

## HAVE-FUN: Human Avatar Reconstruction from Few-Shot Unconstrained Images

Xihe Yang<sup>1,2,†\*</sup> Xingyu Chen<sup>1,✉\*</sup> Daiheng Gao<sup>4</sup> Shaohui Wang<sup>1,5,†</sup>

Xiaoguang Han<sup>2,3</sup> Baoyuan Wang<sup>1,✉</sup>

<sup>1</sup>Xiaobing.AI <sup>2</sup>SSE, CUHKSZ <sup>3</sup>FNii, CUHKSZ <sup>4</sup>Freelancer <sup>5</sup>Tsinghua University

<https://seanchenxy.github.io/HaveFunWeb/>




Figure 1. Given a few images with various viewpoints and articulated poses, our approach can reconstruct an animatable human avatar.

### Abstract

As for human avatar reconstruction, contemporary techniques commonly necessitate the acquisition of costly data and struggle to achieve satisfactory results from a small number of casual images. In this paper, we investigate this task from a few-shot unconstrained photo album. The reconstruction of human avatars from such data sources is challenging because of limited data amount and dynamic articulated poses. For handling dynamic data, we integrate a skinning mechanism with deep marching tetrahedra (DMTet) to form a drivable tetrahedral representation, which drives arbitrary mesh topologies generated by the DMTet for the adaptation of unconstrained images. To effectively mine instructive information from few-shot data, we devise a two-phase optimization method with few-shot reference and few-shot guidance. The former focuses on aligning avatar identity with reference images, while the latter aims to generate plausible appearances for unseen regions. Overall, our framework, called HaveFun, can under-

take avatar reconstruction, rendering, and animation. Extensive experiments on our developed benchmarks demonstrate that HaveFun exhibits substantially superior performance in reconstructing the human body and hand.

### 1. Introduction

Human avatar reconstruction has experienced rapid development in recent years and shown great potential for applications of AR/VR, metaverse, etc [5, 68]. One of the challenges in this field is the data acquisition. Previous works typically require an expensive setup for multi-view RGB video [53], textured scan video [63], or multi-view static images [60]. Recently, there is a tendency to utilize easily accessible data sources for this task, such as monocular RGB video [74] or image set [75].

This leads to a pivotal question: “Is it possible to leverage a cheaper data source to reconstruct human avatars?” Intuitively, the cheaper data should be characterized by its limited quantity and the unconstrained nature of human behaviors it captures. As this paper will demonstrate, the answer leans toward the affirmative, and we refer to such a data source as few-shot unconstrained images. These data

\*Equal contribution; ✉ Corresponding author.

†This work was done when Xihe Yang and Shaohui Wang were interns at Xiaobing.AI, led by Xingyu Chen.

can be obtained from either a personal photo album or key frames of a video capture. Thereby, we are inspired to explore the human avatar reconstruction task with the few-shot unconstrained images.

From a technical perspective, we carefully design a 3D representation for the difficulties caused by the aforementioned data setting. When it comes to modeling dynamic humans, a popular idea is the use of a dynamic neural radiance field (NeRF) [49]. Nevertheless, despite various dynamic designs being reported [4, 55, 74, 79], accurately driving a volume space using limited data remains challenging. Recently, deep marching tetrahedra (DMTet) [64] proposes a hybrid method to produce a triangle mesh through a differentiable process. Due to the ease of mesh deformation, we are motivated to model dynamic humans with the DMTet representation. Specifically, we integrate skinning weights and blendshapes defined by SMPLX [51] with the static DMTet to create a drivable tetrahedral representation for adapting the unconstrained data. Furthermore, conducive to the few-shot task, the driveable representation allows the use of canonical SMPLX shape as the prior to make the tetrahedral grid well-initialized. In this manner, we extend the traditional static-scene DMTet to model the articulated human body under a dynamic condition.

In terms of learning from a limited amount of data, we employ the score distillation sampling (SDS) [54, 70] technique to generate plausible textures for unseen regions. Different from existing SDS-based image-to-3D tasks that rely on textual descriptions or captions [48, 67], we directly utilize the image as the prompt [44] so that faithful visual feature can be preserved. We refer to the SDS-based optimization as a few-shot guidance. Additionally, we incorporate traditional reconstruction optimization, namely few-shot reference, in our pipeline.

Our overall framework is called HaveFun, **H**uman **A**Avatar **r**Econstruction from **F**ew-shot **U**nconstrained images. Note that the human body and hand have distinct properties. The body is characterized by intricate geometry (*e.g.*, hair/cloth) and remarkable facial features, while the hand exhibits smooth bare geometry and subtle palm wrinkles. To evaluate our framework for the human body and hand, we develop benchmarks for them with the assistance of XHumans [63] and DART [16]. Remarkably, HaveFun effectively addresses both scenarios of the human body and hand. As a result, our approach can reconstruct human avatars with few-shot (as few as 2) dynamic images and achieve realistic rendering quality. Besides, we can perform avatar animation in various unseen human poses.

Our main contributions are summarized as follows:

- We propose a novel framework, termed HaveFun, to solve the challenging problem of human avatar reconstruction from few-shot unconstrained images.
- We explore a drivable tetrahedral representation for ar-

ticulated human motion and an SDS loss for non-static human reconstruction.

- We develop benchmarks for the few-shot dynamic body/hand reconstruction task. Extensive evaluations indicate our method outperforms previous one-shot [26] or video-based [30] approaches by a large margin.

We believe our endeavors would enhance the practical significance of this research area, paving a new way for human avatar reconstruction and real-world applications.

## 2. Related Work

**Avatar reconstruction of the human body.** Prominent techniques for human body reconstruction necessitate expensive data acquisition, *e.g.*, multi-view RGB video [9, 20, 27, 28, 38, 39, 52, 53, 62, 71, 79, 84, 86, 87], monocular RGB video [7, 18, 19, 30, 33, 34, 69, 74, 81, 85], textured scan video [21, 63], or image set [75]. Also, many of them utilize NeRF or its variations as the 3D representation and craft a motion field to connect the gap between body articulation and the canonical NeRF space. For example, HumanNeRF [74] created a personalized avatar by training an inverse skinning field using monocular video data. LISA [9] utilized multi-view video data to learn hand appearance and incorporated a multi-layer perceptron (MLP) to forecast skinning weights for hand animation. PersonNeRF [75] gathered hundreds of images of a human individual and generated an avatar with disentangled attributes. Differing from the majority of prior studies, our focus lies in addressing the few-shot body reconstruction challenge. Moreover, we explore articulation-friendly tetrahedral grid as the 3D representation.

**Few-shot human avatar creation.** Many research efforts rely on pre-trained generative models [3, 6, 12, 14, 32, 66, 76] and accomplish one-shot reconstructions through GAN inversion techniques [13, 57]. This line of research typically requires only a single image to recover a latent representation that aligns with the ground truth. Nevertheless, these pipelines often struggle to accurately represent the data that falls outside the GAN distribution [1, 72]. Instead of GAN inversion, another approach involves directly extracting image features and predicting a pixel-aligned implicit field to represent the human [10, 60, 61, 77, 78]. As a groundbreaking work, PIFu [60] employed MLPs to model both the occupancy value and color of the human body from one or several images. Despite the few-shot setting in the inference phase, the training of pixel-aligned methods still demands a large-size image set. Further, they treat the human as a static scene, neglecting the dynamic nature of the human body. By contrast, this paper can handle few-shot dynamic human images without the need for additional auxiliary data collection.

**Sparse-view 3D reconstruction.** When it comes to few-shot reconstruction, traditional methods typically optimize a NeRF using sparse-view data with the aid of geometry regularization, semantic consistency, depth supervision, *etc.* [11, 29, 35, 46, 50, 80]. For example, RegNeRF [50] introduced a patch regularizer to mitigate geometry artifacts and employed a log-likelihood model to ensure multi-view appearance consistency. In contrast to this line, we utilize a pre-trained diffusion model to regularize unseen textures from novel viewpoints. Moreover, in contrast to static scenes typically handled by related works, our approach enables dynamic body reconstruction using sparse-view data.

**Avatar creation with text-based priors.** Recently, text-to-3D task has gained popularity, thanks to the pre-trained language-vision models [22, 56]. Many reports have achieved remarkable performance using the diffusion model [41, 48, 54, 67, 70]. As a pioneering effort, DreamFusion [54] proposed SDS loss to optimize a NeRF with text prompts. Furthermore, text-guided strategies have been incorporated into the field of human avatar creation [2, 23–26, 31, 36, 40, 82]. These approaches allow for the generation of a realistic 3D human appearance that aligns with text semantics. While the text-based paradigm is effective in creating famous characters, it is more challenging when it comes to reconstructing actual human individuals. For example, TeCH [26] accomplished one-shot human avatar reconstruction with 5 stages of VQA [37] caption, DreamBooth [59] fine-tuning, geometry optimization, geometry post-processing, and texture optimization. In contrast, our method can be trained in an end-to-end manner. Hence, we believe our approach is inherently more elegant, precise, and accurate than text-based methods for the reconstruction task since we directly employ image features as guidance without the potential ambiguity caused by image captions.

### 3. Approach

Given a personalized unconstrained photo album  $\mathcal{I} = \{\mathbf{I}_i\}_{i=1}^N$  ( $N \leq 8$ ), this paper aims to reconstruct a 3D representation  $\mathcal{G}$  for free-viewpoint rendering and free-pose animation.  $\mathcal{G}$  takes explicit viewpoint  $\mathbf{R} \in \mathbb{R}^{d_{\mathbf{R}}}$ , human articulated poses  $\boldsymbol{\theta} \in \mathbb{R}^{d_{\boldsymbol{\theta}}}$ , and expression coefficients  $\boldsymbol{\psi} \in \mathbb{R}^{d_{\boldsymbol{\psi}}}$  as the input and generates a 2D image  $\hat{\mathbf{I}}$ :

$$\mathcal{G} : (\mathbf{R}, \boldsymbol{\theta}, \boldsymbol{\psi}) \in \mathbb{R}^{d_{\mathbf{R}}} \times \mathbb{R}^{d_{\boldsymbol{\theta}}} \times \mathbb{R}^{d_{\boldsymbol{\psi}}} \rightarrow \hat{\mathbf{I}} \in \mathbb{R}^{H \times W \times 3}, \quad (1)$$

where  $H = W = 512$  denote the image height and width. The 3D representation  $\mathcal{G} = \{\mathcal{M}, \mathcal{C}, \mathbf{W}, \mathbf{E}\}$  includes a triangular mesh  $\mathcal{M}$ , a texture field  $\mathcal{C}$ , skinning weights  $\mathbf{W}$ , and expression blendshapes  $\mathbf{E}$ . As shown in Fig. 2, we build a drivable tetrahedral representation as the core of  $\mathcal{G}$  to produce  $\{\mathcal{M}, \mathcal{C}, \mathbf{W}, \mathbf{E}\}$  (Sec. 3.2). Additionally, we design two phases of few-shot reference and few-shot guidance to train our framework (Sec. 3.3). We describe each part in detail below.

### 3.1. Preliminaries

**Deep marching tetrahedra.** DM Tet is a hybrid representation for 3D geometry, denoted as  $(\mathbf{V}^t, \mathbf{T})$ , where  $\mathbf{V}^t, \mathbf{T}$  are vertices and the tetrahedral indices, respectively. Each tetrahedron  $\mathbf{t} \in \mathbf{T}$  is represented by four vertices  $\{\mathbf{v}_a^t, \mathbf{v}_b^t, \mathbf{v}_c^t, \mathbf{v}_d^t\}$ . Each vertex  $\mathbf{v}^t = (x, y, z) \in \mathbf{V}^t$  has a 3D position vector and a signed distance value  $s$ . If two vertices in a tetrahedron have different signs of  $s$  (e.g.,  $\mathbf{v}_a^t$  with  $s_a < 0$  and  $\mathbf{v}_b^t$  with  $s_b > 0$ ), they can determine a vertex  $\mathbf{v}^m$  in triangular mesh  $\mathcal{M}$ :

$$\mathbf{v}^m = \frac{(\mathbf{v}_a^t + \delta\mathbf{v}_a^t) \cdot s_b^t - (\mathbf{v}_b^t + \delta\mathbf{v}_b^t) \cdot s_a}{s_b - s_a}, \quad (2)$$

where  $\delta\mathbf{v}$  represents the estimated vertex displacement. As a result, the triangular mesh can be generated using the differentiable volume subdivision method. However, the DM Tet cannot represent an articulated dynamic object.

**Viewpoint-conditioned diffusion model.** Zero123 [44] and follow-up works [42, 43, 45, 65] introduce a denoising diffusion model that leverage posed CLIP embedding, involving both the visual CLIP feature [56] and the camera viewpoint  $\delta\mathbf{R}$ , as the conditioning elements. Consistent with the traditional diffusion model [22], Zero123 has a forward sampling process:

$$\mathbf{z}_t = \sqrt{\bar{\alpha}_t} \mathbf{I} + \sqrt{(1 - \bar{\alpha}_t)} \boldsymbol{\epsilon}, \boldsymbol{\epsilon} \sim \mathcal{N}(0, 1), \quad (3)$$

where  $\bar{\alpha}_t$  is a hyperparameter and  $\mathbf{z}_t$  is noising image at the  $t$ -th step. Subsequently, the noise prediction model  $\hat{\boldsymbol{\epsilon}}$  is optimized to estimate the added noise  $\boldsymbol{\epsilon}$ :

$$\min \mathbb{E}_{t, \boldsymbol{\epsilon}} \|\boldsymbol{\epsilon} - \hat{\boldsymbol{\epsilon}}(\mathbf{z}_t, t, \text{CLIP}(\mathbf{I}), \delta\mathbf{R})\|_2^2. \quad (4)$$

After training, the model can generate samples from an arbitrary viewpoint given an image  $\mathbf{I}$ :

$$\mathbf{I}_{\delta\mathbf{R}} = \text{Zero123}(\mathbf{I}, \delta\mathbf{R}). \quad (5)$$

Though free-viewpoint data can be generated using the pure 2D pipeline, it lacks rigorous 3D consistency and is unable to control articulated human poses.

**Score distillation sampling.** While the diffusion model exhibits remarkable proficiency in generating 2D images, it is constrained by its inherent incapacity to produce a 3D representation directly. For 3D object generation from text prompt, DreamFusion proposes the SDS loss [54] that optimized a 3D representation with diffusion guidance, which can be formulated as

$$\nabla_{\boldsymbol{\eta}} \mathcal{L}_{SDS} \triangleq \mathbb{E}_{t, \boldsymbol{\epsilon}} \left[ w(t) (\hat{\boldsymbol{\epsilon}}(\mathbf{z}_t, t, y) - \boldsymbol{\epsilon}) \frac{\partial \mathbf{x}}{\partial \boldsymbol{\eta}} \right], \quad (6)$$

where  $\boldsymbol{\eta}$  represents the optimization parameters;  $w$  is a weighting function that depends on the timestep; and  $y$  is the text-conditioned feature.

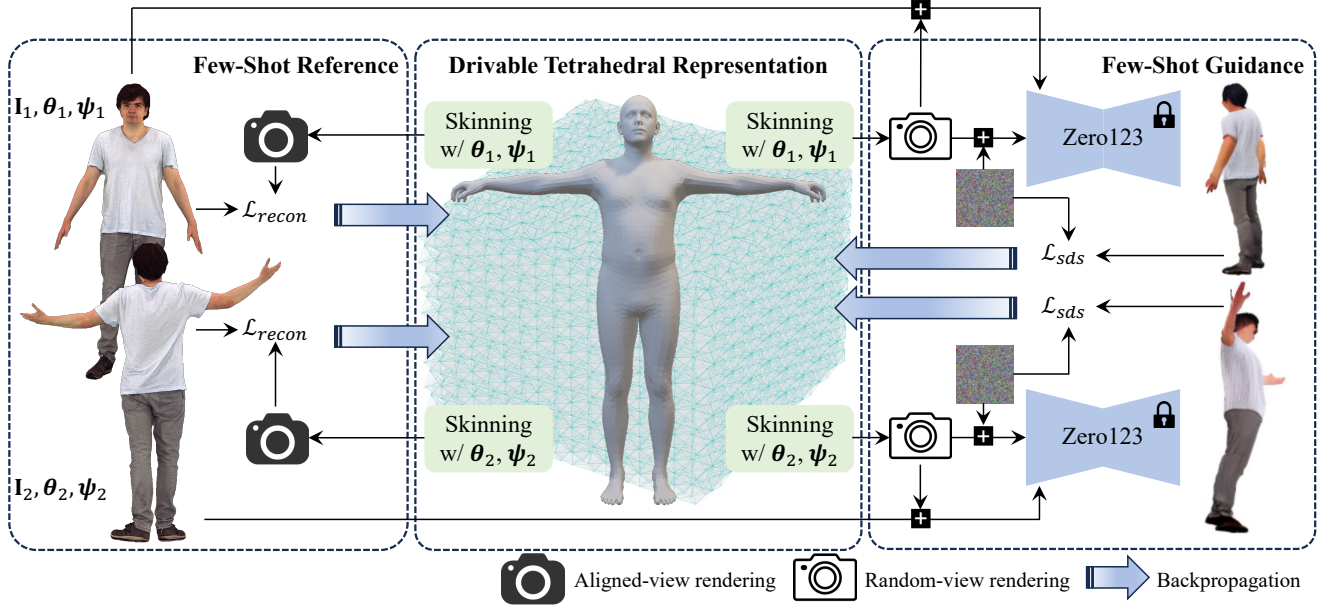


Figure 2. Overview of HaveFun framework. Based on the DMTet, we design a drivable tetrahedral representation with the skinning mechanism. In terms of optimization, we employ loss functions based on reference-data reconstruction and SDS guidance to create human avatars from few-shot unconstrained images.

### 3.2. Drivable Tetrahedral Representation

Following DMTet [64], we employ a hybrid representation for 3D geometry. As introduced in Sec 3.1, this method depends on a pre-defined tetrahedral grid, along with learnable vertex displacements  $\delta\mathbf{V}^t = \{\delta\mathbf{v}^t\}$  and signed distance values  $\mathbf{S} = \{s\}$ , to represent an arbitrary 3D geometry. Based on its volume subdivision method, the triangular mesh  $\mathcal{M} = (\mathbf{V}^m, \mathbf{F})$  can be obtained with vertices  $\mathbf{V}^m = \{\mathbf{v}^m\}$  and faces  $\mathbf{F}$ .

Nevertheless, the tetrahedral grid can only represent a static scene, while the human body has dynamic articulated poses. To tackle human articulated motion, we introduce the skinning mechanism [47] for the generated mesh. As for each vertex  $\mathbf{v}^m$ , we find the nearest triangle on a parametric mesh [51, 58] with vertices  $\{\mathbf{v}_1^p, \mathbf{v}_2^p, \mathbf{v}_3^p\}$ . The skinning weights and expression blendshapes of  $\mathbf{v}^m$  can be formulated as follows,

$$\begin{aligned} \mathbf{W}_{\mathbf{v}^m} &= u\mathbf{W}_{\mathbf{v}_1^p}^p + v\mathbf{W}_{\mathbf{v}_2^p}^p + \gamma\mathbf{W}_{\mathbf{v}_3^p}^p \\ \mathbf{E}_{\mathbf{v}^m} &= u\mathbf{E}_{\mathbf{v}_1^p}^p + v\mathbf{E}_{\mathbf{v}_2^p}^p + \gamma\mathbf{E}_{\mathbf{v}_3^p}^p, \end{aligned} \quad (7)$$

where  $u, v, \gamma$  denote the barycentric coordinates of the projection of  $\mathbf{v}^m$  onto the face;  $\mathbf{W}^p, \mathbf{E}^p$  are the skinning weights and blendshapes defined by SMPLX [51] or MANO [58]. With the retrieval of  $\mathbf{W}, \mathbf{E}$ , mesh vertices can be deformed to the posed space:

$$\tilde{\mathbf{v}}^m = \sum_{b=1}^B \mathbf{W}_{\mathbf{v}^m, b} \mathbf{G}_b(\boldsymbol{\theta}, \mathbf{J}) \mathbf{G}(\mathbf{0}, \mathbf{J})^{-1} (\mathbf{v}^m + \mathbf{E}_{\mathbf{v}^m} \boldsymbol{\psi}), \quad (8)$$

where  $\mathbf{G}$  is the kinematic transformation matrix,  $b$  indexes articulated bones, and  $\mathbf{J}$  denotes bone joints. Please refer to SMPL [47] for more details of the skinning mechanism.

Considering the limited quantity of training data and various articulated poses, the initialization of the tetrahedral grid is important yet non-trivial. Benefiting from the aforementioned drivable mechanism, we can use a canonical SMPLX template mesh for the initialization, as shown in Fig. 2. This approach allows us to incorporate human geometry prior into the 3D representation.

In addition, a texture field is adopted for colored appearance. Following Get3D [17], we find the surface points  $\mathbf{P}_s$  on  $\mathcal{M}$  that align with pixels by the rasterization of the deformed mesh  $(\{\tilde{\mathbf{v}}^m\}, \mathbf{F})$ . Then, the texture field  $\mathcal{C}$  with MLPs can predict RGB values  $\hat{\mathbf{I}}$  as follows,

$$\mathcal{C}(\mathbf{P}_s) : \mathbf{P}_s \in \mathbb{R}^{H \times W \times 3} \rightarrow \hat{\mathbf{I}} \in \mathbb{R}^{H \times W \times 3}. \quad (9)$$

### 3.3. Optimization

We use few-shot images  $\{\mathbf{I}_i\}_{i=1}^N$  to optimize the drivable tetrahedral representation with the optimization parameter  $\eta_{\mathcal{G}} = \{\delta\mathbf{V}^t, \mathbf{S}, \mathcal{C}\}$ . For each image, we use off-the-shelf tools [5, 51] to obtain parametric geometry  $\{\boldsymbol{\theta}_i, \boldsymbol{\psi}_i\}_{i=1}^N$  with pose and expression coefficients. Given aligned or novel viewpoints  $\mathbf{R}^{align}, \mathbf{R}^{novel}$ , images can be rendered with our model:

$$\begin{aligned} \hat{\mathbf{I}}_i^{align} &= \mathcal{G}(\mathbf{R}_i^{align}, \boldsymbol{\theta}_i, \boldsymbol{\psi}_i) \\ \hat{\mathbf{I}}_i^{novel} &= \mathcal{G}(\mathbf{R}_i^{novel}, \boldsymbol{\theta}_i, \boldsymbol{\psi}_i). \end{aligned} \quad (10)$$

**Few-shot reference for body.** In order to enhance the ability to express human body characteristics with few-shot images  $\{\mathbf{I}_i\}_{i=1}^N$ , we use off-the-shelf tools [15, 77] to estimate the mask  $\{\mathbf{M}_i\}_{i=1}^N$ , normal  $\{\mathbf{N}_i\}_{i=1}^N$ , and depth  $\{\mathbf{D}_i\}_{i=1}^N$ . Then, we design reconstruction losses as follows,

$$\begin{aligned} \mathcal{L}_{texture} &= \text{LPIPS}(\hat{\mathbf{I}}_i^{align}, \mathbf{I}_i) + \|\hat{\mathbf{I}}_i^{align} - \mathbf{I}_i\|_2^2 \\ \mathcal{L}_{normal} &= 1 - \langle \hat{\mathbf{N}}_i^{align}, \mathbf{N}_i \rangle \\ \mathcal{L}_{depth} &= \frac{\text{cov}(\hat{\mathbf{D}}_i^{align}, \mathbf{D}_i)}{\sigma_{\hat{\mathbf{D}}_i^{align}} \sigma_{\mathbf{D}_i}} \\ \mathcal{L}_{mask} &= \|\hat{\mathbf{M}}_i^{align} - \mathbf{M}_i\|_2^2, \end{aligned} \quad (11)$$

where  $\langle \cdot, \cdot \rangle$  represents cosine similarity, LPIPS computes perceptual similarity [83].  $\mathcal{L}_{depth}$  is formulated as pearson correlation coefficient [8].

Therefore,  $\mathcal{L}_{recon} = \mathcal{L}_{texture} + \mathcal{L}_{normal} + \mathcal{L}_{depth} + \mathcal{L}_{mask}$  for body avatar reconstruction. In addition, we render high-resolution (*i.e.*,  $256 \times 256$ ) hand/head regions to compute  $\mathcal{L}_{recon}^{hand}$ ,  $\mathcal{L}_{recon}^{head}$  using Eq. (11) and add them to the final  $\mathcal{L}_{recon}$ .

**Few-shot reference for hand.** The properties of the hand are distinctive from the body, *i.e.*, subtle wrinkles and bare geometry. Hence, the benefits from normal/depth supervision are limited and we only use  $\mathcal{L}_{texture}$  and  $\mathcal{L}_{mask}$  in Eq. (11) for hand optimization. In addition, we design a Laplacian constraint for geometry smoothness as follows,

$$\mathcal{L}_{lap} = \|\mathbf{L}_{\mathcal{M}} \cdot \hat{\mathbf{n}}_{\mathcal{M}}\|_2, \quad (12)$$

where  $\mathbf{L}_{\mathcal{M}}$  and  $\hat{\mathbf{n}}_{\mathcal{M}}$  denote Laplacian matrix and vertex normals for the generated triangle mesh  $\mathcal{M}$ . Finally,  $\mathcal{L}_{recon} = \mathcal{L}_{texture} + \mathcal{L}_{mask} + \lambda_{lap} \mathcal{L}_{lap}$  for hand avatar creation.

**Few-shot guidance.** Few-shot images cannot cover complete visual features of the human body because of sparse viewpoints and articulated self-occlusion. Therefore, we improve random-view rendering with the diffusion prior.

Specifically, pre-trained Zero123 [44] is employed as the prior model, which takes image features and viewpoints as the condition to produce random-viewpoint images, as described in Sec. 3.1. Given  $\hat{\mathbf{I}}_i^{novel}$ , we first use Eq. (3) to sample noising image  $\mathbf{z}_{i,t}$  with gaussian noise  $\epsilon$  and compute relative viewpoint  $\delta \mathbf{R}$  based on  $\mathbf{R}_i^{align}$  and  $\mathbf{R}^{novel}$ . To generate gradients to optimize the 3D representation, we employ SDS loss derived by Zero123 as follows,

$$\nabla_{\eta_{\mathcal{G}}} \mathcal{L}_{sds} = \mathbb{E}_{t, \epsilon} \left[ w(t) \left( \hat{\epsilon}(\mathbf{z}_{i,t}, t, \text{CLIP}(\mathbf{I}_i), \delta \mathbf{R}) - \epsilon \right) \frac{\partial \mathbf{x}}{\partial \eta_{\mathcal{G}}} \right]. \quad (13)$$

Overall,  $\mathcal{L} = \mathcal{L}_{recon} + \lambda_{sds} \mathcal{L}_{sds}$  is employed to optimize the 3D representation  $\mathcal{G}$  for creating human avatars.



Figure 3. Training data for ablation studies in Figs. 4 and 5. Blue and green boxes indicate 2- and 4-shot training data, respectively.

## 4. Experiments

### 4.1. Datasets

We build new dataset benchmarks to investigate the task of few-shot dynamic human reconstruction. The purpose of our dataset is to generate training data with casual human poses and conduct multi-view evaluations under the canonical human pose. To present sufficient information with few-shot data, extremely articulated self-occlusion is neglected.

For quantitative metrics, we report LPIPS [83], PSNR, and SSIM [73] to reflect rendering quality.

**FS-XHumans.** XHumans [63] offer 3D clothed human scans with 20 different identities and various poses. For each identity, we choose 8 scans with different poses and render them from different viewpoints to create our training data. Furthermore, as XHumans do not provide canonical-pose data, we select the scan from their dataset that most closely matches the A-pose and render it from 24 spherically distributed viewpoints for evaluation.

**FS-DART.** We utilize DART [16], a hand texture model, to generate 100 hand identities with distinct shapes and textures. To acquire training poses, we collect real hand poses using a monocular reconstruction method [5]. Each hand identity has 8 training data with varying poses and viewpoints. For evaluation, we render zero-pose hand samples from 24 spherically distributed viewpoints.

Please refer to *suppl. material* for more dataset details.

### 4.2. Ablation Studies

Considering that the primary information can be presented on the palm and back of the hand, we use the 2-shot setting for hand ablation studies. In contrast, the lateral body holds important information, so we employ the 4-shot setting to

| Method                                    | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
|-------------------------------------------|-----------------|-----------------|--------------------|
| <i>4-shot FS-XHumans</i>                  |                 |                 |                    |
| $\lambda_{sds} = 0$                       | 24.54           | 0.9522          | 0.0379             |
| $\lambda_{sds} = 0.01$                    | <b>25.64</b>    | <b>0.9627</b>   | <b>0.0347</b>      |
| $\lambda_{sds} = 0.05$                    | 25.10           | 0.9562          | 0.0373             |
| $\lambda_{sds} = 1$                       | 24.56           | 0.9413          | 0.0403             |
| <i>2-shot FS-DART</i>                     |                 |                 |                    |
| $\lambda_{sds} = 0, \lambda_{lap} = 1$    | 25.62           | 0.9532          | 0.0627             |
| $\lambda_{sds} = 0.01, \lambda_{lap} = 1$ | 26.22           | 0.9635          | <b>0.0559</b>      |
| $\lambda_{sds} = 0.05, \lambda_{lap} = 1$ | <b>26.61</b>    | <b>0.9673</b>   | 0.0577             |
| $\lambda_{sds} = 1, \lambda_{lap} = 1$    | 26.01           | 0.9656          | 0.0736             |
| $\lambda_{sds} = 0.05, \lambda_{lap} = 0$ | 26.13           | 0.9648          | 0.0725             |
| $\lambda_{sds} = 0.05, \lambda_{lap} = 5$ | 25.69           | 0.9633          | 0.0602             |

Table 1. The effects of SDS and Laplacian normal losses.




Figure 4. Ablation studies on the few-shot body reconstruction task. Zoom in to see details.

study the performance on the human body. Finally, we report our  $N$ -shot ( $N = 2, 4, 8$ ) results. The training data used for Figs. 4 and 5 are shown in Fig. 3.

**Effect of few-shot guidance.** The SDS loss is utilized to facilitate the reconstruction of unseen regions. Therefore,




Figure 5. Ablation studies on the few-shot hand reconstruction task. Zoom in to see details.

the effect of SDS should be compatible with unseen areas. That is, in contrast to the 4-shot setting, the 2-shot reconstruction exhibits a heightened reliance on the SDS loss due to the lack of information. As shown in Table 1,  $\lambda_{sds}$  offers different optimal choices for varying training data amounts.

The effect of SDS loss is also demonstrated in Fig. 4. As indicated by arrows,  $\lambda_{sds} = 0$  prevents the model from presenting reasonable unseen textures. Besides, over-size SDS constraints would lead to the problem of color distortion, thereby harming the reconstruction quality. As a result, we use  $\lambda_{sds} = 0.05, 0.01, 0.01$  for 2-, 4-, and 8-shot reconstruction settings.

**Effect of Laplacian normal constraint.** Instead of ground-truth normal/depth supervision, we design a Laplacian normal loss to regularize hand geometry. As illustrated in Fig. 5,  $\lambda_{lap} = 0$  induces fractured geometry, and  $\lambda_{lap} = 5$  can produce an over-smooth geometry. These two situations are not suited to represent the hand. That is, the hand has subtle palm wrinkles and texture. If the geometry can effectively capture wrinkles, it could contribute to generating faithful textures. Surprisingly, our experiment shows that the proper Laplacian normal constraint can achieve this. As shown,  $\lambda_{lap} = 1$  promotes our model to represent detailed hand shapes. Benefiting from the informative geometry, the texture reconstruction is more accurate, as indicated by arrows.

| Method            | PSNR $\uparrow$         | SSIM $\times 10^2 \uparrow$ | LPIPS $\times 10^2 \downarrow$ |
|-------------------|-------------------------|-----------------------------|--------------------------------|
| <i>FS-XHumans</i> |                         |                             |                                |
| SelfRecon         | 19.9/20.7               | 92.7/94.3                   | 6.5/6.3                        |
| TeCH              | 21.0                    | 92.4                        | 6.5                            |
| HaveFun (ours)    | 24.0/25.6/ <b>26.8</b>  | 95.5/96.3/ <b>96.7</b>      | 4.2/3.5/ <b>3.0</b>            |
| <i>FS-DART</i>    |                         |                             |                                |
| SelfRecon         | 20.8/21.3/21.7          | 92.0/92.4/92.8              | 9.6/9.0/8.7                    |
| HaveFun (ours)    | 26.6/ <b>26.7</b> /26.3 | 96.7/ <b>96.8</b> /96.7     | 5.8/5.5/ <b>5.2</b>            |

Table 2. Comparison of few-shot human reconstruction. For FS-Xhumans, TeCH [26] only supports single-image reconstruction. SelfRecon [30] results follow the format of “8-/100-shot”, whereas our metrics are provided for “2-/4-/8-shot” tasks. For FS-DART, all metrics are for “2-/4-/8-shot” tasks.

**Ablation study on  $N$ -shot reconstruction.** Our framework can reconstruct a human avatar with an arbitrary quantity of images. We argue that one-shot data cannot furnish adequate information for human reconstruction, so we demonstrate 2-, 4-, and 8-shot tasks as examples. Referring to our results in Table 2, 2-shot metrics are very close to those of the 8-shot task, indicating our approach is adept at reconstructing humans from minimal data amount. When it comes to the hand reconstruction task, some 2-shot results are even better than those of the 8-shot task. The primary reason is that the hand feature is predominantly influenced by the hand palm or back, with the lateral hand contributing only minimal additional information.

Similar conclusions can be observed in Figs. 4 and 5. Visually, the 8-shot task shows the best quality with faithful hair geometries, cloth textures, palm wrinkles, *etc.*

### 4.3. Comparison with Prior Arts

We compare our approach with SelfRecon [30] and TeCH [26]. The former can take dynamic images or videos as the input, whereas the latter uses the DMTet and SDS prior for static human reconstruction. We present typical results in Figs. 6 and 7. As shown, SelfRecon tends to generate overly smooth human appearances and inherently fails to learn effectively from limited data. Moreover, due to the absence of expression handling, SelfRecon still struggles to produce a plausible portrait, even when trained using video (100-shot) data. Text-based TeCH can produce detailed texture and sound geometry with a single image. However, TeCH cannot perform a faithful reconstruction. As shown in the first row of Fig. 7, facial identity cannot be preserved by TeCH. Furthermore, features from the text caption would be improperly introduced to the avatar. Referring to the second row of Fig. 7, the front head is aligned with the reference image while the lateral head is dominated by the BLIP [37] caption of “caucasian”. Hence, since the caption cannot perfectly describe visual details, the text-based approach is akin to a generative method rather than a reconstructive one. In contrast, thanks to pure visual prompts



Figure 6. Comparison of hand reconstruction on FS-DART. See *suppl. material* for training data.

from few-shot images, our method can perform a faithful avatar reconstruction for the body, face, and hand.

It is worthwhile to revisit our 2- and 8-shot results again. As shown in the first row of Fig. 7, the 2-shot result is better than the 8-shot one in terms of portrait reconstruction. This discrepancy arises from the diverse expressions presented in training data. That is, the 2-shot task relies on a single image to describe the face, while the 8-shot task involves lateral body data with various facial expressions. Due to the imperfect expression blendshapes from SMPLX, the performance of 8-shot portrait reconstruction from data with diverse expressions is constrained. In addition, the 8-shot reconstruction yields enhanced results for hands and lateral textures owing to the expanded visibility of data.

For quantitative comparison, our method achieves the best results in all metrics, as shown in Table 2, indicating superior rendering quality of the HaveFun framework.

### 4.4. Applications

**Animation results.** Thanks to the drivable tetrahedral representation, our method can perform free-pose articulated deformation for the body and hand so that complex human motion can be presented, as shown in Fig. 8.

**Reconstructing human avatars with real-world casual capture.** The FS-DART is a synthetic dataset, and the FS-XHumans provides real human images but captured in a stu-




Figure 7. Comparison of body reconstruction on FS-XHumans, where TeCH is a 1-shot method and SelfRecon is illustrated with 8-shot or video (100-shot) data training. See *suppl. material* for training data.



Figure 8. Articulated animation of human avatars.




Figure 9. Avatar reconstruction from real-world captured data.

dio. Therefore, we acquire real-world unconstrained images to validate our approach. As shown in Fig. 9, based on 2 or 4 images, our method excels in canonical-space avatar reconstruction and free-pose animation.

## 5. Conclusions

This paper poses a novel research problem of human avatar reconstruction from few-shot unconstrained images. We propose a HaveFun framework with a drivable tetrahedral representation to solve this issue. To optimize our 3D representation, we design a two-phase method with few-shot reference and few-shot guidance. In addition, we develop evaluation benchmarks for the human body and hand. As a result, our approach can produce animatable human avatars with superior rendering quality, which we believe enables a new way for real-world avatar creation.

**Acknowledgment** The work was supported in part by the Basic Research Project No.HZQB-KCZYZ-2021067 of Hetao Shenzhen-HK S&T Cooperation Zone, Guangdong Provincial Outstanding Youth Project No. 2023B1515020055, the National Key R&D Program of China with grant No.2018YFB1800800, by Shenzhen Outstanding Talents Training Fund 202002, by Guangdong Research Projects No.2017ZT07X152 and No.2019CX01X104, by Key Area R&D Program of Guangdong Province (Grant No.2018B030338001), by the Guangdong Provincial Key Laboratory of Future Networks of Intelligence (Grant No.2022B1212010001), and by Shenzhen Key Laboratory of Big Data and Artificial Intelligence (Grant No.ZDSYS201707251409055). It is also partly supported by NSFC-62172348, Shenzhen General Project No. JCYJ20220530143604010 and China National Postdoctoral Program for Innovative Talents No. BX2023004.



## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN: How to embed images into the stylegan latent space? In *ICCV*, 2019. 2
- [2] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K Wong. DreamAvatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv:2304.00916*, 2023. 3
- [3] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 2
- [4] Xu Chen, Yufeng Zheng, Michael J Black, Otmar Hilliges, and Andreas Geiger. SNARF: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 2
- [5] Xingyu Chen, Yufeng Liu, Yajiao Dong, Xiong Zhang, Chongyang Ma, Yanmin Xiong, Yuan Zhang, and Xiaoyan Guo. Mobrecon: Mobile-friendly hand mesh reconstruction from monocular image. In *CVPR*, 2022. 1, 4, 5
- [6] Xingyu Chen, Yu Deng, and Baoyuan Wang. Mimic3d: Thriving 3d-aware gans via 3d-to-2d imitation. In *ICCV*, 2023. 2
- [7] Xingyu Chen, Baoyuan Wang, and Heung-Yeung Shum. Hand avatar: Free-pose hand animation and rendering from monocular video. In *CVPR*, 2023. 2
- [8] Israel Cohen, Yiteng Huang, Jingdong Chen, Jacob Benesty, Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. Pearson correlation coefficient. *Noise reduction in speech processing*, pages 1–4, 2009. 5
- [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. 2
- [10] Enric Corona, Mihai Zanfir, Thimo Alldieck, Eduard Gabriel Bazavan, Andrei Zanfir, and Cristian Sminchisescu. Structured 3d features for reconstructing controllable avatars. In *CVPR*, 2023. 2
- [11] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised NeRF: Fewer views and faster training for free. In *CVPR*, 2022. 3
- [12] Yu Deng, Jiaolong Yang, Jianfeng Xiang, and Xin Tong. GRAM: Generative radiance manifolds for 3D-aware image generation. In *CVPR*, 2022. 2
- [13] Yu Deng, Baoyuan Wang, and Heung-Yeung Shum. Learning detailed radiance manifolds for high-fidelity and 3D-consistent portrait synthesis from monocular image. In *CVPR*, 2023. 2
- [14] Zijian Dong, Xu Chen, Jinlong Yang, Michael J Black, Otmar Hilliges, and Andreas Geiger. AG3D: Learning to Generate 3D Avatars from 2D Image Collections. In *ICCV*, 2023. 2
- [15] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *ICCV*, 2021. 5
- [16] 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. 2, 5
- [17] Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, and Sanja Fidler. Get3D: A generative model of high quality 3D textured shapes learned from images. In *NeurIPS*, 2022. 4
- [18] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular RGB videos. In *CVPR*, 2022. 2
- [19] Chen Guo, Tianjian Jiang, Xu Chen, Jie Song, and Otmar Hilliges. Vid2Avatar: 3D avatar reconstruction from videos in the wild via self-supervised scene decomposition. In *CVPR*, 2023. 2
- [20] Zhiyang Guo, Wengang Zhou, Min Wang, Li Li, and Houqiang Li. HandNeRF: Neural radiance fields for animatable interacting hands. In *CVPR*, 2023. 2
- [21] Hsuan-I Ho, Lixin Xue, Jie Song, and Otmar Hilliges. Learning locally editable virtual humans. In *CVPR*, 2023. 2
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 3
- [23] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. Avatarclip: Zero-shot text-driven generation and animation of 3D avatars. *ACM TOG*, 41(4):1–19, 2022. 3
- [24] Yukun Huang, Jianan Wang, Ailing Zeng, He Cao, Xi-anbiao Qi, Yukai Shi, Zheng-Jun Zha, and Lei Zhang. Dreamwaltz: Make a scene with complex 3D animatable avatars. *arXiv:2305.12529*, 2023.
- [25] Yangyi Huang, Hongwei Yi, Weiyang Liu, Haofan Wang, Boxi Wu, Wenxiao Wang, Binbin Lin, Debing Zhang, and Deng Cai. One-shot implicit animatable avatars with model-based priors. In *ICCV*, 2023.
- [26] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Jiaxiang Tang, Deng Cai, and Justus Thies. TeCH: Text-guided reconstruction of lifelike clothed humans. In *3DV*, 2024. 2, 3, 7
- [27] Mustafa Işık, Martin Rünz, Markos Georgopoulos, Taras Khakhulin, Jonathan Starck, Lourdes Agapito, and Matthias Nießner. HumanRF: High-fidelity neural radiance fields for humans in motion. *ACM TOG*, 42(4):1–12, 2023. 2
- [28] Shun Iwase, Shunsuke Saito, Tomas Simon, Stephen Lombardi, Timur Bagautdinov, Rohan Joshi, Fabian Prada, Takaaki Shiratori, Yaser Sheikh, and Jason Saragih. RelightableHands: Efficient neural relighting of articulated hand models. In *CVPR*, pages 16663–16673, 2023. 2
- [29] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *ICCV*, 2021. 3
- [30] Boyi Jiang, Yang Hong, Hujun Bao, and Juyong Zhang. SelfRecon: Self reconstruction your digital avatar from monocular video. In *CVPR*, 2022. 2, 7
- [31] Ruixiang Jiang, Can Wang, Jingbo Zhang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. AvatarCraft: Transforming text into neural human avatars with parameterized shape and pose control. In *ICCV*, 2023. 3

- [32] Suyi Jiang, Haoran Jiang, Ziyu Wang, Haimin Luo, Wenzheng Chen, and Lan Xu. HumanGen: Generating human radiance fields with explicit priors. In *CVPR*, 2023. 2
- [33] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, 2022. 2
- [34] Korrawe Karunratanakul, Sergey Prokudin, Otmar Hilliges, and Siyu Tang. Harp: Personalized hand reconstruction from a monocular rgb video. In *CVPR*, 2023. 2
- [35] Mijeong Kim, Seonguk Seo, and Bohyung Han. InfoNeRF: Ray entropy minimization for few-shot neural volume rendering. In *CVPR*, 2022. 3
- [36] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Fieraru, and Cristian Sminchisescu. DreamHuman: Animatable 3d avatars from text. *arxiv:2306.09329*, 2023. 3
- [37] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 3, 7
- [38] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. TAVA: Template-free animatable volumetric actors. In *ECCV*, 2022. 2
- [39] Zhe Li, Zerong Zheng, Yuxiao Liu, Boyao Zhou, and Yebin Liu. Posevocab: Learning joint-structured pose embeddings for human avatar modeling. In *ACM SIGGRAPH Conference Proceedings*, 2023. 2
- [40] Tingting Liao, Hongwei Yi, Yuliang Xiu, Jiayang Tang, Yangyi Huang, Justus Thies, and Michael J Black. Tada! text to animatable digital avatars. *arXiv:2308.10899*, 2023. 3
- [41] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3D: High-resolution text-to-3D content creation. In *CVPR*, 2023. 3
- [42] Yukang Lin, Haonan Han, Chaoqun Gong, Zunnan Xu, Yachao Zhang, and Xiu Li. Consistent123: One image to highly consistent 3d asset using case-aware diffusion priors. *arXiv:2309.17261*, 2023. 3
- [43] Minghua Liu, Chao Xu, Haiyan Jin, Linghao Chen, Zexiang Xu, Hao Su, et al. One-2-3-45: Any single image to 3d mesh in 45 seconds without per-shape optimization. *arXiv:2306.16928*, 2023. 3
- [44] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3D object. In *ICCV*, 2023. 2, 3, 5
- [45] Yuan Liu, Cheng Lin, Zijiao Zeng, Xiaoxiao Long, Lingjie Liu, Taku Komura, and Wenping Wang. SyncDreamer: Generating multiview-consistent images from a single-view image. *arXiv:2309.03453*, 2023. 3
- [46] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *ECCV*, 2022. 3
- [47] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015. 4
- [48] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. RealFusion: 360 reconstruction of any object from a single image. In *CVPR*, 2023. 2, 3
- [49] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 2
- [50] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Reg-NeRF: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, 2022. 3
- [51] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019. 2, 4
- [52] Sida Peng, Juntong Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 2
- [53] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 1, 2
- [54] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2023. 2, 3
- [55] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural radiance fields for dynamic scenes. In *CVPR*, 2021. 2
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3
- [57] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM TOG*, 42(1):1–13, 2022. 2
- [58] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM TOG*, 2017. 4
- [59] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3
- [60] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 1, 2
- [61] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 2
- [62] 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. 2

- [63] Kaiyue Shen, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Julien Valentin, Jie Song, and Otmar Hilliges. X-Avatar: Expressive human avatars. In *CVPR*, 2023. 1, 2, 5
- [64] Tianchang Shen, Jun Gao, Kangxue Yin, Ming-Yu Liu, and Sanja Fidler. Deep marching tetrahedra: a hybrid representation for high-resolution 3D shape synthesis. In *NeurIPS*, 2021. 2, 4
- [65] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: A single image to consistent multi-view diffusion base model. *arXiv:2310.15110*, 2023. 3
- [66] Jingxiang Sun, Xuan Wang, Lizhen Wang, Xiaoyu Li, Yong Zhang, Hongwen Zhang, and Yebin Liu. Next3D: Generative neural texture rasterization for 3D-aware head avatars. In *CVPR*, 2023. 2
- [67] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3D: High-fidelity 3D creation from a single image with diffusion prior. In *ICCV*, 2023. 2, 3
- [68] Xiao Tang, Tianyu Wang, and Chi-Wing Fu. Towards accurate alignment in real-time 3D hand-mesh reconstruction. In *ICCV*, 2021. 1
- [69] 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. 2
- [70] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2D diffusion models for 3D generation. In *CVPR*, 2023. 2, 3
- [71] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. Arah: Animatable volume rendering of articulated human SDFs. In *ECCV*, 2022. 2
- [72] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-fidelity GAN inversion for image attribute editing. In *CVPR*, 2022. 2
- [73] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 2004. 5
- [74] 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. 1, 2
- [75] Chung-Yi Weng, Pratul P Srinivasan, Brian Curless, and Ira Kemelmacher-Shlizerman. PersonNeRF: Personalized reconstruction from photo collections. In *CVPR*, 2023. 1, 2
- [76] Zhangyang Xiong, Di Kang, Derong Jin, Weikai Chen, Linchao Bao, Shuguang Cui, and Xiaoguang Han. Get3DHuman: Lifting StyleGAN-Human into a 3D generative model using pixel-aligned reconstruction priors. In *ICCV*, 2023. 2
- [77] Yuliang Xiu, Jinlong Yang, Dimitrios Tzionas, and Michael J. Black. ICON: Implicit clothed humans obtained from normals. In *CVPR*, 2022. 2, 5
- [78] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J. Black. ECON: Explicit clothed humans optimized via normal integration. In *CVPR*, 2023. 2
- [79] Tianhan Xu, Yasuhiro Fujita, and Eiichi Matsumoto. Surface-aligned neural radiance fields for controllable 3D human synthesis. In *CVPR*, 2022. 2
- [80] Wenqi Yang, Guanying Chen, Chaofeng Chen, Zhenfang Chen, and Kwan-Yee K. Wong. PS-NeRF: Neural inverse rendering for multi-view photometric stereo. In *ECCV*, 2022. 3
- [81] Zhengming Yu, Wei Cheng, Xian Liu, Wayne Wu, and Kwan-Yee Lin. MonoHuman: Animatable human neural field from monocular video. *CVPR*, 2023. 2
- [82] Huichao Zhang, Bowen Chen, Hao Yang, Liao Qu, Xu Wang, Li Chen, Chao Long, Feida Zhu, Kang Du, and Min Zheng. Avatarverse: High-quality & stable 3D avatar creation from text and pose. *arXiv:2308.03610*, 2023. 3
- [83] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 5
- [84] 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. 2
- [85] 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. 2
- [86] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, 2022. 2
- [87] Zerong Zheng, Xiaochen Zhao, Hongwen Zhang, Boning Liu, and Yebin Liu. Avatarrex: Real-time expressive full-body avatars. *ACM TOG*, 42(4), 2023. 2