

Hand Avatar: Free-Pose Hand Animation and Rendering from Monocular Video

Xingyu Chen Baoyuan Wang Heung-Yeung Shum
 Xiaobing.AI




Figure 1. Demonstration of HandAvatar. (a) Personalized hand rendering. From left to right: hand mesh, compositional occupancy, albedo, illumination, shaded appearance, and ground truth; (b) three groups of texture editing in terms of lighting, albedo, and shadow (by altering the self-occlusion effect of thumb); (c) free-pose hand animation and rendering.

Abstract

We present *HandAvatar*, a novel representation for hand animation and rendering, which can generate smoothly compositional geometry and self-occlusion-aware texture. Specifically, we first develop a *MANO-HD* model as a high-resolution mesh topology to fit personalized hand shapes. Sequentially, we decompose hand geometry into per-bone rigid parts, and then re-compose paired geometry encodings to derive an across-part consistent occupancy field. As for texture modeling, we propose a self-occlusion-aware shading field (*Self*). In *Self*, drivable anchors are paved on the *MANO-HD* surface to record albedo information under a wide variety of hand poses. Moreover, directed soft occupancy is designed to describe the ray-to-surface relation, which is leveraged to generate an illumination field for the disentanglement of pose-independent albedo and pose-dependent illumination. Trained from monocular video data, our *HandAvatar* can perform free-pose hand animation and rendering while at the same time achieving superior appearance fidelity. We also demonstrate that

HandAvatar provides a route for hand appearance editing. Project website: <https://seanchenxy.github.io/HandAvatarWeb>.

1. Introduction

Human avatars [5, 16, 19, 20, 27, 74] have been vigorously studied for years. However, there has been limited research that particularly focuses on hand avatars [9]. Due to the nature of distinctive properties (e.g., serious self-occlusion and contact) between the hand and the rest of the human parts (i.e., face, head, and body), it is essential to investigate a specialized representation tailored for modeling both the hand geometry and texture.

Traditional pipeline tends to adopt texture maps and colored mesh for hand appearance modeling [7, 11, 12, 24, 43], but developing an elaborate personalized hand mesh and texture map usually requires expensive scan data [54] and artistic knowledge. Recently, the neural rendering technique has gained raising attention, where neural radiance field (NeRF) [32] has been adapted to represent humans by predicting geometry and texture properties for an arbitrary

3D point query [9, 10, 20, 25, 37, 39, 40, 48, 51, 58, 62–64, 69, 72, 75]. Compared to the conventional mesh-texture pipeline, NeRF is cheap in training data collection and superior in rendering fidelity. Despite the huge success of human body and face modeling, neural rendering-based hand representation [9] remains much less explored. The hand is highly articulated such that the complex hand motion brings difficulties for neural rendering. Firstly, the deformation of hand geometry is hard to model. When coping with large and complex hand deformations (*e.g.*, self-contact), previous skinning-based methods can hardly find accurate skinning weights for an arbitrary query [3, 6, 18, 19, 31, 35, 39, 47, 62, 74], while part-aware methods usually suffer from across-part inconsistency issue [17, 21, 30, 60]. Secondly, hand texture is hard to model because of the highly articulated structure. For example, articulated hand motion induces serious self-occlusion so that different hand poses lead to noticeable variations in illumination and shadow patterns. Illumination is important for realistic rendering, but we are not aware of any prior work in estimating illumination caused by articulated self-occlusion.

Motivated by the above challenges, we propose HandAvatar for animatable realistic hand rendering. Considering different difficulties in geometry and texture modeling, we follow the idea of inverse graphics [76] to disentangle hand geometry, albedo, and illumination. At first, we employ explicit mesh to depict hand shapes. However, the popular hand mesh model, *i.e.*, MANO [44], only provides a coarse mesh with 778 vertices, whose shape fitting capacity is limited. Therefore, we design a super-resolution version of MANO with 12,337 vertices and 24,608 faces, namely MANO-HD, which can fit personalized hand shapes with per-vertex displacements. Additionally, massive existing MANO-annotated data can be seamlessly represented by MANO-HD. For introducing mesh-based hand shape to the volume rendering pipeline [32], we propose a local-pair occupancy field (PairOF), where every two part-level geometry encodings are reassembled according to physical connections to yield an across-part consistent field. As for hand texture, we propose a self-occlusion-aware shading field (Self). Self is comprised of an albedo field and an illumination field. The albedo field resorts to anchors that are uniformly paved on MANO-HD surfaces, each of which holds positional and albedo encodings to model a small hand region. The illumination field is to cope with articulated self-occlusion, where directed soft occupancy is designed to estimate illumination and shadow patterns.

MANO-HD and PairOF are pre-trained with MANO parameter annotations, then they cooperate with Self in end-to-end training on monocular video data. Finally, with hand pose as the input, our HandAvatar can perform hand animation and rendering. We evaluate our approach on the InterHand2.6M dataset [34] and achieve high-fidelity ge-

ometry and texture for free-pose hand animation. We also demonstrate that it is convenient to edit hand appearance in HandAvatar as shown in Fig. 1. Therefore, our main contributions are summarized as follows:

- We propose a HandAvatar framework, the first method for neural hand rendering with self-occluded illumination.
- We develop MANO-HD and a local-pair occupancy field that fit hand geometry with personalized shape details.
- We propose a self-occlusion-aware shading field that can render hand texture with faithful shadow patterns.
- Our framework is end-to-end developed for free-pose realistic hand avatars. Extensive evaluations indicate our method outperforms prior arts by a large margin.

2. Related Work

Articulated Human Geometry. Parametric human models [23, 28, 38, 44] have developed for years, where mesh can be inferred given pose and shape parameters. Specifically, as a common-used hand model, MANO [44] can produce a hand mesh with 778 vertices and 1,538 faces. This mesh template is too coarse so its representation capacity is largely limited. Gyeongsik *et al.* [33] proposed DeepHandMesh to generate dense and high-fidelity hand mesh, but brought restricted generalization as multi-view depth data was required for training. In contrast, our MANO-HD is a general high-resolution hand mesh model so that all existing MANO-annotated data can be seamlessly represented using MANO-HD. Meanwhile, MANO-HD can fit personalized hand shapes with monocular RGB video data.

Mesh suffers drawbacks of discontinuity and unalterable topology structure. To remedy this issue, recent research tends to explore implicit human geometry [2, 22, 31, 46, 67, 73], which has the advantages of flexibility and continuity. For example, GraspField [22] leveraged the signed distance field (SDF) to describe hand-object contact. However, implicit geometry is poor in free-pose animation when compared to explicit mesh, so the articulated driving of implicit human geometry is widely studied. As reported in [3, 35, 39, 62], a posed-space query can be transformed back to canonical space with linear blend skinning and inverse skinning weights. The inverse skinning paradigm fails to deal with self-contact, where a query can match multiple canonical-space points. Then, forward skinning deformation is designed to transform canonical-space points to posed space with an iterative root finding method [6, 19, 31, 47, 74], but the iterative optimization algorithm could hurt end-to-end network training. By and large, per-bone rigid transformations can compose a large motion space with the difficulty of optimizing accurate skinning weights for an arbitrary 3D point query. With the aid of parametric models [28, 44], another idea of deformation between posed and canonical spaces is to leverage the surface

motion [25, 53, 64]. For a query, a mesh-surface point is found according to Euclidean distance as the reference, then the query deformation is set the same as that of the reference point. Yuan *et al.* [68] and Garbin *et al.* [13] pointed out that the reference from triangular mesh is not accurate enough and proposed to use tetrahedral mesh. However, the deformation of tetrahedral mesh is hard to cooperate with popular human priors [23, 28, 44] and could be potentially slow [49]. Without requiring motion approximation, part-aware methods [17, 21, 30, 60] are developed by fusing part-wise geometries. NASA [17] divided the body into per-bone parts, then a query was deformed into each part space with unambiguous rigid transformation for decoding of part-level occupancy. NASA can describe complex deformations owing to accurate query motion, but information between body parts is ignored. To relieve this issue, COAP [30] encoded connected parts with PointNet [42], but it still incurred non-smooth part connections. Instead of part-wise modeling, we propose a part-pair-wise decoder to generate across-part consistent geometry.

Human Texture. Previously, the primary focus on hand texture is the texture map and colored mesh [7, 11, 12, 24, 43]. Although many high-quality texture maps are explored, the design of personalized texture maps usually requires expensive scan data [54] and artistic knowledge [11]. In contrast, LISA [9] employed radiance field [65] to learn hand appearance from multi-view images and introduced color parameters for texture generalization. Different from LISA, we design a monocular method for the convenience of data collection. Moreover, we pay attention to detailed personalized textures including albedo and illumination. Because of the aforementioned difficulty in implicit deformation, the learned texture on the human surface could be blurred [19, 62]. To enhance surface texture representation, local representations are developed with explicit mesh as the guidance. NeuralBody [40] attached latent codes to mesh vertices, which can diffuse into space with sparse convolution [15]. NeuMesh [8] also put color features on mesh vertices, and achieved an editable radiance field. Furthermore, mesh-guided local volume [27] and local radiance field [75] were designed. We follow the local modeling paradigm and uniformly place anchors on MANO-HD surface using barycentric sampling to trace local information.

Human Inverse Rendering. Most methods model human appearance with entangled geometry, albedo, and illumination [62]. Meanwhile, there has been a surge of interest in human inverse rendering, the purpose of which is to extract intrinsic components (*i.e.*, geometry, material, and illumination) from RGB data [36, 55–57, 76]. For example, GAN2X [36] designed an unsupervised framework to model albedo and specular properties of non-Lambertian material, then rendered face with Phong shading [41]. With a similar pur-

pose, HyFRIS-Net [76] disentangled albedo and illumination with an inverse 3DMM model to achieve a considerably improved quality of face rendering. S2HAND [7] simultaneously estimated camera pose, colored mesh, and lighting condition to form a photometric loss for hand pose estimation, but its rendering quality was coarse without a detailed appearance. Although the inverse rendering technique on the human face has been becoming a well-studied issue, the knowledge cannot be trivially transferred to hand tasks. Different from the face, the hand is characterized by articulated self-occlusion. Illumination and shadow caused by self-occlusion have not yet been discussed in prior works, and thus we fill this gap for hand inverse graphics.

Illumination in Radiance Field. The existing literature on NeRF-based illumination technique [4, 50, 52, 71] is to estimate source light condition or surface reflection property (*i.e.*, bidirectional reflectance distribution function, BRDF). For example, NeRV [50] took as input a set of images under known lighting to predict BRDF, and achieved novel-view rendering with arbitrary lighting conditions. NeLF [52] designed a lighting estimation module, and then performed face relighting. NeRF-OSR [45] collected multi-view outdoor images to predict albedo and shadow maps. In contrast to the prior art, we dedicate modeling illumination under the condition of articulated self-occlusion.

3. Method

Fig. 2 illustrates the overview pipeline of our HandAvatar system, including MANO-HD (Sec. 3.1), PairOF (Sec. 3.2), and Self (Sec. 3.3). Table 1 also provides the list of symbol notations and their definitions used in this paper.

3.1. MANO-HD

Mesh Subdivision. MANO [44] deforms hand mesh with shape parameter β and pose parameter $\theta \in \mathbb{R}^{B \times 3}$ ($B = 16$ indicates the number of per-bone parts). For lifting mesh resolution, we uniformly subdivide MANO template mesh by adding new vertices on edge middle points [16]. This operation increases the vertex number to 12,337 and the face amount to 24,608 (see Fig. 3). Then, the skinning weights of added vertices are given with the average of their seminal vertices. To eliminate artifacts during skinning, we optimize upsampled skinning weights for better dynamic performance. Please see the *suppl. material* for details.

Shape Fitting. Although MANO-HD has a high-resolution template, its shape is still bounded by β . Hence, when modeling personalized hand mesh, we get rid of β and use a multi-layer perceptron (MLP) to derive a refined shape $\tilde{\mathbf{V}} = \tilde{\mathbf{V}} + \mathcal{M}_{shape}([\mathcal{P}(\tilde{\mathbf{V}}), \theta])$, where $\tilde{\mathbf{V}}$, $\mathcal{P}(\cdot)$, $[\cdot]$ denote the MANO-HD template vertices, positional encoding, and concatenation. The MLP can be trained with monocular

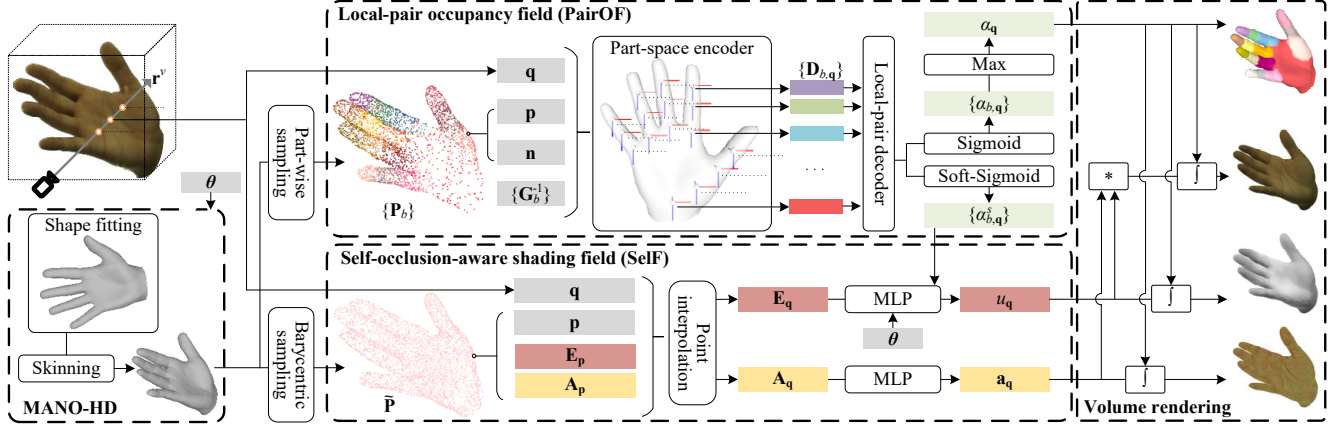


Figure 2. HandAvatar overview. Given hand pose, MANO-HD produces personalized mesh, while PariOF yields accordingly occupancy field. Self estimates albedo and illumination fields under self-occlusion. Then, hand appearance is synthesized by volume rendering.

θ	hand pose parameter	β	hand shape parameter
\mathbf{P}	point cloud	\mathbf{r}	ray direction
b	index of per-bone parts	\mathbf{n}	sampled point normal
\mathbf{q}	query point position	\mathbf{p}	sampled point position
\mathbf{G}_b	bone transformation matrix	$\mathbf{D}_{b,\mathbf{q}}$	part geometry encoding
\mathbf{E}	positional encoding	\mathbf{A}	albedo encoding
u	illumination value	\mathbf{a}	albedo value in RGB
α	occupancy value	α^s	soft occupancy value

Table 1. Symbolical notations.




Figure 3. Templates of MANO (left) and MANO-HD (right).

video data and IoU loss $\mathcal{L}_{shape} = 1 - \text{IoU}(\mathcal{D}(\mathcal{S}(\tilde{\mathbf{V}})), \mathbf{S}^*)$, where $\mathcal{D}, \mathcal{S}, \mathbf{S}^*$ are silhouette rendering [26], linear blend skinning, and ground-truth silhouette.

3.2. Local-Pair Occupancy Field

Given the query point $\mathbf{q} \in \mathbb{R}^3$, PairOF predicts the occupancy value $\alpha_{\mathbf{q}}$ to describe whether it locates in ($\alpha_{\mathbf{q}} > 0.5$) or out of ($\alpha_{\mathbf{q}} < 0.5$) the surface. Hence, the hand surface can be formulated as $\{\mathbf{q} | \alpha_{\mathbf{q}} = 0.5\}$.

Part-Space Encoder. Following NASA [17], we divide hand mesh into per-bone part meshes and uniformly sample N^p points on part mesh faces to obtain point clouds $\mathbf{P}_b = \{\mathbf{p} \in \mathbb{R}^3\}$ and normals $\mathbf{N}_b = \{\mathbf{n} \in \mathbb{R}^3\}$. Points, normals, and query are transferred back to part canonical spaces with $\hat{\mathbf{P}}_b = \mathbf{G}_b^{-1}\mathbf{P}_b$, $\hat{\mathbf{N}}_b = \mathbf{R}_b^{-1}\mathbf{N}_b$, $\hat{\mathbf{q}}_b = \mathbf{G}_b^{-1}\mathbf{q}$, where \mathbf{G} is bone transformation matrix and \mathbf{R} is the rotation component of \mathbf{G} . Following COAP [30], $\hat{\mathbf{P}}_b, \hat{\mathbf{N}}_b$ are fed to

a PointNet \mathcal{Q}_{part} to extract latent geometry features. The part geometry encoding is ultimately formulated by concatenating PointNet representation and canonicalized query, *i.e.*, $\mathbf{D}_{b,\mathbf{q}} = [\mathcal{Q}_{part}([\hat{\mathbf{P}}_b, \hat{\mathbf{N}}_b]), \hat{\mathbf{q}}_b]$.

Local-Pair Decoder. To explore inter-part relations, some previous works fuse per-bone features according to the kinematic tree to form structured representations [1, 14, 31]. Our intuition is that the local part shape is not related to the kinematic tree, instead, there is a strong relationship between parts. Hence, we define local pair as two parts that are physically connected. Then, we propose local-pair decoder \mathcal{Q}_{pair} based on PointNet to fuse each paired encodings and predict occupancy value:

$$\alpha_{b,\mathbf{q}} = \sigma(\max\{\mathcal{Q}_{pair}(\{\mathbf{D}_{b,\mathbf{q}}, \mathbf{D}_{b',\mathbf{q}}\}) | b' \in \mathbb{P}(b)\}) \quad (1)$$

where $\mathbb{P}(b)$ selects locally paired parts that has a physical connection with part b ; σ is the sigmoid function; $\alpha_{b,\mathbf{q}}$ is part-level occupancy value. Through the part-pair-wise decoding, the part boundaries become blurred. Therefore, the intuition behind the maximum operator in Eq. (1) is to yield a union of part-level geometries that extend to the connections. Finally, The global occupancy value is given by fusing part-level values, *i.e.*, $\alpha_{\mathbf{q}} = \max\{\alpha_{b,\mathbf{q}}\}_{b=1}^B$.

Pre-Training. With free MANO parameter annotations, we pre-train PairOF as a prior model to endow PairOF with prior knowledge of 3D hands. With a hand mesh inferred by MANO-HD, we sample point clouds with N^t points as training data [30]. The objective is to minimize the mean squared error between ground truth α^* and predicted occupancy values, *i.e.*, $\mathcal{L}_{PairOF} = \frac{1}{N^t} \sum_{\mathbf{q}} (\alpha_{\mathbf{q}} - \alpha_{\mathbf{q}}^*)^2$.

3.3. Self-Occlusion-Aware Shading Field

Given a query \mathbf{q} , Self predicts its albedo and illumination values, which are rendered with the volumetric method.




Figure 4. Illustration of self-occluded illumination along the kinematic tree of the forefinger. The deepest purple positions contribute to directed soft occupancy value.

Volume Rendering. For a ray casting on view direction \mathbf{r}^v , we uniformly sample N^q queries $\{\mathbf{q}_i\}_{i=1}^{N^q}$, each of which has occupancy α , albedo \mathbf{a} , and illumination value u . We render neural fields with the volumetric method [32]:

$$\mathbf{X}_{\mathbf{r}^v} = \sum_{i=1}^{N^q} \left(\prod_{j=1}^{i-1} (1 - \alpha_{\mathbf{q}_j}) \right) \alpha_{\mathbf{q}_i} \mathbf{X}_{\mathbf{q}_i}. \quad (2)$$

When \mathbf{X} equates to \mathbf{a} , u , or $u\mathbf{a}$, we obtain the albedo value, illumination value, or shaded RGB color of a pixel.

Albedo Field. Albedo describes the intrinsic color of the material, which is invariant *w.r.t.* hand pose, illumination, *etc.* Motivated by the invariant property, we fix anchors on the MANO-HD surface, whose relative geodesic locations are independent from varying poses. To this end, we uniformly sample point clouds $\tilde{\mathbf{P}}$ with N^a points on MANO-HD template mesh and represent them with barycentric coordinates. Compared to directly using vertex as the anchor [8, 27, 40], our barycentric anchors are more uniform to cover the hand surface. Then, we develop albedo encodings $\mathbf{A} \in \mathbb{R}^{N^a \times D^a}$ with random initialization and attach them to anchors. Given hand pose, anchors can be re-sampled based on deformed vertices and fixed barycentric coordinates to form deformed points clouds $\tilde{\mathbf{P}}$. For a query \mathbf{q} , we find N^n nearest points in $\tilde{\mathbf{P}}$ and interpolate \mathbf{A} using inverse Euclidean distances as the weights. Thereby, we obtain the albedo encoding $\mathbf{A}_{\mathbf{q}} \in \mathbb{R}^{D^a}$ and then fed it to an MLP to predict the albedo value $\mathbf{a}_{\mathbf{q}} \in \mathbb{R}^3$, *i.e.*, $\mathbf{a}_{\mathbf{q}} = \mathcal{M}_{albedo}(\mathbf{A}_{\mathbf{q}})$.

Directed Soft Occupancy. For self-occluded illumination estimation, we require the near-far relationship for a bone part. That is, the illumination of a query \mathbf{q} is affected by self-occlusion when \mathbf{q} is close to multiple parts. Although occupancy value can describe the relation between \mathbf{q} and parts, the value is nearly binary so can only depict inside-outside relations. Hence, a soft factor τ is introduced to the sigmoid function to soften the occupancy value:

$$\sigma^s(x) = \frac{1}{1+e^{-\tau x}}, \quad 0 < \tau < 1. \quad (3)$$

Soft occupancy $\alpha_{b,\mathbf{q}}^s$ is derived by replacing σ in Eq. (1) with σ^s . Further, instead of modeling a single query, we de-

sign *directed soft occupancy* to reflect the near-far relation between a ray casting and an articulated part. For a ray casting \mathbf{r} that can reach \mathbf{q} , the directed soft occupancy $\alpha_{b,\mathbf{q},\mathbf{r}}^s$ is defined as the maximal value on \mathbf{r} before the ray hits \mathbf{q} . For discretization, we uniformly sample queries $\{\mathbf{q}_i\}_{i=1}^{N^q}$ on a ray casting \mathbf{r} , and compute directed soft occupancy as

$$\alpha_{b,\mathbf{q},\mathbf{r}}^s = \max\{\alpha_{b,\mathbf{q}_i}^s | \mathbf{q}_i \leq \mathbf{q}\}, \quad (4)$$

where $\mathbf{q}_i \leq \mathbf{q}$ selects queries that the ray traverses before reaching \mathbf{q} . For example, $\alpha_{b_3,\mathbf{q},\mathbf{r}}^s$ equals to $\alpha_{b_3}^s$ of the deepest purple query in Fig. 4.

Illumination Field. It is well known that the illumination effects come with light-source distribution, irradiance, and reflectance. Independent from self-occlusion, reflectance is the material property, which is not our focus. Affected by self-occlusion, some ambient lighting rays could be occluded such that the irradiance could be changed. Thereby, the problem is formulated as estimating irradiance of an outside query \mathbf{q} ($\alpha_{b,\mathbf{q}}^s < 0.5$), which indicates the energy amount that can reach \mathbf{q} . To this end, we use the hand pose θ and query location as cues. Similar to the albedo encodings, positional encodings $\mathbf{E} = \mathcal{P}(\tilde{\mathbf{P}})$ are attached to anchors, and we obtain $\mathbf{E}_{\mathbf{q}}$ with interpolation as the surface-calibrated location of \mathbf{q} . Nevertheless, self-occlusion is quite intractable for θ and $\mathbf{E}_{\mathbf{q}}$, so we leverage directed soft occupancy to enhance the awareness of self-occlusion.

As shown in Fig. 4, the articulated structure prohibits a portion of energy from arriving \mathbf{q} . Apparently, the situation of energy occlusion around a ray direction is implied in a set of directed soft occupancy $\{\alpha_{b,\mathbf{q},\mathbf{r}}^s\}_{b=1}^B$. That is, if a ray casting is close to multiple parts before hitting \mathbf{q} , the illumination of \mathbf{q} shall be impacted by self-occlusion. Prohibited by a limited memory budget, we cannot consider spherically distributed ray directions, and thus the number of ray castings is imperative to be reduced. Our institution is that (1) the selected ray should be able to arrive \mathbf{q} (*i.e.*, $\alpha_{b,\mathbf{q},\mathbf{r}}^s < 0.5$) such that can elaborate the near-far relations for all articulated parts; (2) an articulated part can only affect the illumination around it, where the query is close to the part (*i.e.*, $\alpha_{b,\mathbf{q}}^s \rightarrow 0.5$). Meanwhile, we have $\alpha_{b,\mathbf{q},\mathbf{r}}^s \geq \alpha_{b,\mathbf{q}}^s$ from Eq. (4). Thereby, the variation caused by ray directions is minor, and we use $\{\alpha_{b,\mathbf{q},\mathbf{r}^v}^s\}_{b=1}^B$ as the guidance to estimate irradiance of \mathbf{q} , where \mathbf{r}^v is the view direction. Without introducing extra ray castings, we significantly reduce computational costs by leveraging the ray casting and queries on the view direction.

Finally, we use an MLP to predict the illumination value, *i.e.*, $u_{\mathbf{q}} = \mathcal{M}_{illum}([\theta, \mathbf{E}_{\mathbf{q}}, [\alpha_{b,\mathbf{q},\mathbf{r}^v}^s]_{b=1}^B])$.

Optimization. The training of Self is based on reconstruction loss functions, including LPIPS [70] loss and l_1 error between the rendered image \mathbf{C} and the ground truth \mathbf{C}^* , *i.e.*, $\mathcal{L}_{Self} = \mathcal{L}_{LPIPS}(\mathbf{C}, \mathbf{C}^*) + \|\mathbf{C} - \mathbf{C}^*\|_1$.

Method	#Param	Guided mesh	IoU (%) \uparrow	$Lap.$ \downarrow	$Cham.$ \downarrow
COAP [30]	138K	MANO	94.08	2.371	8.564
COAP	138K	MANO-HD	94.01	2.348	7.694
COAP*	287K	MANO-HD	95.08	2.339	7.405
PairOF \dagger	256K	MANO-HD	96.06	2.288	6.998
PairOF (ours)	237K	MANO-HD	96.32	2.281	7.151

Table 2. Effects of MANO-HD and PairOF. *: w/ wider-MLP decoder; \dagger : w/ Transformer-based decoder.



Figure 5. Effects of MANO-HD and PairOF. (a) COAP w/ MANO. (b) COAP w/ MANO-HD. (c) PairOF w/ MANO-HD.

4. Experiments

4.1. Implementation Details and Metrics

Pre-Training of PairOF. We adopt all right-hand annotations in InterHand2.6M [34] for pre-training, whose training/test set contains 875,530/565,611 samples. Learnable parameters includes Q_{part} , Q_{pair} . We set $N^t = N^p = 256$, and the training process is to minimize \mathcal{L}_{PairOF} .

End-to-End Training. With personalized monocular video, we optimize \mathcal{M}_{shape} , Q_{pair} , \mathcal{M}_{albedo} , \mathcal{M}_{illum} and \mathbf{A} in an end-to-end manner. Video data are selected from InterHand2.6M dataset [34]. The objective is to minimize $\mathcal{L}_{shape} + \mathcal{L}_{PairOF} + \mathcal{L}_{Self}$. Hyperparameters in Self are set as $N^q = 64$, $N^a = 4096$, $N^n = 4$, $D^a = 128$, $\tau = 0.05$. The rendering resolution is 256×256 . Please see the *suppl. material* for data selection, pre-processing, and more training details.

Metrics Following COAP [30], IoU is used to evaluate the occupancy field. We also employ Laplacian smooth ($Lap.$) and Chamfer distance ($Cham.$) to evaluate the mesh quality, the latter of which is formulated as the minimal distance between the vertices extracted from occupancy field [29] and the guided mesh faces. $Lap.$ and $Cham.$ are presented in 10^{-4} m. Consistent with HumanNeRF [62], we report LPIPS [70], PSNR, and SSIM [61] to reflect image




Figure 6. Attention map for fusing part-wise geometry encodings. The rows with red indices contribute to global occupancy value.

similarity as the metrics of rendering quality. All evaluation data are with novel poses that are unseen in training.

4.2. Evaluation on Geometry Performance

Comparison with Prior Arts PairOF and COAP [30] use the same encoder but different decoders, so their comparison can reveal the effect of our part-pair-wise decoding. The results of COAP are from the officially released code, and we re-train models on InterHand2.6M dataset. Because our local-pair decoder is larger than that of COAP, we enlarge MLP width for comparable model size (denoted as COAP*). Referring to Table 2, an occupancy field guided by MANO-HD has a smoother surface (lower $Lap.$) and higher fidelity (lower $Cham.$), so there are benefits of MANO-HD over MANO in guiding an implicit function. Moreover, PairOF can improve all metrics by a large margin. As shown in Fig. 5, MANO-HD can improve overall smoothness, while PairOF exhaustively eliminates non-smooth part connections to achieve cross-part consistency.

Comparison with Transformer-Based Decoder To verify the local-pair prior knowledge in PairOF, the Transformer [59] technique is employed as the decoder, where self-attention can adaptively fuse part-wise geometry without inductive prior. As shown in Table 2, the local-pair decoder performs on par with the Transformer-based decoder. To unveil the effect of self-attention, we delve deep into attention-based feature fusion based on two representative hand poses (*i.e.*, flat and fist poses). At first, we extract mesh vertices from the occupancy field, each of which comes with respective attention maps. Then, we find the part that has the maximal occupancy value for each vertex and gather vertices into groups accordingly. Each group of vertices can reflect the property of a bone part, and we show their average attention maps in Fig. 6. Because of the maximum operator in part-wise geometry fusion, only one part contributes

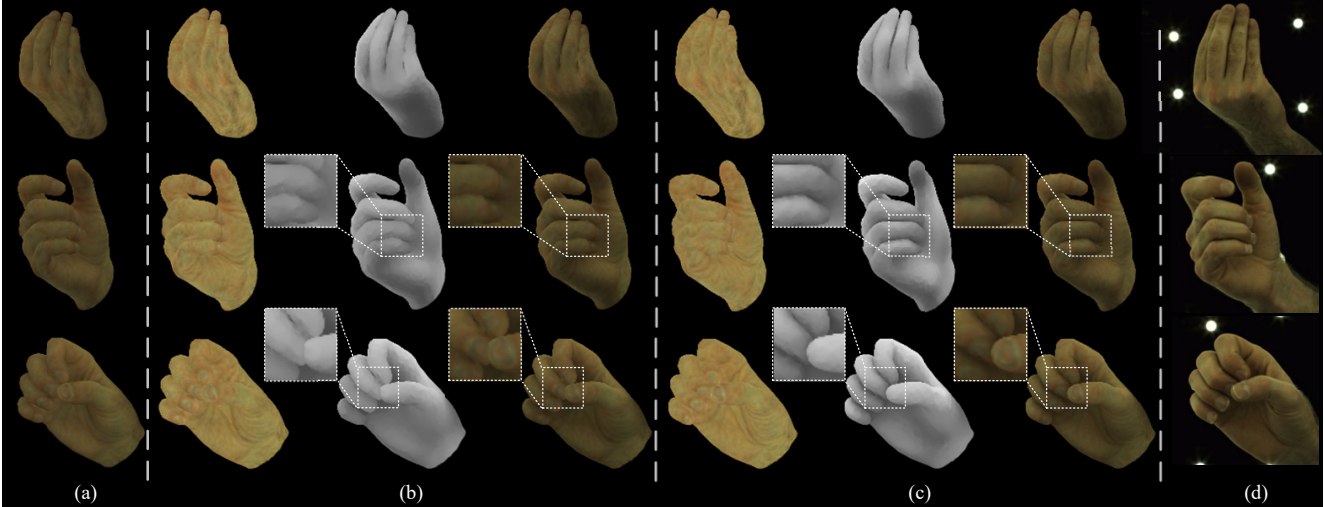


Figure 7. Effects of the disentangled albedo and illumination fields in Self. (a) Coupled albedo and illumination. (b,c) Disentangled albedo and illumination; directed soft occupancy is not involved in (b); from left to right: albedo, illumination, shaded image. (d) Ground truth.

to the global occupancy value. That is, for part b , we should focus on the b th row (red number in Fig. 6). Consequently, attention-based fusion is consistent with our local-pair design. For example, the attention map fuses parts b_0, b_{10}, b_{11} to evolve the encoding of part b_{10} . Therefore, the design concept of the local-pair decoder is evident. Nevertheless, the Transformer-based decoder is not efficient enough because the attention map contains *vertical* patterns instead of *diagonal* ones. That is, meaningless computations are introduced by the Transformer, despite they do not contribute to global occupancy. For example, referring to “fist pose” and “bone part 12” in Fig. 6, the attention map integrates encodings of parts b_0, b_{12} (instead of b_1) for the prediction of part b_1 . The reason behind this is that the *inside* property is exclusively enhanced, and vertices belonging to part b_{12} also have inside properties to part b_0 under the fist pose. Refer *suppl. material* for more details and part indices.

4.3. Evaluation on Rendering Quality

Ablation Study on Shape Fitting We fit personalized hand shape with \mathcal{M}_{shape} . For comparison, \mathcal{M}_{shape} is replaced with β -based shape fitting, where β is the annotation in InterHand2.6M dataset. As shown in Table 3, β -based shape induces poor rendering quality, while our method brings a significant improvement. Therefore, the shape fitting capacity of our proposed MANO-HD is confirmed.

Ablation Study on Self Referring to Fig. 7(a) and (b), the disentanglement of albedo and illumination fields can improve the rendering reality by introducing shadow patterns. Moreover, directed soft occupancy can further elevate the illumination representation through ray-based occlusion estimation. Referring to Fig. 7(c), it is remarkable that the

Shape fit.	Illum.	Dir. occ.	LPIPS ↓	PSNR ↑	SSIM ↑
			0.1268	26.53	0.8692
✓			0.1113	27.32	0.8830
✓	✓		0.1063	28.02	0.8903
✓	✓	✓	0.1035	28.23	0.8941

Table 3. Effects of shape fitting (Shape fit.), illumination field (Illum.), and directed soft occupancy (Dir. occ.) on rendering quality. Results are from InterHand2.6M *test/Capture0*.

shadow patterns on the palm and fingers are more faithful with fewer artifacts when compared to Fig. 7(b). In Table 3, the illumination field and directed soft occupancy lead to quantitative improvements in rendering metrics, indicating our Self is conducive to realistic rendering.

Comparison with Prior Arts We compare HandAvatar with previous monocular methods HumanNeRF [62] and SelfRecon [19], both of which are from officially released codes and re-trained on InterHand2.6M dataset. SelfRecon uses the surface-based rendering [66] method, and its representation of texture detail is not good enough, as shown in Fig. 8. HumanNeRF and our method leverage the volume rendering method [32], which can produce realistic hand texture. However, limited by inverse skinning deformation, HumanNeRF cannot cope with self-contact that commonly occurs in hand animation. As shown in Fig. 8, HumanNeRF has corrupted geometry when fingers contact with each other. In contrast, benefiting from our PairOF, HandAvatar has the advantage of free-pose animation while at the same time maintaining geometry fidelity. In addition, both SelfRecon and HumanNeRF employ entangled albedo and illumination for color prediction, so the shadow on hand is hard to be aware of, as shown in the 3rd-8th

Method	<i>test/Capture0</i>			<i>test/Capture1</i>			<i>val/Capture0</i>		
	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑
SelfRecon [19]	0.1421	26.38	0.8786	0.1389	25.18	0.8758	0.1490	25.78	0.8687
HumanNeRF [62]	0.1145	27.64	0.8836	0.1177	26.31	0.8803	0.1192	27.80	0.8816
ours	0.1035	28.23	0.8941	0.1076	26.56	0.8902	0.1062	28.04	0.8900

Table 4. Rendering quality comparison among our HandAvatar and prior arts on the InterHand2.6M dataset.



Figure 8. Visualization results of SelfRecon [19], HumanNeRF [62], and our HandAvatar on free-pose animation and rendering.

rows of Fig. 8. On the contrary, our HandAvatar leads to a superiority in modeling illumination and shadow patterns caused by articulated self-occlusion. This is attributed to the disentanglement of albedo and illumination brought by




Figure 9. Visualization results of LISA [9] and our HandAvatar. The result of LISA is cropped from their original paper.

our proposed Self. Further, referring to the white arrows in Fig. 8, we can present varying shadow intensities for different self-occlusion situations, thanks to our directed soft occupancy. For quantitative comparison, we train models on three video sequences and achieve the best results in all metrics, as shown in Table 4.

As the most related work to this paper, LISA [9] is trained on non-released multi-view data, and its models/codes remain unavailable. Thereby, we compare LISA based on the result reported in their original paper. As shown in Fig. 9, LISA has difficulty in capturing accurate hand pose with a learnable skinning-based deformation. Besides the faithful shape and pose reconstruction, our rendered texture details are more realistic than that of LISA.

In addition to rendering fidelity, HandAvatar also provides a route for appearance editing as shown in Fig. 1.

5. Conclusions

In this work, we present a novel hand representation called HandAvatar for free-pose animation and rendering. First, we extend MANO to MANO-HD as a high-resolution topology structure to improve the shape-fitting capacity of hand mesh. Subsequently, PairOF with a local-pair decoder is developed, which can generate an across-part consistent occupancy field. Furthermore, we propose Self, the first approach to model hand texture under articulated self-occlusion, to disentangle hand albedo and illumination. Extensive experiments demonstrate our superior results on free-pose hands animation and rendering. We believe our method paves a new way for dynamic hand representation.

Limitation and future works For affordable computational costs, we use the directed soft occupancy on view direction to estimate the irradiance. This could lead to view-direction-dependent shadow patterns. Thus, the improved illumination field is worthy of ongoing exploration.

References

- [1] Emre Aksan, Manuel Kaufmann, and Otmar Hilliges. Structured prediction helps 3D human motion modelling. In *ICCV*, 2019.
- [2] Thiemo Alldieck, Hongyi Xu, and Cristian Sminchisescu. imGHUM: Implicit generative models of 3d human shape and articulated pose. In *ICCV*, 2021.
- [3] Bharat Lal Bhatnagar, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. LoopReg: Self-supervised learning of implicit surface correspondences, pose and shape for 3d human mesh registration. In *NeurIPS*, 2020.
- [4] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. NeRD: Neural reflectance decomposition from image collections. In *ICCV*, 2021.
- [5] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhoefer, Shun-Suke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shou-I Yu, et al. Authentic volumetric avatars from a phone scan. *ACM TOG*, 2022.
- [6] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021.
- [7] Yujin Chen, Zhigang Tu, Di Kang, Linchao Bao, Ying Zhang, Xuefei Zhe, Ruizhi Chen, and Junsong Yuan. Model-based 3D hand reconstruction via self-supervised learning. In *CVPR*, 2021.
- [8] Chong Bao and Bangbang Yang, Zeng Junyi, Bao Hujun, Zhang Yinda, Cui Zhaopeng, and Zhang Guofeng. NeuMesh: Learning disentangled neural mesh-based implicit field for geometry and texture editing. In *ECCV*, 2022.
- [9] Enric Corona, Tomas Hodan, Minh Vo, Francesc Moreno-Noguer, Chris Sweeney, Richard Newcombe, and Lingni Ma. LISA: Learning implicit shape and appearance of hands. In *CVPR*, 2022.
- [10] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. GRAM: Generative radiance manifolds for 3D-aware image generation. In *CVPR*, 2022.
- [11] Daiheng Gao, Yuliang Xiu, Kailin Li, Lixin Yang, Feng Wang, Peng Zhang, Bang Zhang, Cewu Lu, and Ping Tan. DART: Articulated hand model with diverse accessories and rich textures. In *NeurIPS*, 2022.
- [12] Daiheng Gao, Xindi Zhang, Xingyu Chen, Andong Tan, Bang Zhang, Pan Pan, and Ping Tan. CycleHand: Increasing 3D pose estimation ability on in-the-wild monocular image through cyclic flow. In *ACM MM*, 2022.
- [13] Stephan J Garbin, Marek Kowalski, Virginia Estellers, Stanislaw Szymanowicz, Shideh Rezaeifar, Jingjing Shen, Matthew Johnson, and Julien Valentin. VolTeMorph: Real-time, controllable and generalisable animation of volumetric representations. *arXiv:2208.00949*, 2022.
- [14] Georgios Georgakis, Ren Li, Srikrishna Karanam, Terrence Chen, Jana Košecká, and Ziyang Wu. Hierarchical kinematic human mesh recovery. In *ECCV*, 2020.
- [15] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3D semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 2018.
- [16] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular RGB videos. In *CVPR*, 2022.
- [17] Timothy Jeruzalski, Boyang Deng, Mohammad Norouzi, John P Lewis, Geoffrey Hinton, and Andrea Tagliasacchi. NASA: Neural articulated shape approximation. In *ECCV*, 2020.
- [18] Timothy Jeruzalski, David IW Levin, Alec Jacobson, Paul Lalonde, Mohammad Norouzi, and Andrea Tagliasacchi. NiLBS: Neural inverse linear blend skinning. *arXiv:2004.05980*, 2020.
- [19] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. SelfRecon: Self reconstruction your digital avatar from monocular video. In *CVPR*, 2022.
- [20] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. NeuMan: Neural human radiance field from a single video. *arXiv:2203.12575*, 2022.
- [21] Korrawe Karunratanakul, Adrian Spurr, Zicong Fan, Otmar Hilliges, and Siyu Tang. A skeleton-driven neural occupancy representation for articulated hands. *arXiv:2109.11399*, 2021.
- [22] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *3DV*, 2020.
- [23] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.*, 36(6), 2017.
- [24] Yuwei Li, Longwen Zhang, Zesong Qiu, Yingwenqi Jiang, Nianyi Li, Yuexin Ma, Yuyao Zhang, Lan Xu, and Jingyi Yu. NIMBLE: A non-rigid hand model with bones and muscles. *ACM TOG*, 2022.
- [25] Lingjie Liu, Marc Habermann, Viktor Rudnev, Kripasindhu Sarkar, Jiatao Gu, and Christian Theobalt. Neural actor: Neural free-view synthesis of human actors with pose control. *ACM TOG*, 2021.
- [26] Shichen Liu, Tianye Li, Weikai Chen, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *ICCV*, 2019.
- [27] Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *ACM TOG*, 2021.
- [28] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015.
- [29] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.
- [30] Marko Mihajlovic, Shunsuke Saito, Aayush Bansal, Michael Zollhoefer, and Siyu Tang. COAP: Compositional articulated occupancy of people. In *CVPR*, 2022.
- [31] Marko Mihajlovic, Yan Zhang, Michael J Black, and Siyu Tang. LEAP: Learning articulated occupancy of people. In *CVPR*, 2021.
- [32] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF:

- Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [33] Gyeongsik Moon, Takaaki Shiratori, and Kyoung Mu Lee. DeepHandMesh: A weakly-supervised deep encoder-decoder framework for high-fidelity hand mesh modeling. In *ECCV*, 2020.
- [34] Gyeongsik Moon, Shoou-I Yu, He Wen, Takaaki Shiratori, and Kyoung Mu Lee. Interhand2.6M: A dataset and baseline for 3D interacting hand pose estimation from a single RGB image. In *ECCV*, 2020.
- [35] Atsuhiko Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, 2021.
- [36] Xingang Pan, Ayush Tewari, Lingjie Liu, and Christian Theobalt. GAN2X: Non-lambertian inverse rendering of image gans. In *3DV*, 2022.
- [37] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, pages 5865–5874, 2021.
- [38] 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*, pages 10975–10985, 2019.
- [39] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021.
- [40] 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.
- [41] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.
- [42] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017.
- [43] Neng Qian, Jiayi Wang, Franziska Mueller, Florian Bernard, Vladislav Golyanik, and Christian Theobalt. HTML: A parametric hand texture model for 3D hand reconstruction and personalization. In *ECCV*, 2020.
- [44] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM TOG*, 2017.
- [45] Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. NeRF for outdoor scene relighting. In *ECCV*, 2022.
- [46] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019.
- [47] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, 2021.
- [48] Ruizhi Shao, Hongwen Zhang, He Zhang, Mingjia Chen, Yan-Pei Cao, Tao Yu, and Yebin Liu. DoubleField: Bridging the neural surface and radiance fields for high-fidelity human reconstruction and rendering. In *CVPR*, 2022.
- [49] Olga Sorkine and Marc Alexa. As-rigid-as-possible surface modeling. In *Symposium on Geometry processing*, 2007.
- [50] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. NeRV: Neural reflectance and visibility fields for relighting and view synthesis. In *CVPR*, 2021.
- [51] Shih-Yang Su, Frank Yu, Michael Zollhöfer, and Helge Rhodin. A-NeRF: Articulated neural radiance fields for learning human shape, appearance, and pose. In *NeurIPS*, 2021.
- [52] Tiancheng Sun, Kai-En Lin, Sai Bi, Zexiang Xu, and Ravi Ramamoorthi. NeLF: Neural light-transport field for portrait view synthesis and relighting. In *Eurographics Symposium on Rendering*, 2021.
- [53] Gusi Te, Xiu Li, Xiao Li, Jinglu Wang, Wei Hu, and Yan Lu. Neural capture of animatable 3d human from monocular video. In *ECCV*, 2022.
- [54] Ten24. <https://ten24.info>.
- [55] Ayush Tewari, Florian Bernard, Pablo Garrido, Gaurav Bharaj, Mohamed Elgharib, Hans-Peter Seidel, Patrick Pérez, Michael Zollhofer, and Christian Theobalt. FML: Face model learning from videos. In *CVPR*, 2019.
- [56] Ayush Tewari, Hans-Peter Seidel, Mohamed Elgharib, Christian Theobalt, et al. Learning complete 3D morphable face models from images and videos. In *CVPR*, 2021.
- [57] Ayush Tewari, Michael Zollhofer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Christian Theobalt. MoFA: Model-based deep convolutional face autoencoder for unsupervised monocular reconstruction. In *ICCV*, 2017.
- [58] Edgar Tretschk, Ayush Tewari, Vladislav Golyanik, Michael Zollhöfer, Christoph Lassner, and Christian Theobalt. Non-rigid neural radiance fields: Reconstruction and novel view synthesis of a dynamic scene from monocular video. In *ICCV*, 2021.
- [59] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.
- [60] Shaofei Wang, Andreas Geiger, and Siyu Tang. Locally aware piecewise transformation fields for 3d human mesh registration. In *CVPR*, 2021.
- [61] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004.
- [62] 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.
- [63] Hongyi Xu, Thiemo Alldieck, and Cristian Sminchisescu. H-NeRF: Neural radiance fields for rendering and temporal reconstruction of humans in motion. *NeurIPS*, 2021.
- [64] Tianhan Xu, Yasuhiro Fujita, and Eiichi Matsumoto. Surface-aligned neural radiance fields for controllable 3D human synthesis. In *CVPR*, 2022.

- [65] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *NeurIPS*, 2021.
- [66] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. In *NeurIPS*, 2020.
- [67] Tarun Yenamandra, Ayush Tewari, Florian Bernard, Hans-Peter Seidel, Mohamed Elgharib, Daniel Cremers, and Christian Theobalt. i3DMM: Deep implicit 3D morphable model of human heads. In *CVPR*, 2021.
- [68] Yu-Jie Yuan, Yang-Tian Sun, Yu-Kun Lai, Yuewen Ma, Rongfei Jia, and Lin Gao. NeRF-Editing: geometry editing of neural radiance fields. In *CVPR*, 2022.
- [69] Jiakai Zhang, Xinhang Liu, Xinyi Ye, Fuqiang Zhao, Yanshun Zhang, Minye Wu, Yingliang Zhang, Lan Xu, and Jingyi Yu. Editable free-viewpoint video using a layered neural representation. *ACM TOG*, 2021.
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [71] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. NeRFactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM TOG*, 2021.
- [72] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. HumanNeRF: Efficiently generated human radiance field from sparse inputs. In *CVPR*, 2022.
- [73] Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. Imface: A nonlinear 3D morphable face model with implicit neural representations. In *CVPR*, 2022.
- [74] Yufeng Zheng, Victoria Fernández Abrevaya, Marcel C Bühler, Xu Chen, Michael J Black, and Otmar Hilliges. I M Avatar: Implicit morphable head avatars from videos. In *CVPR*, 2022.
- [75] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, 2022.
- [76] Yongjie Zhu, Chen Li, Si Li, Boxin Shi, and Yu-Wing Tai. Hybrid face reflectance, illumination, and shape from a single image. *IEEE TPAMI*, pages 5002–5015, 2021.