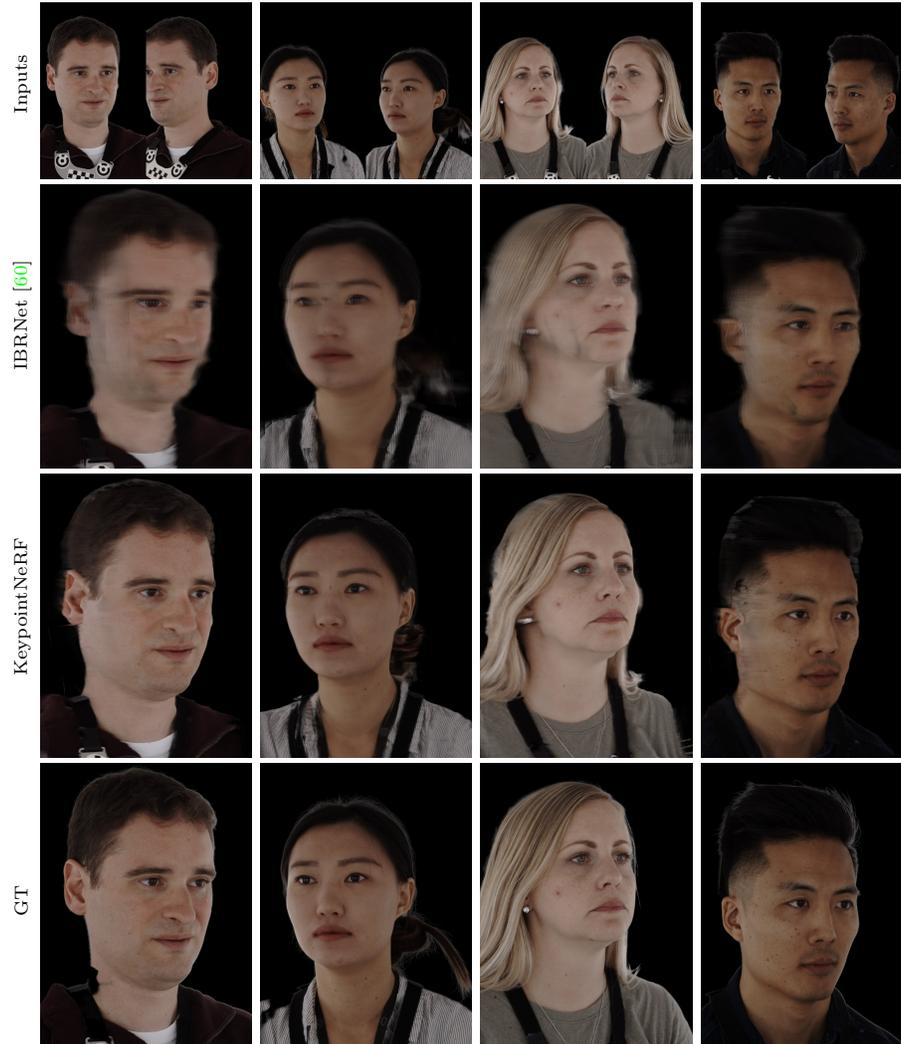


Fig. 3. Studio Capture Results. Reconstruction results on held-out subjects from only two input views. Best viewed in electronic format.

compared to the baselines and are of significantly higher quality. This improvement is confirmed by the quantitative evaluation (Tab. 1) which further indicates that the proposed distance weighting of the relative spatial encoding improves the reconstruction quality. The third-best performing method is our pipeline without any spatial encoding. However, such a method does not generalize well to other capture systems as we will demonstrate in the next section on in-the-wild captured data.

Table 2. Robustness to Different Noise Levels. Perturbing our keypoints vs. perturbing head position in the canonical xyz encoding used in [46]. Our proposed encoding demonstrates significantly slower performance degradation as the noise increases

Noise level [mm]	Canonical xyz encoding		Our keypoint encoding	
	SSIM↑	PSNR↑	SSIM↑	PSNR↑
no noise	83.65	27.05	85.19	27.64
1	82.79	26.24	85.20	27.64
2	82.26	26.05	85.08	27.62
3	81.58	25.48	84.96	27.59
4	80.92	25.08	84.95	27.56
5	80.36	25.16	84.80	27.51
10	76.62	22.23	83.69	27.10
15	73.33	20.26	82.33	26.47
20	70.40	18.87	81.17	25.77

Robustness to Different Noise Levels. We evaluate the robustness of our relative spatial encoding and the encoding in canonical space proposed by Raj et al. [46] by adding different noise levels to the estimated keypoints and the head center respectively. The results reported in Tab. 2 show that our proposed encoding based on keypoints is significantly more robust compared to modeling in an object-specific canonical space. Note that the canonical encoding requires head template fitting for which we used the ground truth estimation from all views and which, in practice, is very erroneous or even infeasible from two views alone.

Dynamic scenes. Although our model is trained only with a neutral face, it generalizes well to dynamic expressions and outperforms the baseline methods. We evaluate the trained models on 38 test subjects performing eight different expressions and report results in Tab. 3.

Table 3. Dynamic expressions. Our model is more accurate than the baselines

								SSIM↑ PSNR↑		
IBRNet								IBRNet [60]	82.64	26.79
								Ours (no keypoints)	84.97	27.14
Ours								Ours	85.31	27.30

6.2 Reconstruction from In-the-wild Captures

Setup. To tackle the problem of reconstructing humans in the wild, we acquire a small dataset of subjects by taking several photos with an iPhone and estimate camera parameters. We directly use intrinsic from manufacturing information of iPhone and extrinsic is computed by multi-view RGB-D fitting as in [9].



Fig. 4. In-the-wild Captures. Our approach produces high-quality reconstructions from three iPhone images, while all baselines show significant artifacts.

We evaluate the reconstruction methods trained on the studio captured data in Sec. 6.1 without any retraining.

As input, all methods take three 1920×1024 -resolution images of a person and predict a radiance field that is then rendered from novel views. In Fig. 4, we display rendered novel views of IBRNet [60], our method without any spatial encoding, and our method with the proposed spatial encoding. The baseline methods produce significantly worse results with lots of blur and cloudy artifacts, whereas our method can reliably reconstruct the human heads. This improvement is quantitatively supported by computed SSIM and PSNR on novel held-out views of the visualized four subjects (Tab. 4). This experiment demonstrates that our relative spatial encoding is the crucial component for cross-dataset generalization. Please see the supplementary material for more visualizations.

Table 4. In-the-wild Captures.

	SSIM \uparrow	PSNR \uparrow
IBRNet [60]	81.74	18.45
Ours (no keypoints)	79.50	19.79
KeypointNeRF	86.73	25.29

6.3 Reconstruction of Human Bodies

Additionally, we demonstrate that our method is suitable for reconstructing full volumetric human bodies without relying on template fitting of parametric human bodies [32]. We use the public ZJU [43] dataset in order to follow the experimental setup used in [27], so that we could closely compare our method’s ability to reconstruct human bodies to the current state-of-the-art method without changing any experimental variables. We follow the standard training-test split of frames and use a total of seven subjects for training and three for validation. At inference time, all methods use three input views. We compare our method with the generalizable volumetric methods: pixelNeRF [67], PVA [46], the current state-of-the-art Neural Human Performer (NHP) [27], and our method without weighting the relative spatial encoding in Eq. 3. We report results on unseen identities for 438 novel views in Table 5 and side-by-side qualitative comparisons with NHP in Fig. 5. The results demonstrate that weighting the spatial encoding benefits reconstruction of human bodies as well. Our method is on par with the significantly more complex NHP, which relies on the accurate registration of the SMPL body model [32] and temporal feature fusion, whereas ours only requires skeleton keypoints.

7 Conclusion

We present a simple yet highly effective approach for generating high-fidelity volumetric humans from as few as two input images. The key to our approach is a novel spatial encoding based on relative information extracted from 3D keypoints. Our approach outperforms state-of-the-art methods for head reconstruction and better generalizes to challenging out-of-domain inputs, such as selfies captured in the wild by an iPhone. Since our approach does not require

References

1. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018) [3](#)
2. Alldieck, T., Xu, H., Sminchisescu, C.: imghum: Implicit generative models of 3d human shape and articulated pose. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
3. Alldieck, T., Zanfir, M., Sminchisescu, C.: Photorealistic monocular 3d reconstruction of humans wearing clothing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1506–1515 (2022) [8](#)
4. Angelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM SIGGRAPH 2005 Papers, pp. 408–416 (2005) [4](#)
5. Athar, S., Shu, Z., Samaras, D.: Flame-in-nerf : Neural control of radiance fields for free view face animation. arXiv preprint arXiv:2108.04913 (2021) [4](#)
6. Bansal, A., Chen, X., Russell, B., Gupta, A., Ramanan, D.: Pixelnet: Representation of the pixels, by the pixels, and for the pixels. arXiv preprint arXiv:1702.06506 (2017) [2](#)
7. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th annual conference on Computer graphics and interactive techniques. pp. 187–194 (1999) [1](#)
8. Buehler, M.C., Meka, A., Li, G., Beeler, T., Hilliges, O.: Varitex: Variational neural face textures. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2021) [4](#)
9. Cao, C., Simon, T., Kim, J.K., Schwartz, G., Zollhoefer, M., Saito, S., Lombardi, S., en Wei, S., Belko, D., i Yu, S., Sheikh, Y., Saragih, J.: Authentic volumetric avatars from a phone scan. ACM Transactions on Graphics (TOG) (2022) [4](#), [11](#)
10. Cao, C., Wu, H., Weng, Y., Shao, T., Zhou, K.: Real-time facial animation with image-based dynamic avatars. ACM Transactions on Graphics **35**(4) (2016) [3](#)
11. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7291–7299 (2017) [2](#), [3](#), [6](#)
12. Chatziagapi, A., Athar, S., Moreno-Noguer, F., Samaras, D.: Sider: Single-image neural optimization for facial geometric detail recovery. arXiv preprint arXiv:2108.05465 (2021) [4](#)
13. Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2021) [4](#)
14. Chibane, J., Bansal, A., Lazova, V., Pons-Moll, G.: Stereo radiance fields (srf): Learning view synthesis from sparse views of novel scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (jun 2021) [4](#)
15. Gafni, G., Thies, J., Zollhofer, M., Niessner, M.: Dynamic neural radiance fields for monocular 4d facial avatar reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [2](#), [4](#)
16. Gao, C., Shih, Y., Lai, W.S., Liang, C.K., Huang, J.B.: Portrait neural radiance fields from a single image. arXiv preprint arXiv:2012.05903 (2020) [2](#), [4](#)
17. Grassal, P.W., Prinzler, M., Leistner, T., Rother, C., Nießner, M., Thies, J.: Neural head avatars from monocular rgb videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18653–18664 (2022) [4](#)

18. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003) [3](#), [6](#)
19. He, T., Xu, Y., Saito, S., Soatto, S., Tung, T.: Arch++: Animation-ready clothed human reconstruction revisited. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2021) [2](#), [4](#)
20. Hu, L., Saito, S., Wei, L., Nagano, K., Seo, J., Fursund, J., Sadeghi, I., Sun, C., Chen, Y.C., Li, H.: Avatar digitization from a single image for real-time rendering. *ACM Transactions on Graphics (ToG)* **36**(6), 1–14 (2017) [3](#)
21. Huang, Z., Xu, Y., Lassner, C., Li, H., Tung, T.: Arch: Animatable reconstruction of clothed humans. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [2](#), [4](#)
22. Ichim, A.E., Bouaziz, S., Pauly, M.: Dynamic 3d avatar creation from hand-held video input. *ACM Transactions on Graphics (ToG)* **34**(4), 1–14 (2015) [3](#)
23. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016) [7](#)
24. Ke, Z., Sun, J., Li, K., Yan, Q., Lau, R.W.: Modnet: Real-time trimap-free portrait matting via objective decomposition. In: AAAI (2022) [8](#)
25. Kim, H., Garrido, P., Tewari, A., Xu, W., Thies, J., Nießner, M., Pérez, P., Richardt, C., Zollöfer, M., Theobalt, C.: Deep video portraits. *ACM Transactions on Graphics (TOG)* **37**(4), 163 (2018) [4](#)
26. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proceedings of International Conference on Learning Representations (2015) [8](#)
27. Kwon, Y., Kim, D., Ceylan, D., Fuchs, H.: Neural human performer: Learning generalizable radiance fields for human performance rendering. In: Advances in Neural Information Processing Systems (NeurIPS). Curran Associates, Inc. (2021) [2](#), [3](#), [4](#), [5](#), [8](#), [9](#), [13](#), [14](#)
28. Lombardi, S., Saragih, J., Simon, T., Sheikh, Y.: Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018) [4](#)
29. Lombardi, S., Simon, T., Saragih, J., Schwartz, G., Lehrmann, A., Sheikh, Y.: Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)* (2019) [1](#), [4](#)
30. Lombardi, S., Simon, T., Schwartz, G., Zollhoefer, M., Sheikh, Y., Saragih, J.: Mixture of volumetric primitives for efficient neural rendering. arXiv preprint arXiv:2103.01954 (2021) [4](#)
31. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015) [2](#)
32. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* **34**(6), 1–16 (2015) [1](#), [4](#), [13](#)
33. Martin-Brualla, R., Pandey, R., Yang, S., Pidlypenskyi, P., Taylor, J., Valentin, J., Khamis, S., Davidson, P., Tkach, A., Lincoln, P., et al.: Lookingood: Enhancing performance capture with real-time neural re-rendering. arXiv preprint arXiv:1811.05029 (2018) [4](#)
34. Matusik, W., Buehler, C., Raskar, R., Gortler, S.J., McMillan, L.: Image-based visual hulls. In: ACM SIGGRAPH. pp. 369–374 (2000) [1](#)
35. Meka, A., Pandey, R., Haene, C., Orts-Escolano, S., Barnum, P., David-Son, P., Erickson, D., Zhang, Y., Taylor, J., Bouaziz, S., et al.: Deep relightable textures: volumetric performance capture with neural rendering. *ACM Transactions on Graphics (TOG)* **39**(6), 1–21 (2020) [4](#)

36. Mihajlovic, M., Saito, S., Bansal, A., Zollhoefer, M., Tang, S.: COAP: Compositional articulated occupancy of people. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2022) [4](#)
37. Mihajlovic, M., Weder, S., Pollefeys, M., Oswald, M.R.: Deepsurfels: Learning on-line appearance fusion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14524–14535 (2021) [4](#)
38. Mihajlovic, M., Zhang, Y., Black, M.J., Tang, S.: LEAP: Learning articulated occupancy of people. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [4](#)
39. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* (2019) [6](#)
40. 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. pp. 405–421. Springer (2020) [1](#), [4](#), [5](#), [6](#), [8](#)
41. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: European conference on computer vision. pp. 483–499. Springer (2016) [2](#), [7](#)
42. Peng, S., Dong, J., Wang, Q., Zhang, S., Shuai, Q., Zhou, X., Bao, H.: Animatable neural radiance fields for modeling dynamic human bodies. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14314–14323 (2021) [4](#)
43. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [2](#), [4](#), [8](#), [9](#), [13](#), [14](#)
44. Prokudin, S., Black, M.J., Romero, J.: Smpipix: Neural avatars from 3d human models. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1810–1819 (2021) [4](#)
45. Raj, A., Tanke, J., Hays, J., Vo, M., Stoll, C., Lassner, C.: Anr: Articulated neural rendering for virtual avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3722–3731 (2021) [4](#)
46. Raj, A., Zollhoefer, M., Simon, T., Saragih, J., Saito, S., Hays, J., Lombardi, S.: Pva: Pixel-aligned volumetric avatars. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [2](#), [4](#), [5](#), [9](#), [11](#), [13](#), [14](#)
47. Ramon, E., Triginer, G., Escur, J., Pumarola, A., Garcia, J., i Nieto, X.G., Moreno-Noguer, F.: H3d-net: Few-shot high-fidelity 3d head reconstruction. *arXiv preprint arXiv:2107.12512* (2021) [4](#)
48. Rebain, D., Matthews, M., Yi, K.M., Lagun, D., Tagliasacchi, A.: Lolnerf: Learn from one look. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1558–1567 (2022) [4](#)
49. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2019) [2](#), [3](#), [4](#), [5](#), [9](#)
50. Saito, S., Simon, T., Saragih, J., Joo, H.: Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) [3](#), [4](#), [5](#), [9](#)

51. Schwarz, K., Liao, Y., Niemeyer, M., Geiger, A.: Graf: Generative radiance fields for 3d-aware image synthesis. In: *Advances in Neural Information Processing Systems (NeurIPS)*. Curran Associates, Inc. (2020) [8](#)
52. Shao, R., Zhang, H., Zhang, H., Cao, Y., Yu, T., Liu, Y.: Doublefield: Bridging the neural surface and radiance fields for high-fidelity human rendering. *arXiv preprint arXiv:2106.03798* (2021) [7](#)
53. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014) [8](#)
54. Tewari, A., Fried, O., Thies, J., Sitzmann, V., Lombardi, S., Sunkavalli, K., Martin-Brualla, R., Simon, T., Saragih, J., Nießner, M., et al.: State of the art on neural rendering. In: *Computer Graphics Forum*. vol. 39, pp. 701–727. Wiley Online Library (2020) [3](#)
55. Tewari, A., Thies, J., Mildenhall, B., Srinivasan, P., Tretschk, E., Wang, Y., Lassner, C., Sitzmann, V., Martin-Brualla, R., Lombardi, S., et al.: Advances in neural rendering. *arXiv preprint arXiv:2111.05849* (2021) [3](#)
56. Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: *CVPR 2011*. pp. 1521–1528. IEEE (2011) [3](#)
57. Vlastic, D., Baran, I., Matusik, W., Popović, J.: Articulated mesh animation from multi-view silhouettes. *ACM Transactions on Graphics* **27**(3), 97 (2008) [1](#)
58. Vlastic, D., Peers, P., Baran, I., Debevec, P., Popović, J., Rusinkiewicz, S., Matusik, W.: Dynamic shape capture using multi-view photometric stereo. *ACM Transactions on Graphics* **28**(5), 174 (2009) [1](#)
59. Wang, L., Chen, Z., Yu, T., Ma, C., Li, L., Liu, Y.: Faceverse: a fine-grained and detail-controllable 3d face morphable model from a hybrid dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20333–20342 (2022) [4](#)
60. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4690–4699 (2021) [4](#), [5](#), [7](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#)
61. Wang, S., Mihajlovic, M., Ma, Q., Geiger, A., Tang, S.: Metaavatar: Learning animatable clothed human models from few depth images. *Advances in Neural Information Processing Systems* **34** (2021) [4](#)
62. Wang, S., Schwartz, K., Geiger, A., Tang, S.: ARAH: Animatable volume rendering of articulated human sdf. In: *European conference on computer vision (2022)* [4](#)
63. Wang, Z., Bagautdinov, T., Lombardi, S., Simon, T., Saragih, J., Hodgins, J., Zollhofer, M.: Learning compositional radiance fields of dynamic human heads. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2021) [4](#)
64. Weng, C.Y., Curless, B., Srinivasan, P.P., Barron, J.T., Kemelmacher-Shlizerman, I.: Humannerf: Free-viewpoint rendering of moving people from monocular video. *arXiv preprint arXiv:2201.04127* (2022) [4](#), [8](#)
65. Xu, H., Alldieck, T., Sminchisescu, C.: H-nerf: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *Advances in Neural Information Processing Systems* **34** (2021) [4](#)
66. Xu, H., Bazavan, E.G., Zanfir, A., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Ghum & ghuml: Generative 3d human shape and articulated pose models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6184–6193 (2020) [4](#)

67. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021) [2](#), [3](#), [4](#), [5](#), [13](#), [14](#)
68. Yu, T., Zheng, Z., Guo, K., Zhao, J., Dai, Q., Li, H., Pons-Moll, G., Liu, Y.: Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7287–7296 (2018) [1](#)
69. Zhao, H., Zhang, J., Lai, Y.K., Zheng, Z., Xie, Y., Liu, Y., Li, K.: High-fidelity human avatars from a single rgb camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15904–15913 (2022) [4](#)
70. Zheng, Y., Abrevaya, V.F., Bühler, M.C., Chen, X., Black, M.J., Hilliges, O.: I M Avatar: Implicit morphable head avatars from videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (Jun 2022) [4](#)
71. Zheng, Z., Huang, H., Yu, T., Zhang, H., Guo, Y., Liu, Y.: Structured local radiance fields for human avatar modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15893–15903 (2022) [4](#)
72. Zheng, Z., Yu, T., Liu, Y., Dai, Q.: Pamir: Parametric model-conditioned implicit representation for image-based human reconstruction. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021) [4](#)