Efficient Meshy Neural Fields for Animatable Human Avatars

Xiaoke Huang[1]   Yiji Cheng[1]   Yansong Tang[1*]   Xiu Li[1]   Jie Zhou[2]   Jiwen Lu[2]

{[1]Shenzhen International Graduate School, [2]Department of Automation}, Tsinghua University
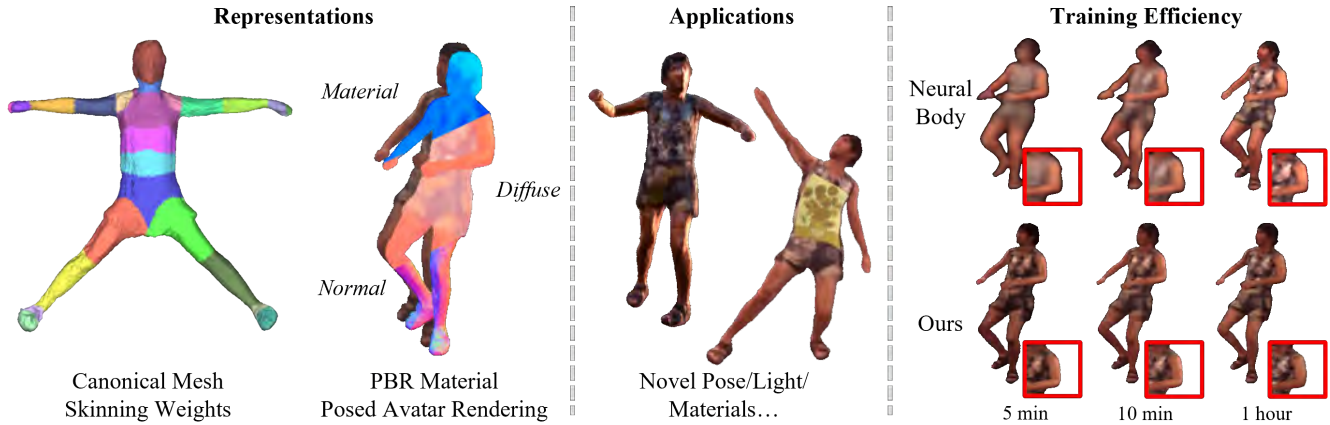
Figure 1: **EMA** efficiently and jointly learns canonical shapes, materials, and motions via differentiable inverse rendering in an end-to-end manner. The method does not require any predefined templates or riggings. The derived avatars are animatable and can be directly applied to the graphics renderer and downstream tasks. All figures are best viewed in color.

Abstract

*Efficiently digitizing high-fidelity animatable human avatars from videos is a challenging and active research topic. Recent volume rendering-based neural representations open a new way for human digitization with their friendly usability and photo-realistic reconstruction quality. However, they are inefficient for long optimization times and slow inference speed; their implicit nature results in entangled geometry, materials, and dynamics of humans, which are hard to edit afterward. Such drawbacks prevent their direct applicability to downstream applications, especially the prominent rasterization-based graphic ones. We present **EMA**, a method that **E**fficiently learns **M**eshy neural fields to reconstruct animatable human **A**vatars. It jointly optimizes explicit triangular canonical mesh, spatial-varying material, and motion dynamics, via inverse rendering in an end-to-end fashion. Each above component is derived from separate neural fields, relaxing the requirement of a template, or rigging. The mesh representation is highly compatible with the efficient rasterization-based renderer, thus our method only takes about an hour of training and can render in real-time. Moreover, only minutes of optimization are enough for plausible reconstruction results. The disentanglement of meshes enables direct downstream applications. Extensive experiments illustrate the very competitive performance and significant speed boost against previous methods. We also showcase applications including novel pose synthesis, material editing, and relighting. The project page: https://xk-huang.github.io/ema/.*

1. Introduction

Recent years have witnessed the rise of human digitization [25, 2, 70, 3, 72]. This technology greatly impacts the entertainment, education, design, and engineering industry. There is a well-developed industry solution for this task. High-fidelity reconstruction of humans can be achieved either with full-body laser scans [77], dense synchronized multi-view cameras [91, 90], or light stages [2]. However, these settings are expensive and tedious to deploy and consist of a complex processing pipeline, preventing the technology's democratization.

Another solution is to view the problem as inverse rendering and learn digital humans directly from custom-

---
*Corresponding author.

collected data. Traditional approaches directly optimize explicit mesh representation [53, 17, 68] which suffers from the problems of smooth geometry and coarse textures [71, 4]. Besides, they require professional artists to design human templates, rigging, and unwrapped UV coordinates. Recently, with the help of volumetric-based implicit representations [58, 65, 57] and neural rendering [43, 50, 83], one can easily digitize a quality-plausible human avatar from video footage [31, 89]. Particularly, volumetric-based implicit representations [58, 70] can reconstruct scenes or objects with much higher fidelity against previous neural renderer [83, 71], and is more user-friendly as it does not need any human templates, pre-set rigging, or UV coordinates. Captured visual footage and corresponding skeleton tracking are enough for training. However, better reconstructions and more friendly usability are at the expense of the following factors. 1) **Inefficiency:** They require longer optimization times (typically tens of hours or days) and inference slowly. Volume rendering [58, 52] formulates images by querying the densities and colors of millions of spatial coordinates. In the training stage, due to memory constraints, only a small fraction of points are sampled which leads to slow convergence speed. 2) **Entangled representations**: The geometry, materials, and motion dynamics are entangled in the neural networks. Due to the implicit nature of neural nets, one can hardly edit one property without touching the others [93, 51]. 3) **Graphics incompatibility**: Volume rendering is incompatible with the current popular graphic pipeline, which renders triangular/quadrilateral meshes efficiently with the rasterization technique. Many downstream applications require mesh rasterization in their workflow (*e.g.*, editing [18], simulation [7], real-time rendering [59], ray-tracing [84]). Although there are approaches [54, 42] can convert volumetric fields into meshes, the gaps from discrete sampling degrade the output quality in terms of both meshes and textures.

To address these issues, we present **EMA**, a method based on **E**fficient **M**eshy neural fields to reconstruct animatable human **A**vatars. Our method enjoys flexibility from implicit representations and efficiency from explicit meshes, yet still maintains high-fidelity reconstruction quality. Given video sequences and the corresponding pose tracking, our method digitizes humans in terms of canonical triangular meshes, physically-based rendering (PBR) materials, and skinning weights *w.r.t.* skeletons. We jointly learn the above components via inverse rendering [43, 14, 13] in an end-to-end manner. Each of them is derived from a separate neural field, which relaxes the requirements of a preset human template, rigging, or UV coordinates. Specifically, we predict a canonical mesh out of a signed distance field (SDF) by differentiable marching tetrahedra [79, 20, 19, 61], then we extend the marching tetrahedra [79] for spatial-varying materials by utilizing a

neural field to predict PBR materials *on the mesh surfaces* after rasterization [61, 26, 43]. To make the canonical mesh animatable, we take another neural field to model the forward linear blend skinning for the meshes. Given a posed skeleton, the canonical mesh is then transformed into the corresponding poses. Finally, we shade the mesh with a rasterization-based differentiable renderer [43] and train our models with a photo-metric loss. After training, we export the mesh with materials and discard the neural fields.

There are several merits of our method design. 1) **Efficiency**: Powered by efficient mesh rendering, our method can render in real-time. Besides, the training speed is boosted as well, since we compute loss holistically on the whole image and the gradients only flow on the mesh surface. In contrast, volume rendering takes limited pixels for loss computation and back-propagates the gradients in the whole space. Our method only needs about an hour of training and minutes of optimization are enough for plausible avatar reconstruction. 2) **Disentangled representations**: Our shape, materials, and motion modules are disentangled naturally by design, which facilitates editing. Besides, Canonical meshes with forward skinning modeling handle the out-of-distribution poses better. 3) **Graphics compatibility**: Our derived mesh representation is compatible with the prominent graphic pipeline, which leads to instant downstream applications (*e.g.*, the shape and materials can be edited directly in design software [18]). To further improve reconstruction quality, we additionally optimize image-based environment lights and non-rigid motions.

We conduct extensive experiments on standards benchmarks H36M [30] and ZJU-MoCap [70]. Our method achieves very competitive performance for novel view synthesis, generalizes better for novel poses, and significantly improves both training time and inference speed against previous arts. Our research-oriented code reaches real-time inference speed (100+ FPS for rendering 512×512 images). We in addition showcase applications including novel pose synthesis, material editing, and relighting.

2. Related Works

Explicit Representations for Human Modeling: It is intuitive to model the surfaces of humans with mesh. However, humans are highly varied in both shape and appearance and have a large pose space, which all contribute to a high dimensional space. Researchers first model humans with limited clothes. One of the prevalent methods is parametric models [6, 53, 68, 74, 81]. Fitting humans from scans is inapplicable in real-world applications. Thus, [34, 9, 41, 96, 95, 40, 82] estimate the human surface from images or videos. To model the clothed human, [71, 5, 4] deform the template human vertices in canonical T-pose. However, these methods are prone to capturing coarse geometry due to the limited capacity of the
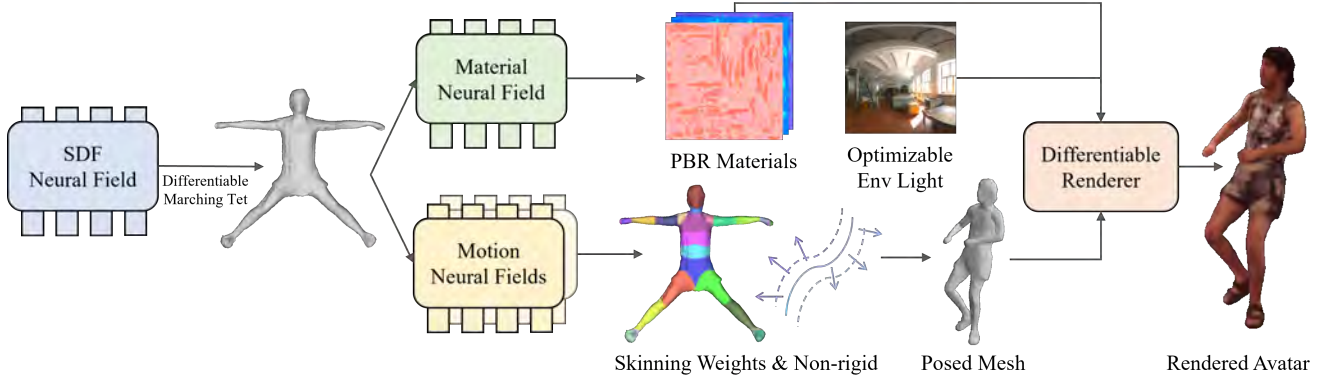
Figure 2: **The pipeline of EMA**. EMA jointly optimizes canonical shapes, materials, lights, and motions via efficient differentiable inverse rendering. The canonical shapes are attained firstly through the differentiable marching tetrahedra [19, 79, 61], which converts SDF fields into meshes. Next, it queries PBR materials, including diffuse colors, roughness, and specularity on the mesh surface. Meanwhile, the skinning weights and per-vertices offsets are predicted on the surface as well, which are then applied to the canonical meshes with the guide of input skeletons. Finally, a rasterization-based differentiable renderer takes in the posed meshes, materials, and environment lights, and renders the final avatars efficiently.

deformation layer. Besides, the textures are modeled with sphere harmonics which are far from photo-realistic. Our method takes the mesh as our core representation to enable efficient training and rendering and realize the topological change of shape and photo-realistic texture via neural fields. **Implicit Representations for Human Modeling:** Implicit representations [65, 57, 58] model the objects in a continuous manner, whose explicit entity cannot be attained directly. Specifically, Signed Distance Function [65], Occupancy Field [57] and Radiance Field [58] are all parametrized by neural networks. Given full-body scans as 3D supervision, [75, 76, 28, 29, 5] learned the SDFs or occupancy fields directly from images, which could predict photo-realistic human avatars in inference phrase. [70, 81, 49, 69, 45, 31, 12, 88, 99, 63, 102] leveraged the radiance field for more photo-realistic human avatars from multi-view images or single-view videos without any 3D supervision. Although implicit representations improve reconstruction quality against explicit ones, they still have drawbacks, *e.g.*, large computation burden or poor geometry.s Besides, volume rendering is incompatible with graphics hardware, thus the outputs are inapplicable in downstream applications without further post-processing. Our method absorbs the merits of implicit representations by using neural networks to predict photo-realistic textures and shape fields, leveraging [79] to convert SDFs to explicit meshes, which are fully compatible with the graphic pipeline. **Hybrid Representations for Human Modeling:** There are two tracks of literature modeling humans with explicit geometry representations and implicit texture representations. One track of literature [38, 101] leveraged neural rendering techniques [83]. Meshes [71, 101, 4, 3] or point clouds [86] are commonly chosen explicit representations. Moreover,

fine-grained geometry and textures are learned by neural networks. However, these methods are either only applicable for novel view synthesis [86] or restricted to self-rotation video captures [4, 3]. Besides, the neural renderers have limitations, *e.g.*, stitching texture [36, 37], and baked textures into the renderer. In contrast, the human avatars learned by our method are compatible with graphics pipeline, which means they are **applicable in downstream tasks**, *e.g.*, re-posing, editing in design software. The other track of literature leveraged neural networks to learn both geometry and textures based on differentiable rendering [43, 14, 13]. It [50, 8, 43, 73, 44, 72] equips traditional graphics pipeline with the ability of error backpropagation, which make scene properties (*i.e,* assets, lights, cameras poses, *etc.*) optimizable through gradient descent *w.r.t* photo-metric loss. Thus, we can learn both geometry and textures that are compatible with existing graphics hardware. However, the geometry optimization process is non-convex and highly unstable [23], so it is hard to give fine-grained geometry details. Besides, the topology of the mesh is fixed leading to limited shape modeling. We leverage [61, 79] to convert SDFs to meshes with differentiable marching tets, and model the motion dynamics of humans with additional neural fields. Our method enjoys flexibility from implicit representations and efficiency brought by explicit meshes, yet still maintains high-fidelity reconstructions.

3. Method

We formulate the problem as inverse rendering and extend [61] to model dynamic actors that are driven solely by skeletons. The canonical shapes, materials, lights, and actor motions are learned jointly in an end-to-end manner.

The rendering happens with an efficient rasterization-based differentiable renderer [43].

Optimization Task: Let Φ denote all the trainable parameters (*i.e.,* SDF values and corresponding vertices offset parameters for canonical geometry, spatial-varying and pose-dependent materials and light probe parameters for shading, and forward skinning weights and non-rigid vertices offset parameters for motion). For a given camera pose $\mathbf{c}$ and a tracked skeleton pose $\mathbf{P}$, we render the image $I_{\Phi}(\mathbf{c}, \mathbf{P})$ with a differentiable renderer, and compute loss with a loss function L, against the reference image $I_{ref}(\mathbf{c}, \mathbf{P})$. The optimization goal is to minimize the empirical risk:

$$\underset{\phi}{\arg\min} \, \mathbb{E}_{\mathbf{c}, \mathbf{P}} \left[L\big(I_{\Phi}(\mathbf{c}, \mathbf{P}), I_{ref}(\mathbf{c}, \mathbf{P})\big)\right]. \quad (1)$$

The parameters Φ are optimized with Adam [39] optimizer. Following [61], our loss function $L = L_{\text{img}} + L_{\text{mask}} + L_{\text{reg}}$ consists of three parts: an image loss L_{img} using ℓ_1 norm on tone mapped color, and mask loss L_{mask} using squared ℓ_2, and regularization losses L_{reg} to improve the quality of canonical geometry, materials, lights, and motion. At each optimization step, our method holistically learns both shape and materials from the whole image, while the volume rendering-based implicit counterparts only learn from limited pixels. Powered by an efficient rasterization-based renderer, our method enjoys both faster convergence and real-time rendering speed. For optimization losses and implementation details, please refer to our supplementary.

3.1. Canonical Geometry

Rasterization-based differentiable renderers take triangular meshes as input, which means the whole optimization process happens over the mesh representation. Previous works [4, 3] require a mesh template to assist optimization as either a good initialization or a shape regularization. The templates have fixed topology under limited resolutions which harm the geometry quality. Besides, to make the learned geometry generalize to novel poses, the underlying geometry representations should lie in a canonical space.

We utilize the differentiable marching tetrahedra [79, 19] algorithm, which converts SDF fields into triangular meshes to model the actors in canonical space. This method enjoys the merit of both template and topology-free from implicit SDF representations and outputs triangle meshes that are directly applicable to rasterization-based renderers.

Let $\mathbf{V}_{\text{tet}}$, $\mathbf{F}_{\text{tet}}$, $\mathbf{T}_{\text{tet}}$ and be the pre-defined vertices, faces, and UV coordinates of the tetrahedra grid. We parameterize both per-tet vertice SDF value $\mathbf{S}$ and vertices offsets $\Delta\mathbf{V}_{\text{tet}}$ with a coordiante-based nerual net:

$$F_{\Phi_{geom}} : (\mathbf{V}_{\text{tet}}) \rightarrow (\mathbf{S}, \Delta\mathbf{V}_{\text{tet}}), \quad (2)$$

The canonical mesh $\mathbf{M}_{c} = (\mathbf{V}_{c}, \mathbf{F}_{c}, \mathbf{T}_{c})$ (*i.e,* canonical mesh vertices, faces, and UV map coordinates) is derived by marching tetrahedra operator Π:

$$\Pi : (\mathbf{V}_{\text{tet}}, \mathbf{F}_{\text{tet}}, \mathbf{T}_{\text{tet}}, \mathbf{S}, \Delta\mathbf{V}_{\text{tet}}) \rightarrow (\mathbf{V}_{c}, \mathbf{F}_{c}, \mathbf{T}_{c}). \quad (3)$$

Specifically, the vertices of the canonical mesh are computed by $\mathbf{v}_{c}^{ij} = \frac{\mathbf{v}_{tet}'^{i} s_j - \mathbf{v}_{tet}'^{j} s_i}{s_j - s_i}$, where $\mathbf{v}_{tet}'^{i} = \mathbf{v}^{i} + \Delta\mathbf{v}^{i}$ and $\text{sign}(s_i) \neq \text{sign}(s_j)$. In other words, only the edges that cross the surface of canonical mesh participate the marching tetrahedra operator, which makes the mesh extraction both computation and memory-efficient.

After training, we can discard the SDF and deformation neural nets $F_{\Phi_{geom}}$ and store the derived meshes. That leads to zero computation overhead in inference time.

3.2. Shading Model

Materials: we use a physically-based material model [56], which is directly applicable to our differentiable renderer. It consists of a diffuse term with an isotopic GGX lobe representing specularity. Concretely, it consists of three parts: 1) diffuse lobe $\mathbf{k}_d$ has four components, *i.e.* RGB color channels and an additional alpha channel; 2) specular lobe comprises a roughness value r for GGX normal distribution function and a metalness factor m which interpolates the sense of reality from plastic to pure metallic appearance. The specular highlight color is given by $\mathbf{k}_s = (1-m) \cdot 0.04 + m \cdot \mathbf{k}_d$. We store the specular lobe into texture $\mathbf{k}_{\text{orm}} = (o, r, m)$ following the convention, where the channel o is unused. To compensate for the global illumination, we alternatively store the ambient occlusion value into o. 3) normal maps $\mathbf{n}$ represents the fine-grained geometry details. The diffues color $\mathbf{k}_d$, texture $\mathbf{k}_{\text{orm}}$, and normal maps $\mathbf{n}$ are parametrized by an neural net:

$$F_{\Phi_{\text{mat}}} : (\mathbf{v}_c, \mathbf{P}) \rightarrow (\mathbf{k}_d, \mathbf{k}_{\text{orm}}, \mathbf{n}). \quad (4)$$

We query the vertices after rasterization and barycentric interpolation. The PBR material is further conditioned on poses to model the pose-dependent shading effect.

Lights: Our method learns a fixed environment light directly from the reference images. The lights are represented as a cube light. Given direction ω_o, We follow the render equation [32] to compute the outgoing radiance $L(\omega_o)$:

$$L(\omega_o) = \int_{\Omega} L_i(\omega_i) f(\omega_i, \omega_o) (\omega_i \cdot \mathbf{n}) d\omega_i, \quad (5)$$

The outgoing radiance is the integral of the incident radiance $L_i(\omega_i)$ and the BRDF $f(\omega_i, \omega_o)$. We do not use spherical Gaussians [13] or spherical harmonics [10, 97] to approximate the image-based lighting. Instead, we follow [61] using the split sum approximation that capable of modeling all-frequency image-based lighting:

$$L(\omega_o) \approx \int_{\Omega} f(\omega_i, \omega_o) (\omega_i \cdot \mathbf{n}) d\omega_i$$
$$\int_{\Omega} L_i(\omega_i) D(\omega_i, \omega_o) (\omega_i \cdot \mathbf{n}) d\omega_i. \quad (6)$$
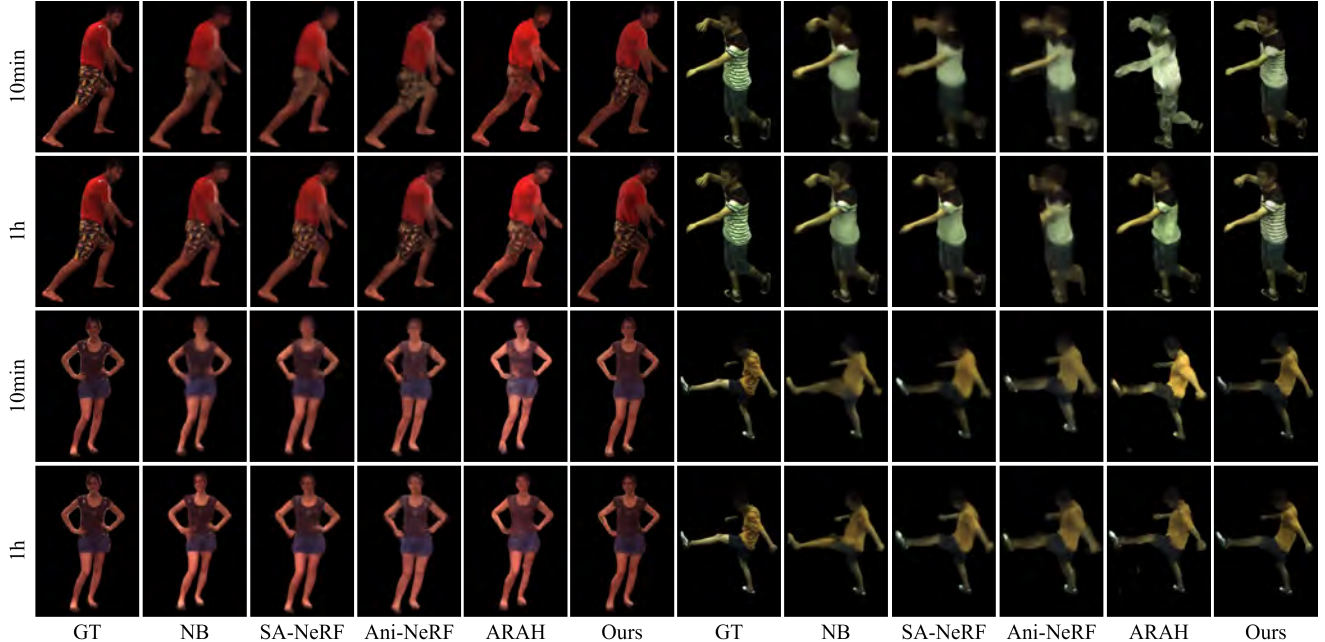
Figure 3: **Qualitative results of novel view synthesis on the H36M and ZJU-MoCap datasets**. [70, 69] generates blurry textures compared with our method. The mesh representations and forward skinning modeling help to improve generalization. Left: Hn36M dataset. Right: ZJU-MoCap dataset. Up: Training pose. Down: Novel pose. **Zoom in for a better view**.

The materials and lights are optimized jointly with geometry and motion modules in an end-to-end manner. The decomposed design of geometry and shading, along with compatibility with the triangle renderer enables editing and content creation instantly after training. For details about light modeling, please refer to the supplementary.

3.3. Motion Model

Since we derived mesh-based actors in canonical space with materials and lights, it is intuitive and natural to choose forward linear skinning as our motion model. Given a skeleton with B bones, the skeleton poses $\mathbf{P} = \{\mathbf{T}_1, \mathbf{T}_2, \ldots, \mathbf{T}_B\}$, where each $\mathbf{T}_i$ represents the transformation on bone i, and the blend skinning weights $\mathbf{W} = \{w_1, w_2, \ldots, w_B\}$, we deform each mesh vertex $\mathbf{v}_c$ in canonical space to the posed vertice $\mathbf{v}_w$ in world space by:

$$\mathbf{v}_w = \text{LBS}(\mathbf{v}_c, \mathbf{P}, \mathbf{W}) = (\sum_{i=1}^{B} w_i \mathbf{T}_i)\mathbf{v}_c, \qquad (7)$$

To compensate for non-rigid cloth dynamics, we further add a layer of pose-dependent non-rigid offsets $\Delta \mathbf{v}_c$ for canonical meshes:

$$\mathbf{v}_w = \text{LBS}(\mathbf{v}_c + \Delta \mathbf{v}_c, \mathbf{P}, \mathbf{W}), \qquad (8)$$

Where the blend skinning weights and the pose-dependent non-rigid offsets are, respectively, parameterized by neural nets whose inputs are canonical mesh vertices:

$$F_{\Phi_{\text{LBS}}} : (\mathbf{v}_c) \rightarrow \mathbf{W}, \qquad (9)$$

$$F_{\Phi_{\text{nr}}} : (\mathbf{v}_c, \mathbf{P}) \rightarrow \Delta \mathbf{v}_c. \qquad (10)$$

Modeling forward skinning is efficient for training as it only forward once in each optimization step, while the volume-based methods [45, 88, 15] solve the root-finding problem for canonical points in every iteration. After training, we can export the skinning weight from neural nets which removes the extra computation burden for inference.

4. Experiments

4.1. Dataset and Metrics

H36M consists of 4 multi-view cameras and uses **marker-based** motion capture to collect human poses. Each video contains a single subject performing a complex action. We follow [69] data protocol which includes subject S1, S5-S9, and S11. The videos are split into two parts: "training poses" for novel view synthesis and "Unseen poses" for novel pose synthesis. Among the video frames, 3 views are used for training, and the rest views are for evaluation. The novel view and novel pose metrics are computed on rest views. We use the same data preprocessing as [69].

ZJU-MoCap records 9 subjects performing complex actions with 23 cameras. The human poses are obtained with a markerless motion capture system. Thus the pose tracking

Table 1: **Quantitative results**. On the marker-based H36M, our method achieves SOTA performance in all optimization durations. While on the markerless ZJU-MoCap, our method is comparable with previous arts. "T.F." means template-free; "Rep." means representation; "T.T" means the training time; ∗ denotes the evaluation on a subset of validation splits.

	T.F.	Rep.	T.T.	H36M				ZJUMOCAP			
				Training pose		Novel pose		Training pose		Novel pose	
				PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
NB [70]		NV	~10 h	23.31	0.902	22.59	0.882	28.10	0.944	23.49	0.885
SA-NeRF [92]		NV	~30 h	24.28	0.909	23.25	0.892	28.27	0.945	24.42	0.902
Ani-NeRF [69]	✓	NV	~10 h	23.00	0.890	22.55	0.880	26.19	0.921	23.38	0.892
ARAH [88]	✓	NV	~48 h	24.79	0.918	23.42	0.896	28.51	0.948	24.63	0.911
Ours	✓	Hybr	~1 h	24.72	0.916	23.64	0.899	26.57	0.901	24.38	0.875
NB		NV		20.58	0.879	20.27	0.867	26.87	0.922	23.67	0.885
SA-NeRF		NV		21.03	0.878	20.71	0.869	24.92	0.882	23.38	0.869
Ani-NeRF	✓	NV	~1 h*	22.54	0.872	21.79	0.856	21.23	0.659	20.65	0.652
ARAH	✓	NV		24.25	0.904	23.61	0.892	26.33	0.924	24.67	0.911
Ours	✓	Hybr		24.83	0.917	23.64	0.899	26.66	0.901	24.64	0.880
NB		NV		20.54	0.863	20.15	0.853	25.37	0.894	23.54	0.873
SA-NeRF		NV		20.81	0.848	20.49	0.841	24.48	0.878	23.75	0.872
Ani-NeRF	✓	NV	~10 m*	20.57	0.822	20.22	0.806	21.17	0.652	21.16	0.656
ARAH	✓	NV		23.83	0.895	23.13	0.884	25.09	0.906	24.21	0.898
Ours	✓	Hybr		24.27	0.909	23.37	0.897	25.51	0.888	24.42	0.878

is rather noisier compared with H36M. Likewise, there are two sets of video frames, "training poses" for novel view synthesis and "Unseen poses" for novel pose synthesis. 4 evenly distributed camera views are chosen for training, and the rest 19 views are for evaluation. Again, the evaluation metrics are computed on rest views. The same data protocol and processing approaches are adopted following [70, 69].

Metrics. We follow the typical protocol in [70, 69] using two metrics to measure image quality: peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM).

4.2. Evaluation and Comparison

Baselines. We compare our method with template-based methods [70, 92] and template-free methods [69, 88]. Here we list the average metric values with different training times to illustrate our very competitive performance and significant speed boost. 1) Tempelate-based methods. Neural Body (NB) [70] learns a set of latent codes anchored to a deformable template mesh to provide geometry guidance. Surface-Aligned NeRF (SA-NeRF) [92] proposes projecting a point onto a mesh surface to align surface points and signed height to the surface. 2) Template-free methods. Animatable NeRF (Ani-NeRF) [69] introduces neural blend weight fields to produce the deformation fields instead of explicit template control. ARAH [88] combines an articulated implicit surface representation with volume rendering and proposes a novel joint root-finding algorithm.

Comparisons with state-of-the-arts. Table 1 illustrates the quantitative comparisons with previous arts. Notably,

our method achieves very competitive performance within much less training time. The previous volume rendering-based counterparts spend tens of hours of optimization time, while our method only takes an hour of training (for previous SOTA method ARAH [88], it takes about 2 days of training). On the marker-based H36M dataset, our method reaches the SOTA performance in terms of novel view synthesis on training poses and outperforms previous SOTA (ARAH [88]) for novel view synthesis on novel poses, which indicates that our method can generalize better on novel poses. The significant boost in training speeds lies in, on the one hand, the core mesh representation which can be rendered efficiently with the current graphic pipeline [43]. On the other hand, the triangular renderer uses less memory. Thus we can compute losses over the whole image to learn the representations holistically. In contrast, previous methods are limited to much fewer sampled pixels in each optimization step.

On the markerless ZJU-Mocap dataset, our method falls behind for training poses novel view synthesis and ranks 3rd place in terms of unseen poses novel view synthesis among the competitors. We argue that the quality of pose tracking results in the performance gaps between the two datasets. The markerless pose tracking data **are much noisier** than the marker-based ones (*e.g.*, the tracked skeleton sequence is jittering, and the naked human [53] rendering is misaligned with human parsings), which makes our performance saturated by harming the multi-view consistency. The problem is even amplified with the holistic loss compu-

tation over the whole pixels. We conduct additional ablation on pose tracking quality on H36M in Sec. 4.3. Besides, our non-rigid modeling is only over the surface (no topology change), which is less powerful than the volume rendering-based ones (with topology change).

We further each method under **the same optimization duration** in Table 1. For the extremely low inference speed of our competitor, we only evaluate at most 10 frames in each subject, and for ZJU MoCap we only choose another 4 evenly spaced cameras as the evaluation views. For both 1 hour and 10 minutes optimization time, our method outperforms other methods for both training poses and unseen poses novel view synthesis on the marker-based H36M dataset. On the markerless ZJU-Mocap dataset, our method is comparable with previous SOTA in terms of PSNR and SSIM for both evaluation splits.

Figure 3 illustrates the qualitative comparisons between our method and previous arts under the same optimization duration. It is worth noting that on both H36M and ZJU-MoCap datasets, our method can synthesize clearer and more fine-grained images against competitors, which raises the misalignment of the quantitative metrics for measuring image similarity.

Rendering Efficiency: We provide the rendering speed of our method against previous methods in Figure 4. Our method reaches real-time inference speed (100+ FPS for rendering 512×512 images), which is hundreds
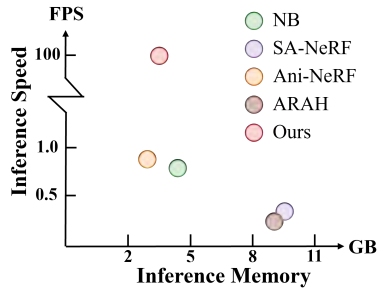

Figure 4: **Rendering Efficiency**

of times faster than the previous ones. Our method takes considerably less memory than the previous ones.

4.3. Ablation Studies

We conduct ablation studies on the H36M S9 subject.
The parametrization type for SDF field. The SDF fields can either be parameterized as either MLPs or value fields. Table 2 and Figure 5 show that using MLP to predict SDF values results in a smoother mesh surface that is watertight. MLP offers extrapolation ability to predict invisible parts and keep the mesh watertight. While directly optimizing SDF value fields leads to a jiggling mesh surface and holes in invisible parts during training (*e.g.*, underarm).
The shading model type in geometry module. We compare PBR shading models with directly predicting RGB colors and PBR without shading specular. Table 2 and Figure 5 show that PBR shading models lead to higher metrics against RGB predications, which indicates that PBR mate-

Table 2: **The ablation on each module from our method**. The mesh tends to be noisy and poor for rendering novel poses without MLP parametrization for the geometry module; Removing the non-rigid module harms the convergence of our model due to the disability to solve multi-view inconsistency; PBR materials improve the overall shading quality by joint modeling both decomposed materials and lighting.

	Training Pose		Novel Pose	
	PSNR	SSIM	PSNR	SSIM
w/o SDF MLP	25.17	0.913	23.37	0.874
w/o Non-rigid	25.03	0.909	23.45	0.877
w/o PBR	25.10	0.914	23.44	0.878
w/o Specular	25.24	0.915	**23.58**	**0.879**
Full	**25.26**	**0.916**	23.52	**0.879**

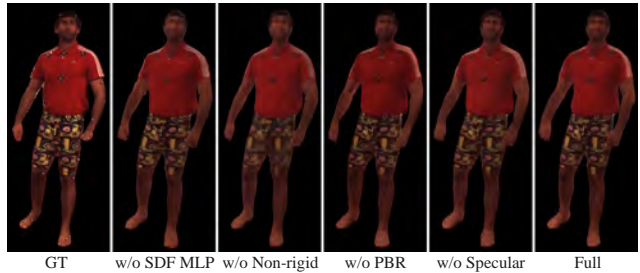GT   w/o SDF MLP   w/o Non-rigid   w/o PBR   w/o Specular   Full

Figure 5: **Qualitative ablation on each module**. The SDF MLP improves the mesh smoothness; non-rigid modeling proves the texture quality by solving the multi-view consistency of cloth dynamics; The PBR materials have a larger capacity for modeling complex materials and lighting against the only-RGB and the no-specular counterparts, which further facilitates both mesh and material learning.

rials can better model complex textures and lights for dynamic humans. Removing the specular term in PBR does not affect the performance much. We conjecture that there is less specularity in human skin and clothes materials.
The impact of the non-rigid net in motion module. As shown in Table 2 and Figures 5, modeling pose-dependent non-rigid dynamics of clothes improves the overall reconstruction quality. It facilitates the aggregation of shading information for multi-view inputs during training.
The impact of human tracking quality. Table 3 (a) and Figure 6 show that using marker-based pose-tracking data can give better results. The same phenomenon has been stated in [69]. Noisy marker-less pose-tracking harms the optimization process by damaging the multi-view consistency and the exact pose for shading optimization, which leads to blurry textures.
The impact of training view amount. Table 3 (b) and Figure 7 reveal that giving one camera of view degrades the
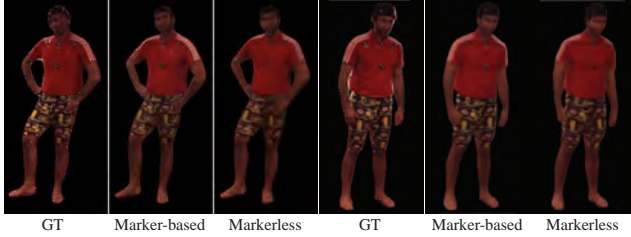
GT Marker-based Markerless GT Marker-based Markerless

Figure 6: **Qualitative results of models trained on poses** from marker-less and marker-based systems.
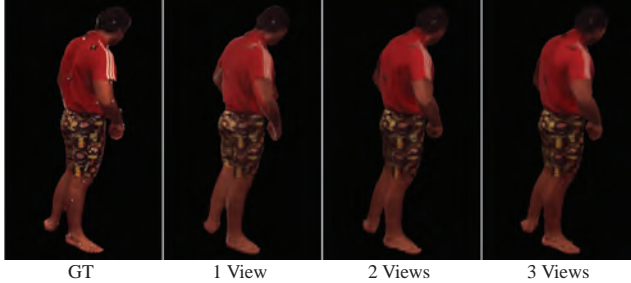

GT 1 View 2 Views 3 Views

Figure 7: **Comparison of models trained with different numbers of camera views** on the subject "S9".

Table 3: **The ablations results on data quality and quantity** on H36M S9 subject, in terms of PSNR and SSIM (higher is better). The better the data quality, the better the reconstruction results.

	Training pose		Novel pose	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
(a) type of pose tracking				
w/o marker	24.73	0.893	22.60	0.853
w/ marker	**25.53**	**0.911**	**23.80**	**0.879**
(b) number of training views				
1 view	25.09	0.906	22.97	0.866
2 views	25.56	**0.911**	**23.76**	**0.878**
3 views	**25.57**	**0.911**	23.67	0.876
(c) number of training frames				
1 frame	20.93	0.817	19.58	0.785
100 frames	23.99	0.882	22.49	0.856
200 frames	**25.27**	**0.905**	**23.32**	**0.873**
800 frames	24.89	0.900	23.16	**0.873**

overall reconstruction quality, and multi-view consistency improves the final results. The model can aggregate multi-view information for better shading optimization, thus leading to clearer surface materials.

The impact of training frame amount. As the number of training frames increases, the rendering quality on novel view and novel pose increases as well (Table 3 (c) and Fig-
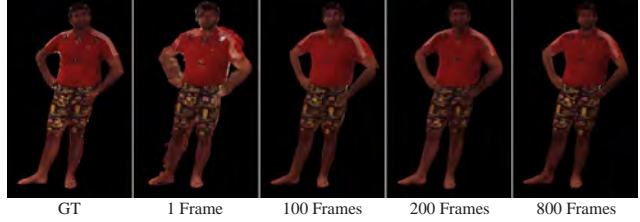

GT 1 Frame 100 Frames 200 Frames 800 Frames

Figure 8: **Comparison of models trained with different numbers of video frames** on the subject "S9".

ure 8). Notice that the reconstruction quality saturated after using a certain amount of training frames, the same results can be observed in [69] as well.

4.4. Applications

After training, we can export mesh representations, which enables instant downstream applications. We showcase two examples of novel pose synthesis, material editing, and human relighting in Figure 1. For more examples, please refer to our supplementary.

5. Discussions and Conclusions

Discussions: Our method leverage mesh as our core representation, which enables us efficiency for both training and rendering. However, the resolution of mesh is fixed in our pipeline, preventing fine-grained geometry and texture reconstruction. One possible solution could be tetrahedra grids sub-division [78, 21, 33]. But it may break the SDF values around the derived meshes since there is no regularization over the whole SDF field. Our non-rigid modeling has less capacity, since we assume there is no topology change of mesh *wrt.* the non-rigid motion. Otherwise, we cannot query materials and motions in the canonical shape. One can solve it via the dense correspondence between the meshes before and after applying non-rigid motions [1, 94], yet such an operation may increase computation drastically.

Conclusions: we present EMA, which learns human avatars through hybrid meshy neural fields efficiently. EMA jointly learns hybrid canonical geometry, materials, lights, and motions via a rasterization-based differentiable renderer. It only requires one hour of training and can render in real-time with a triangle renderer. Minutes of training can produce plausible results. Our method enjoys flexibility from implicit representations and efficiency from explicit meshes. Experiments on the standard benchmark indicate the competitive performance and generalization results of our method. The digitized avatars can be directly used in downstream tasks. We showcase examples including novel pose synthesis, material editing, and human relighting.

References

[1] Naveed Ahmed, Christian Theobalt, Christian Rössl, Sebastian Thrun, and Hans-Peter Seidel. Dense correspondence finding for parametrization-free animation reconstruction from video. In *CVPR*, pages 1–8, 2008.

[2] Oleg Alexander, Mike Rogers, William Lambeth, Jen-Yuan Chiang, Wan-Chun Ma, Chuan-Chang Wang, and Paul E. Debevec. The digital emily project: Achieving a photorealistic digital actor. *IEEE Comput. Graph. Appl.*, 30(4):20–31, 2010.

[3] Thiemo Alldieck, Marcus A. Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Detailed human avatars from monocular video. In *3DV*, pages 98–109, 2018.

[4] Thiemo Alldieck, Marcus A. Magnor, Weipeng Xu, Christian Theobalt, and Gerard Pons-Moll. Video based reconstruction of 3d people models. In *CVPR*, pages 8387–8397, 2018.

[5] Thiemo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3d reconstruction of humans wearing clothing. In *CVPR*, pages 1496–1505, 2022.

[6] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: shape completion and animation of people. *ACM Trans. Graph.*, 24(3):408–416, 2005.

[7] Jan Bender, Matthias Müller, Miguel A. Otaduy, Matthias Teschner, and Miles Macklin. A survey on position-based simulation methods in computer graphics. *Comput. Graph. Forum*, 33(6):228–251, 2014.

[8] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, pages 187–194, 1999.

[9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter V. Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: automatic estimation of 3d human pose and shape from a single image. In *ECCV*, pages 561–578, 2016.

[10] Mark Boss, Raphael Braun, Varun Jampani, Jonathan T. Barron, Ce Liu, and Hendrik P. A. Lensch. Nerd: Neural reflectance decomposition from image collections. In *ICCV*, pages 12664–12674, 2021.

[11] Lu Chen, Jiao Sun, and Wei Xu. FAWA: fast adversarial watermark attack on optical character recognition (OCR) systems. In *ECML*, pages 547–563, 2020.

[12] Mingfei Chen, Jianfeng Zhang, Xiangyu Xu, Lijuan Liu, Yujun Cai, Jiashi Feng, and Shuicheng Yan. Geometry-guided progressive nerf for generalizable and efficient neural human rendering. In *ECCV*, pages 222–239, 2022.

[13] Wenzheng Chen, Huan Ling, Jun Gao, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *NeurIPS*, pages 9605–9616, 2019.

[14] Wenzheng Chen, Joey Litalien, Jun Gao, Zian Wang, Clement Fuji Tsang, Sameh Khamis, Or Litany, and Sanja Fidler. DIB-R++: learning to predict lighting and material with a hybrid differentiable renderer. In *NeurIPS*, pages 22834–22848, 2021.

[15] 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*, pages 11574–11584, 2021.

[16] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1):7–24, 1982.

[17] Haoshu Fang, Shuqin Xie, Yu-Wing Tai, and Cewu Lu. RMPE: regional multi-person pose estimation. In *ICCV*, pages 2353–2362, 2017.

[18] Blender Foundation. Blender.org - Home of the Blender project - Free and Open 3D Creation Software.

[19] Jun Gao, Wenzheng Chen, Tommy Xiang, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Learning deformable tetrahedral meshes for 3d reconstruction. In *NeurIPS*, 2020.

[20] 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. *arXiv:2209.11163*, 2022.

[21] William Gao, April Wang, Gal Metzer, Raymond A. Yeh, and Rana Hanocka. Tetgan: A convolutional neural network for tetrahedral mesh generation. In *BMVC*, page 365, 2022.

[22] Ke Gong, Xiaodan Liang, Yicheng Li, Yimin Chen, Ming Yang, and Liang Lin. Instance-Level Human Parsing via Part Grouping Network. In *ECCV*, pages 805–822, 2018.

[23] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular RGB videos. In *CVPR*, pages 18632–18643, 2022.

[24] 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. *arXiv:2302.11566*, 2023.

[25] Marc Habermann, Weipeng Xu, Michael Zollhöfer, Gerard Pons-Moll, and Christian Theobalt. Deepcap: Monocular human performance capture using weak supervision. In *CVPR*, pages 5051–5062, 2020.

[26] Jon Hasselgren, Nikolai Hofmann, and Jacob Munkberg. Shape, light & material decomposition from images using monte carlo rendering and denoising. *arXiv:2206.03380*, 2022.

[27] Jon Hasselgren, Jacob Munkberg, Jaakko Lehtinen, Miika Aittala, and Samuli Laine. Appearance-Driven Automatic 3D Model Simplification. *arXiv:2104.03989*, 2021.

[28] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH++: animation-ready clothed human reconstruction revisited. In *ICCV*, pages 11026–11036, 2021.

[29] Zeng Huang, Yuanlu Xu, Christoph Lassner, Hao Li, and Tony Tung. ARCH: animatable reconstruction of clothed humans. In *CVPR*, pages 3090–3099, 2020.

[30] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(7):1325–1339, 2014.

[31] Wei Jiang, Kwang Moo Yi, Golnoosh Samei, Oncel Tuzel, and Anurag Ranjan. Neuman: Neural human radiance field from a single video. In *ECCV*, pages 402–418, 2022.

[32] James T. Kajiya. The rendering equation. In *SIGGRAPH*, pages 143–150, 1986.

[33] Nikolai Kalischek, Torben Peters, Jan D. Wegner, and Konrad Schindler. Tetrahedral diffusion models for 3d shape generation. *arXiv:2211.13220*, 2022.

[34] Angjoo Kanazawa, Michael J. Black, David W. Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, pages 7122–7131, 2018.

[35] Brian Karis. Real shading in unreal engine 4. *SIGGRAPH 2013 Course: Physically Based Shading in Theory and Practice*, 2013.

[36] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(12):4217–4228, 2021.

[37] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8107–8116, 2020.

[38] Taras Khakhulin, Vanessa Sklyarova, Victor Lempitsky, and Egor Zakharov. Realistic one-shot mesh-based head avatars. In *ECCV*, pages 345–362, 2022.

[39] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.

[40] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: video inference for human body pose and shape estimation. In *CVPR*, pages 5252–5262, 2020.

[41] Muhammed Kocabas, Chun-Hao P. Huang, Otmar Hilliges, and Michael J. Black. PARE: part attention regressor for 3d human body estimation. In *ICCV*, pages 11107–11117, 2021.

[42] François Labelle and Jonathan Richard Shewchuk. Isosurface stuffing: fast tetrahedral meshes with good dihedral angles. *ACM Trans. Graph.*, 26(3):57, 2007.

[43] Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph.*, 39(6):194:1–194:14, 2020.

[44] Christoph Lassner and Michael Zollhöfer. Pulsar: Efficient sphere-based neural rendering. In *CVPR*, pages 1440–1449, 2021.

[45] Ruilong Li, Julian Tanke, Minh Vo, Michael Zollhöfer, Jürgen Gall, Angjoo Kanazawa, and Christoph Lassner. TAVA: template-free animatable volumetric actors. In *ECCV*, pages 419–436, 2022.

[46] Ruilong Li, Sha Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, pages 13381–13392, 2021.

[47] Xiaoting Li, Lingwei Chen, Jinquan Zhang, James R. Larus, and Dinghao Wu. Watermarking-based defense against adversarial attacks on deep neural networks. In *International Joint Conference on Neural Networks, IJCNN 2021, Shenzhen, China, July 18-22, 2021*, pages 1–8, 2021.

[48] Hsueh-Ti Derek Liu, Francis Williams, Alec Jacobson, Sanja Fidler, and Or Litany. Learning smooth neural functions via lipschitz regularization. In *SIGGRAPH*, pages 31:1–31:13, 2022.

[49] 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 Trans. Graph.*, 40(6):219:1–219:16, 2021.

[50] Shichen Liu, Weikai Chen, Tianye Li, and Hao Li. Soft rasterizer: A differentiable renderer for image-based 3d reasoning. In *ICCV*, pages 7707–7716, 2019.

[51] Steven Liu, Xiuming Zhang, Zhoutong Zhang, Richard Zhang, Jun-Yan Zhu, and Bryan Russell. Editing conditional radiance fields. In *ICCV*, pages 5753–5763, 2021.

[52] Stephen Lombardi, Tomas Simon, Jason M. Saragih, Gabriel Schwartz, Andreas M. Lehrmann, and Yaser Sheikh. Neural volumes: learning dynamic renderable volumes from images. *ACM Trans. Graph.*, 38(4):65:1–65:14, 2019.

[53] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):248:1–248:16, 2015.

[54] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *SIGGRAPH*, pages 163–169, 1987.

[55] Yi Ma, Stefano Soatto, Jana Kosecka, and S. Shankar Sastry. *An Invitation to 3-D Vision*. Springer, 2004.

[56] Stephen McAuley, Stephen Hill, Naty Hoffman, Yoshiharu Gotanda, Brian E. Smits, Brent Burley, and Adam Martinez. Practical physically-based shading in film and game production. In *SIGGRAPH*, pages 10:1–10:7, 2012.

[57] Lars M. Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, pages 4460–4470, 2019.

[58] 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. *Commun. ACM*, 65(1):99–106, 2022.

[59] Tomas Möller, Eric Haines, and Nathaniel Hoffman. *Real-time rendering, 3rd Edition*. Peters, 2008.

[60] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022.

[61] Jacob Munkberg, Wenzheng Chen, Jon Hasselgren, Alex Evans, Tianchang Shen, Thomas Müller, Jun Gao, and Sanja Fidler. Extracting triangular 3d models, materials, and lighting from images. In *CVPR*, pages 8270–8280, 2022.

[62] Thanh Nguyen, Quoc Viet Hung Nguyen, Dung Tien Nguyen, Duc Thanh Nguyen, Thien Huynh-The, Saeid Nahavandi, Thanh Tam Nguyen, Quoc-Viet Pham, and Cuong M. Nguyen. Deep learning for deepfakes creation and detection: A survey. *Comput. Vis. Image Underst.*, 223:103525, 2022.

[63] Atsuhiro Noguchi, Xiao Sun, Stephen Lin, and Tatsuya Harada. Neural articulated radiance field. In *ICCV*, pages 5742–5752, 2021.

[64] Deng Pan, Lixian Sun, Rui Wang, Xingjian Zhang, and Richard O. Sinnott. Deepfake detection through deep learning. In *BDCAT*, pages 134–143, 2020.

[65] Jeong Joon Park, Peter Florence, Julian Straub, Richard A. Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, pages 165–174, 2019.

[66] 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 5845–5854, 2021.

[67] David A. Patterson, Joseph Gonzalez, Quoc V. Le, Chen Liang, Lluis-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training. *arXiv:2104.10350*, 2021.

[68] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. 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.

[69] 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*, pages 14294–14303, 2021.

[70] 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*, pages 9054–9063, 2021.

[71] Sergey Prokudin, Michael J. Black, and Javier Romero. Smplpix: Neural avatars from 3d human models. In *WACV*, pages 1809–1818, 2021.

[72] Amit Raj, Julian Tanke, James Hays, Minh Vo, Carsten Stoll, and Christoph Lassner. ANR: articulated neural rendering for virtual avatars. In *CVPR*, pages 3722–3731, 2021.

[73] Nikhila Ravi, Jeremy Reizenstein, David Novotný, Taylor Gordon, Wan-Yen Lo, Justin Johnson, and Georgia Gkioxari. Accelerating 3d deep learning with pytorch3d. *arXiv:2007.08501*, 2020.

[74] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Trans. Graph.*, 36(6):245:1–245:17, 2017.

[75] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Hao Li, and Angjoo Kanazawa. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, pages 2304–2314, 2019.

[76] Shunsuke Saito, Tomas Simon, Jason M. Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, pages 81–90, 2020.

[77] Shunsuke Saito, Jinlong Yang, Qianli Ma, and Michael J. Black. Scanimate: Weakly supervised learning of skinned clothed avatar networks. In *CVPR*, pages 2886–2897, 2021.

[78] Scott Schaefer, Jan Hakenberg, and Joe D. Warren. Smooth subdivision of tetrahedral meshes. In *SGP*, pages 147–154, 2004.

[79] 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*, pages 6087–6101, 2021.

[80] Michael Stokes, Matthew Anderson, Srinivasan Chandrasekar, and Ricardo Motta. A Standard Default Color Space for the Internet - sRGB, 1996.

[81] 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*, pages 12278–12291, 2021.

[82] Yu Sun, Qian Bao, Wu Liu, Yili Fu, Michael J. Black, and Tao Mei. Monocular, one-stage, regression of multiple 3d people. In *ICCV*, pages 11159–11168, 2021.

[83] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: image synthesis using neural textures. *ACM Trans. Graph.*, 38(4):66:1–66:12, 2019.

[84] Ingo Wald, Will Usher, Nathan Morrical, Laura Lediaev, and Valerio Pascucci. RTX beyond ray tracing: Exploring the use of hardware ray tracing cores for tet-mesh point location. In *HPG*, pages 7–13, 2019.

[85] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the Eurographics Symposium on Rendering Techniques, Grenoble, France, 2007*, pages 195–206, 2007.

[86] Liao Wang, Ziyu Wang, Pei Lin, Yuheng Jiang, Xin Suo, Minye Wu, Lan Xu, and Jingyi Yu. ibutter: Neural interactive bullet time generator for human free-viewpoint rendering. In *ACMMM*, pages 4641–4650, 2021.

[87] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. In *NeurIPS*, pages 27171–27183, 2021.

[88] Shaofei Wang, Katja Schwarz, Andreas Geiger, and Siyu Tang. ARAH: animatable volume rendering of articulated human sdfs. In *ECCV*, pages 1–19, 2022.

[89] 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*, pages 16189–16199, 2022.

[90] Donglai Xiang, Timur M. Bagautdinov, Tuur Stuyck, Fabian Prada, Javier Romero, Weipeng Xu, Shunsuke Saito, Jingfan Guo, Breannan Smith, Takaaki Shiratori, Yaser Sheikh, Jessica K. Hodgins, and Chenglei Wu. Dressing avatars: Deep photorealistic appearance for physically simulated clothing. *ACM Trans. Graph.*, 41(6):222:1–222:15, 2022.

[91] Donglai Xiang, Fabian Prada, Timur M. Bagautdinov, Weipeng Xu, Yuan Dong, He Wen, Jessica K. Hodgins, and Chenglei Wu. Modeling clothing as a separate layer for an animatable human avatar. *ACM Trans. Graph.*, 40(6):199:1–199:15, 2021.

[92] Tianhan Xu, Yasuhiro Fujita, and Eiichi Matsumoto. Surface-aligned neural radiance fields for controllable 3d human synthesis. In *CVPR*, pages 15862–15871, 2022.

[93] 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*, pages 18332–18343, 2022.

[94] Wang Zeng, Wanli Ouyang, Ping Luo, Wentao Liu, and Xiaogang Wang. 3d human mesh regression with dense correspondence. In *CVPR*, pages 7052–7061, 2020.

[95] Hongwen Zhang, Yating Tian, Yuxiang Zhang, Mengcheng Li, Liang An, Zhenan Sun, and Yebin Liu. Pymaf-x: Towards well-aligned full-body model regression from monocular images. *arXiv:2207.06400*, 2022.

[96] Hongwen Zhang, Yating Tian, Xinchi Zhou, Wanli Ouyang, Yebin Liu, Limin Wang, and Zhenan Sun. Pymaf: 3d human pose and shape regression with pyramidal mesh alignment feedback loop. In *ICCV*, pages 11426–11436, 2021.

[97] Kai Zhang, Fujun Luan, Qianqian Wang, Kavita Bala, and Noah Snavely. Physg: Inverse rendering with spherical gaussians for physics-based material editing and relighting. In *CVPR*, pages 5453–5462, 2021.

[98] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020.

[99] Ruiqi Zhang and Jie Chen. NDF: neural deformable fields for dynamic human modelling. In *ECCV*, pages 37–52, 2022.

[100] Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul E. Debevec, William T. Freeman, and Jonathan T. Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40:237:1–237:18, 2021.

[101] Hao Zhao, Jinsong Zhang, Yu-Kun Lai, Zerong Zheng, Yingdi Xie, Yebin Liu, and Kun Li. High-fidelity human avatars from a single RGB camera. In *CVPR*, pages 15883–15892, 2022.

[102] Zerong Zheng, Han Huang, Tao Yu, Hongwen Zhang, Yandong Guo, and Yebin Liu. Structured local radiance fields for human avatar modeling. In *CVPR*, pages 15872–15882, 2022.

Appendix

Thank you for reading our supplementary materials! Here we provide depth descriptions of our method, including details about loss functions (Sec. 6), image-based lighting (Sec. 7), and image-based lighting (Sec. 8). Then we present additional ablations about training views (Sec. 9) and skinning module (Sec. 9). Additional experimental results are illustrated in Sec. 10 and Sec. 12. We showcase application examples in Sec. 11. In the end, we discuss limitations and social impact in Sec. 13. We strongly encourage our readers to view the supplemental video for a more comprehensive visual perception.

6. Loss Functions

Our loss function $L = L_{\text{img}} + L_{\text{mask}} + L_{\text{reg}}$ is composed of three parts: an image loss L_{img} using ℓ_1 norm on tone mapped colors, and mask loss L_{mask} using squared ℓ_2, and regularization losses L_{reg} to improve the quality of canonical geometry, materials, lights, and motion.

Image loss: our renderer utilizes physically-based shading to produce high-dynamic range (HDR) images. Then the complex materials and environmental lights are elaborately optimized. Thus our loss function requires a full range of floating point values. We follow [27, 61, 26] to compute ℓ_1 norm on tone mapped colors. Specifically, we first transform linear radiance values i according to a tone-mapping operator $T(i) = \Gamma(\log(i + 1))$, in which $\Gamma(i)$ is a linear RGB to sRGB transformation function [80]:

$$\Gamma(i) = \begin{cases} 12.92i & i \leq 0.0031308 \\ (1 + a)i^{1/2.4} - a & i > 0.0031308 \end{cases} \quad (11)$$
$$a = 0.055,$$

Mask loss: The renderer [43] renders both the shaded images and the corresponding rasterization masks in a differentiable manner. Therefore, we compute the ℓ_2 norm between the masks and the preprocessed mattings (in both ZJU-MoCap and H36M benchmarks, we use the provided preprocessed subject masks from [70, 22]), akin to the traditional shape-from-silhouette [55] technique. The mask loss is parallel with the image loss, yet facilitates the course of shading optimization by making shape convergence super fast in about a hundred training steps.

Regularizers: We need various priors to encourage the optimization to converge at a place where the geometry, materials, and lighting are well separated and smooth enough [61, 26]. Therefore, we choose to minimize regularization during training.

We introduce smoothness to PBR materials in terms of albedo $\mathbf{k}_d$, specular parameters $\mathbf{k}_{\text{orm}}$, and surface geometry nomral $\mathbf{n}$ as following:

$$L_{\mathbf{k}} = \frac{1}{|\mathbf{x}_{\text{surf}}|} \sum_{\mathbf{x}_{\text{surf}}} |\mathbf{k}(\mathbf{x}_{\text{surf}}) - \mathbf{k}(\mathbf{x}_{\text{surf}} + \epsilon)|, \quad (12)$$

where $|\mathbf{x}_{\text{surf}}|$ is a surface point on the surface in canonical space and $\epsilon \sim \mathcal{N}(0, \sigma = 0.01)$ is a small random offset. We regularize the geometry normal on the surface of the canonical mesh derived from the SDF field for a seek of a smoother surface and avoidance of holes in the surface.

We regularize light by assuming the neutral spectrum in the real world. Specifically, given the per-channel average radiance densities $\bar{c}_i$, we penalize the color shifts as:

$$L_{\text{light}} = \frac{1}{3} \sum_{i=0}^{3} \left| c_i - \frac{1}{3} \sum_{i=0}^{3} c_i \right|, \quad (13)$$

To encourage a watertight surface and reduce floating meshes both inside and outside the subject, we impose regularizations on the SDF field as:

$$L_{\text{sdf}} = \sum_{i,j \in \text{S}_e} H(\sigma(s_i), \text{sign}(s_j)) \\ + H(\sigma(s_j), \text{sign}(s_i)), \quad (14)$$

where S_e is the set of all vertex along their edges in which the signs of the SDF values are different (*i.e.*, $\text{sign}(s_i) \neq \text{sign}(s_j)$). To remove the floating meshes outside the surface, we impose an additional loss. For a triangle surface f extracted by marching tetrahedra, if f is invisible, we encourage its SDF values to be positive as:

$$L_{\text{invis}} = \sum_{i \in \text{S}_{\text{invis}}} H(\sigma(s_i), 1). \quad (15)$$

We weigh the above terms and use the loss for all our experiments:

$$L = L_{\text{image}} + L_{\text{mask}} \\ + \underbrace{\lambda_{\mathbf{k}_d}}_{=0.03} L_{\mathbf{k}_d} + \underbrace{\lambda_{\mathbf{k}_{\text{orm}}}}_{=0.05} L_{\mathbf{k}_{\text{orm}}} + \underbrace{\lambda_{\mathbf{n}}}_{=0.025} L_{\mathbf{n}} \\ + \underbrace{\lambda_{\text{light}}}_{=0.005} L_{\text{light}} + \underbrace{\lambda_{\text{sdf}}}_{=0.02} L_{\text{sdf}} + \underbrace{\lambda_{\text{invis}}}_{=0.01} L_{\text{invis}}. \quad (16)$$

7. Image-based Lighting

The split sum shading model is widely used in real-time rendering [59], giving both realism and efficiency against spherical Gaussians (SG) and spherical harmonics (SH) [10, 13, 97]. We use a differentiable split sum [35] shading model to approximate rendering equation [32] for image-based environment light learning as [61]:

$$L(\omega_o) \approx \int_\Omega f(\omega_i, \omega_o)(\omega_i \cdot \mathbf{n}) d\omega_i \\ \int_\Omega L_i(\omega_i) D(\omega_i, \omega_o)(\omega_i \cdot \mathbf{n}) d\omega_i. \quad (17)$$

where D is the GGX normal distribution function (NDF) [85] in a Cook-Torrance microfacet specular shading model [16]. The first term contributes to the specular BSDF *wrt.* a solid white environment light, which depends solely on the roughness r of the BSDF and the light-surface angles $\cos\theta = \omega_i \cdot \mathbf{n}$. The second term contributes to the integral of the incoming radiance with the GGX normal distribution function, D. Both terms can be pre-computed and filtered to reduce computation [35].

The training parameters are texels of a cube light map whose resolution is $6 \times 512 \times 512$. The pre-integrated lighting for the least roughness values is derived from the base level, and multiple smaller mip levels are constructed from it [35]. Each mip-map is filtered by average-pooling the base level of the current resolution and is convolved with the GGX normal distribution function. The per mip-level filter bounds are pre-computed as well. We leverage a PyTorch implementation with CUDA extensions from [61]. Moreover, a cube map is created to represent the diffuse lighting in a low resolution, akin to the filtered specular probe. It shares the same optimizable parameters and is average-pooled to the mip level with $r = 1$ roughness. The pre-filtering only involves the first term in Eq. 17.

8. Implementation Details

SDF network. We parametrize the SDF field with an MLP to increase surface water-tight and smoothness. We choose the MLP architecture from [58], which consists of 6 frequency bands for positional encoding, and 8 linear layers, each having 256 neurons, followed by ReLU activations. We implicitly regularize the smoothness by increasing the Lipschitz property in the SDF field[48].

Material network. The material model is a small MLP with hash-encoding [60] as the materials query is computationally extensive. The MLP consists of two linear layers, each having 32 neurons, followed by ReLU activations. The hash-encoding has a spatial resolution of 4096 and the rest configures are the same as [61]. To reduce computation, we predict all material channels at once with one backbone network. Besides, we introduce inductive bias of materials of clothed humans in the real world, by providing minimum and maximum values for each materials channel. We follow [100] to limit the albedo $\mathbf{k}_d \in [0.03, 0.8]$, and the roughness $\mathbf{k}_r \in [0.08, 1]$. The texels in the environment light are randomly initialized between $[0.25, 0.75]$.

Motion networks. For the motion module, we use the same MLP architecture as [15, 88], which is similar to our SDF MLP. To resolve the problem where the training pose variation is too limited for skinning field learning (*e.g.*, self-rotation video without any limbs movements), we initialize the MLP with the pre-trained skinning model provided by [88], and impose ℓ_2 norm for the skinning weights logits between our predictions and the ground truth from SMPL [53]. We ablate the design choices in Sec. 9. For the non-rigid modeling, we use another 4-layer ReLU MLP with a 4-frequency-band positional encoding. We also progressively anneal its encoding for 5k iterations as [66]. The weights of the last layer are initialized with a uniform distribution $\mathcal{U}(-10^{-5}, 10^{-5})$, *i.e.* initializing the non-rigid offsets to be close to zero and not interfering with the major optimizations of geometry and materials.

Optimization. We use Adam [39] as our default optimizer. We optimize the subject for 5k steps for 1024×1024 images or 10k steps 512×512 images. We disable the perturbed normal map during optimization as it leads to SDF collapsing abruptly at a certain step (*i.e.*, all SDF values are positive or negative where marching tetrahedra fails). The optimization process takes about an hour on a single NVIDIA GTX3090 GPU. The indicative results with plausible quality appear after a few minutes, which is quite faster than our counterparts [70, 69, 88, 92]. Such superior efficiency could largely accelerate downstream applications. The training visualization is presented in the supplemental video.

Tetrahedra grids. We start with a tetrahedra grid with 128×128 resolution, including 192k tetrahedra and 37k vertices. Each tetrahedron can produce at most 2 triangles by marching tetrahedra algorithm [61, 79, 19]. To increase the resolution of the tetrahedra grid, we subdivide the grid at the 500th step. To avoid the out-of-memory problem caused by the vast amount of floating meshes in the void space at the beginning of training, we pre-train the SDF network to **match a visual hull** of humans in canonical space. The hull could be derived from either the skeleton capsules or the SMPL [53] mesh. Note that we only pre-train for 500 iterations, which leads to **a very coarse shape** akin to the visual hull rather than the given ground truth mesh. The initialized mesh is presented in the training visualization part of the supplemental video.

9. Additional Ablations

Number of Training Views. Table 4 and Figure 9 show that giving one camera of view degrades the overall reconstruction quality, and multi-view consistency improves the final results. The model can aggregate multi-view information for better shading optimization, thus leading to clearer surface materials.

The effect of Skinning Module Design Table 5-6 and Figure 10-12 reveal that the initialization with pre-trained skinning net and the regularization on surface skinning improve the overall reconstruction quality. The initialization provides skinning prior which helps to speed up geometry convergence. From Figure 10-11, the geometry details improve with the initialization under the same training time.

The regularization on surface skinning prevents geometry degradation. Figure 12 indicates that our model can not learn correct canonical geometry without the initialization

Table 4: **Ablation results of training views on the ZJU-MoCap 313 subject**.

ZJU-MoCap 313	Training pose		Novel pose	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
1 view	24.39	0.913	21.45	0.869
2 views	28.06	0.945	22.81	0.888
3 views	28.50	0.956	23.17	0.894
4 views	**29.04**	**0.961**	**23.20**	**0.896**

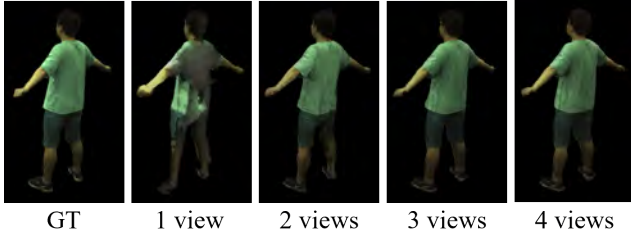
GT   1 view   2 views   3 views   4 views

Figure 9: **Ablation study of training views on the ZJU-MoCap 313 subject.**

and the regularization. The mesh distortion is reduced with the regularization.

Table 5: **The ablation on skinning module of H36M S9 dataset**.

H36M S9	Training Pose		Novel Pose	
	PSNR	SSIM	PSNR	SSIM
w/o skinning init. & reg.	24.88	0.905	21.97	0.851
w/ skinning intialization	26.28	0.926	24.47	0.897
w/ skinning regularization	26.24	0.925	24.34	0.896
Full	**26.29**	**0.926**	**24.53**	**0.899**

Table 6: **The ablation on skinning module of ZJU-MoCap 313 dataset**.

ZJU-MoCap 313	Training Pose		Novel Pose	
	PSNR	SSIM	PSNR	SSIM
w/o skinning init. & reg.	27.46	0.949	20.31	0.831
w/ skinning intialization	28.82	0.958	23.08	0.893
w/ skinning regularization	28.80	0.959	23.14	0.895
Full	**29.05**	**0.961**	**23.27**	**0.897**

Effect of SDF Network The MLP parametrization of the SDF field keeps our surface both water-tight and smooth, as shown in Figure 13.
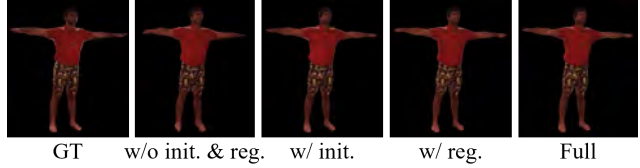

GT   w/o init. & reg.   w/ init.   w/ reg.   Full

Figure 10: **Ablation study of the skinning module on the H36M S9 subject.**

GT   w/o init. & reg.   w/ init.   w/ reg.   Full

Figure 11: **Ablation study of the skinning module on the ZJU-MoCap 313 subject.**

w/o init. & reg.   w/ init.   w/ reg.   Full

Figure 12: **Ablation study of the skinning module on the ZJU-MoCap 313 subject.**
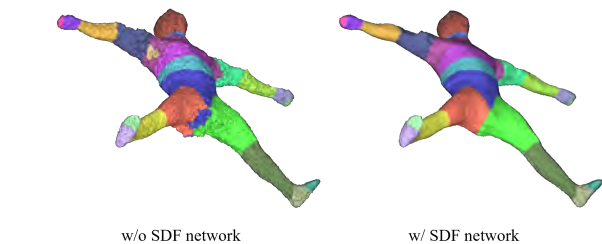

w/o SDF network   w/ SDF network

Figure 13: **Ablation study of SDF field parametrization.**

10. More Comparisons

We present full quantitative comparisons in Table 9, Table 7, Table 10, and Table 8. Meanwhile, more qualitative comparisons are illustrated in Figure 14, Figure 15, and Figure 16.
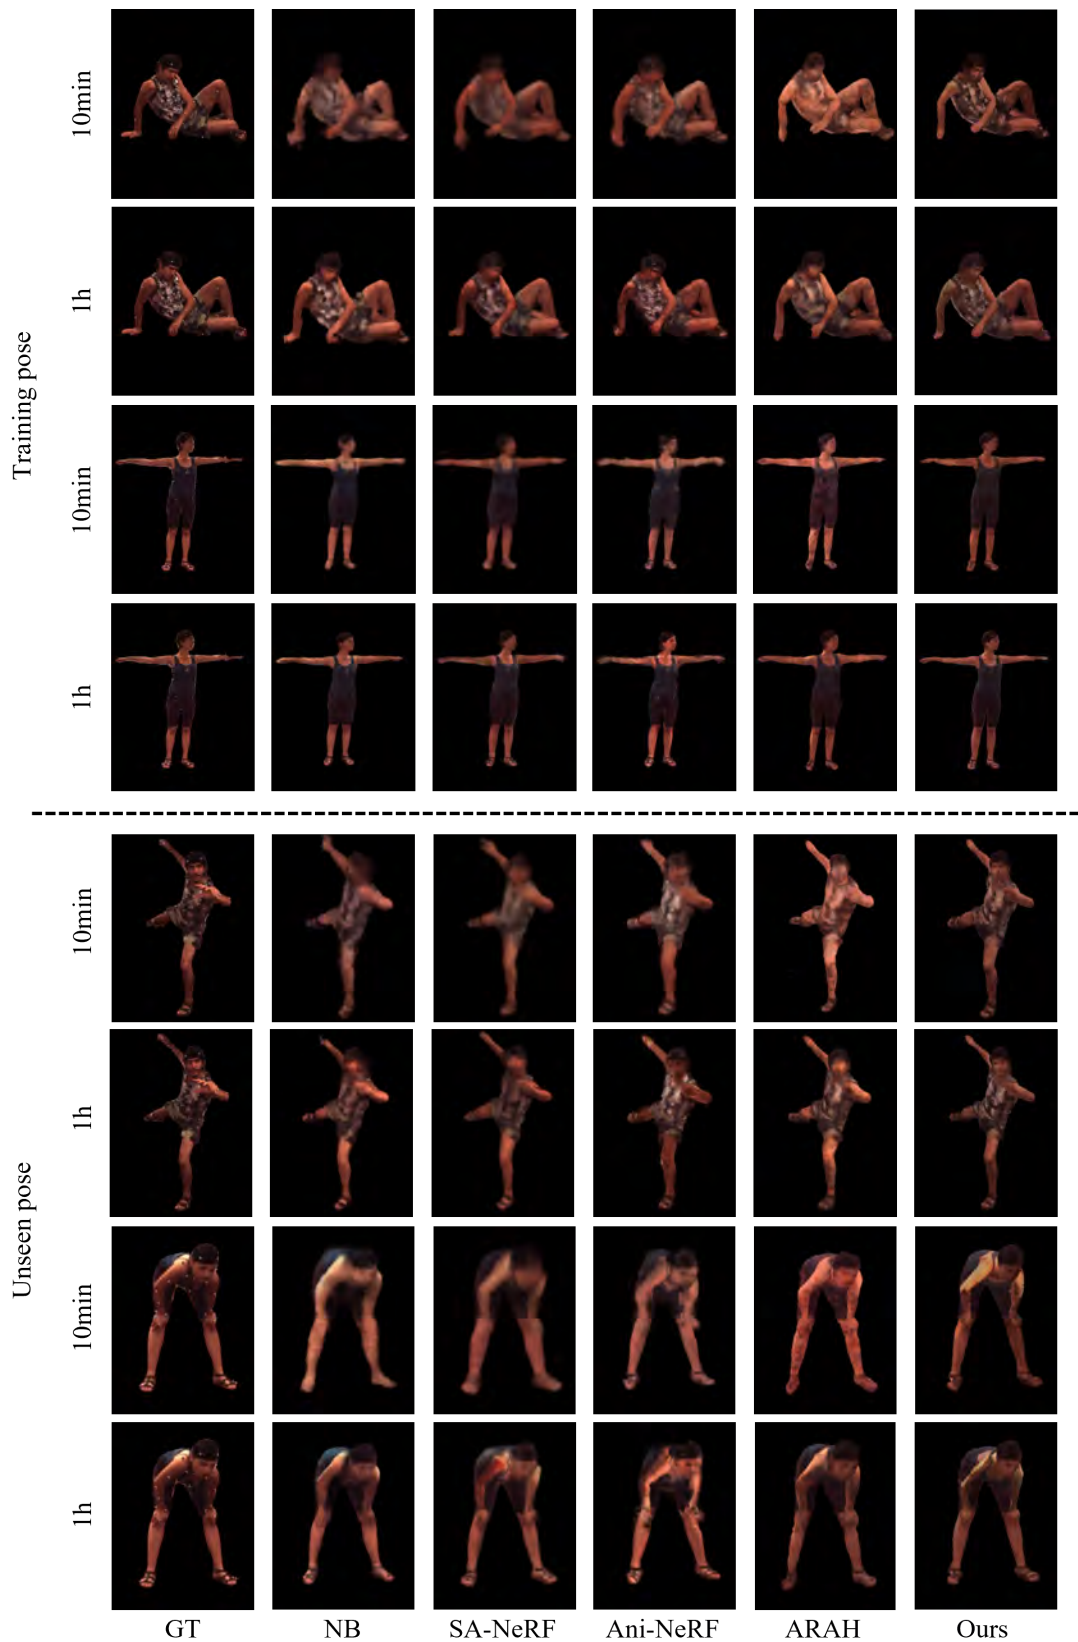
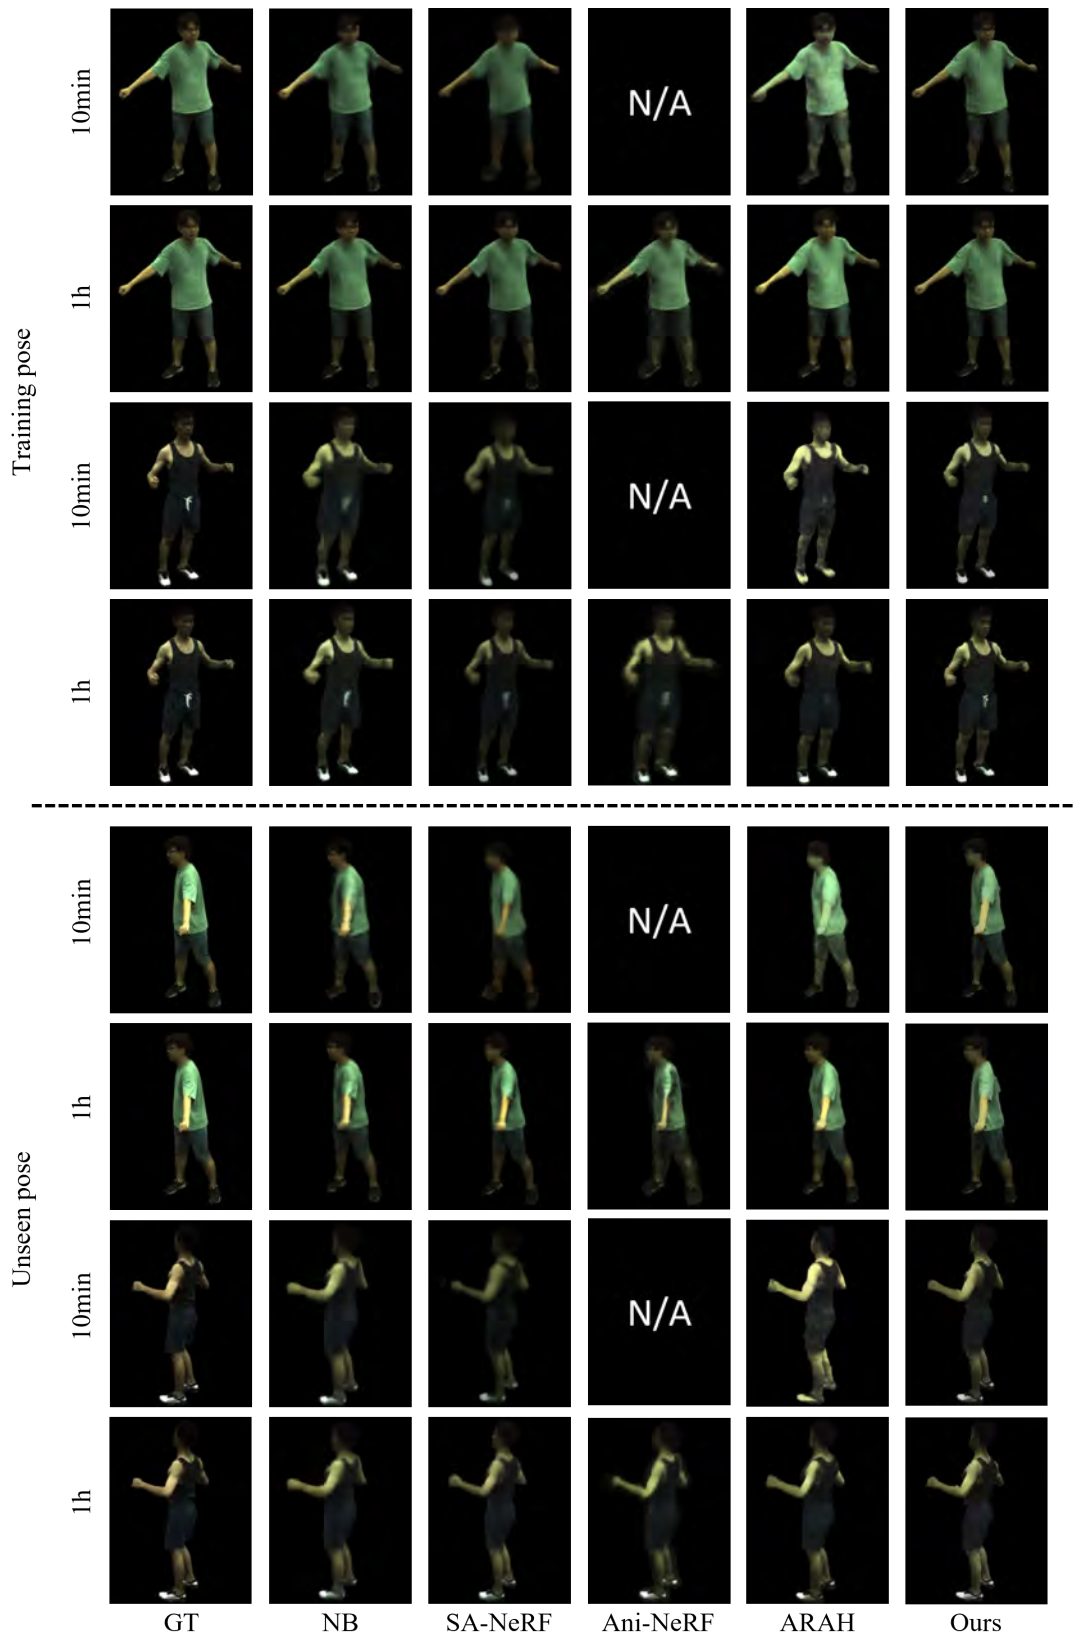Figure 14: **Qualitative results of novel pose synthesis on H36M dataset.** Zoom in for a better view.

Figure 15: **Qualitative results of novel pose synthesis on ZJU-MoCap dataset.** "N/A" denotes nothing to render due to no convergence. Zoom in for a better view.
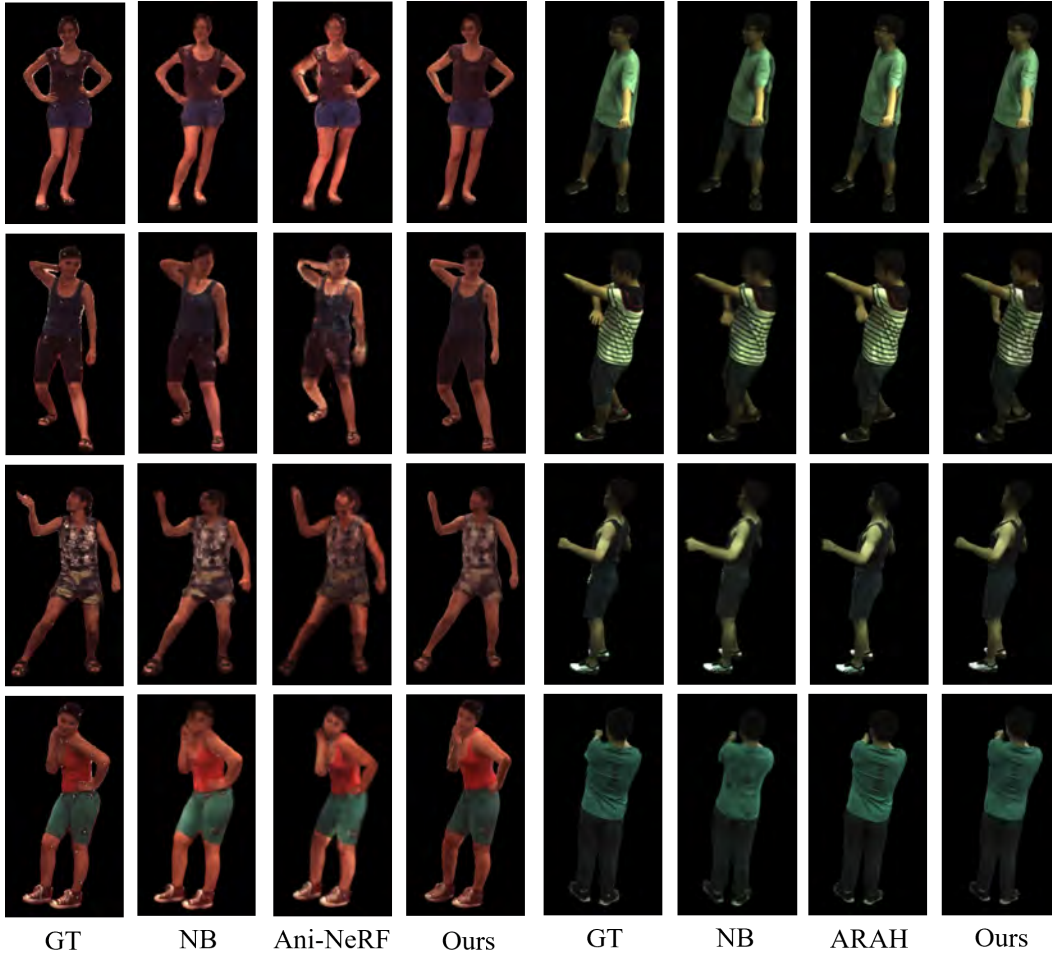
Figure 16: **Qualitative results of novel pose synthesis on H36M and ZJU-MoCap datasets with the full models.** Zoom in for a better view.

Table 7: **Quantitative results of training pose novel view synthesis of H36M dataset**.

| | Training pose | | | | | | | | | |
| | PSNR | | | | | SSIM | | | | |
	NB	SA-NeRF	Ani-NeRF	ARAH	Ours	NB	SA-NeRF	Ani-NeRF	ARAH	Ours
S1	22.87	23.71	22.05	24.45	24.56	0.897	0.915	0.888	0.919	0.919
S5	24.60	24.78	23.27	24.54	24.51	0.917	0.909	0.892	0.918	0.920
S6	22.82	23.22	21.13	24.61	24.55	0.888	0.881	0.854	0.903	0.902
S7	23.17	22.59	22.50	24.31	24.05	0.914	0.905	0.890	0.919	0.916
S8	21.72	24.55	22.75	24.02	23.94	0.894	0.922	0.898	0.921	0.920
S9	24.28	25.31	24.72	26.20	25.99	0.910	0.913	0.908	0.924	0.919
S11	23.70	25.83	24.55	25.43	25.48	0.896	0.917	0.902	0.921	0.915
Average	23.31	24.28	23.00	24.79	24.72	0.902	0.909	0.890	0.918	0.916

11. Applications

We showcase **relighting**, **texture editing**, and **novel poses synthesis** on AIST dataset [46] in Figure 17, Figure 18, and Figure 19 separately. All the above applications are presented in the supplemental video.

Table 8: **Quantitative results of unseen pose novel view synthesis of H36M dataset**.

| | Unseen pose | | | | | | | | | |
| | PSNR | | | | | SSIM | | | | |
	NB	SA-NeRF	Ani-NeRF	ARAH	Ours	NB	SA-NeRF	Ani-NeRF	ARAH	Ours
S1	21.93	22.67	19.96	23.08	23.72	0.873	0.890	0.855	0.899	0.904
S5	23.33	23.27	20.02	22.79	23.13	0.893	0.881	0.840	0.890	0.898
S6	23.26	23.23	23.64	24.04	24.17	0.888	0.888	0.882	0.900	0.903
S7	22.40	22.51	21.76	22.58	22.72	0.888	0.898	0.869	0.891	0.889
S8	20.78	23.06	21.63	22.34	22.71	0.872	0.904	0.877	0.896	0.902
S9	22.87	23.84	21.95	24.36	24.54	0.880	0.889	0.871	0.894	0.895
S11	23.54	24.19	22.55	24.78	24.47	0.879	0.891	0.875	0.902	0.900
Average	22.59	23.25	21.64	23.42	23.64	0.882	0.892	0.867	0.896	0.899

Table 9: **Quantitative results of training pose novel view synthesis of ZJU-MoCap dataset**.

| | Training pose | | | | | | | | | |
| | PSNR | | | | | SSIM | | | | |
	NB	SA-NeRF	Ani-NeRF	ARAH	Ours	NB	SA-NeRF	Ani-NeRF	ARAH	Ours
Twirl(313)	30.56	31.32	29.80	31.60	29.67	0.971	0.974	0.963	0.973	0.947
Taichi(315)	27.24	27.25	23.10	27.00	24.21	0.962	0.962	0.917	0.965	0.919
Swing1(392)	29.44	29.29	28.00	29.50	27.58	0.946	0.946	0.931	0.948	0.899
Swing2(393)	28.44	28.76	26.10	27.70	25.91	0.940	0.941	0.916	0.940	0.890
Swing3(394)	27.58	27.50	27.50	28.90	27.67	0.939	0.938	0.924	0.945	0.902
Warmup(377)	27.64	27.67	24.20	27.80	26.69	0.951	0.954	0.925	0.956	0.926
Punch1(386)	28.60	28.81	25.60	29.20	27.65	0.931	0.931	0.878	0.934	0.881
Punch2(387)	25.79	26.08	25.40	27.00	25.68	0.928	0.929	0.926	0.945	0.908
Kick(390)	27.59	27.77	26.00	27.90	24.08	0.926	0.927	0.912	0.929	0.840
Average	28.10	26.19	28.27	28.51	26.57	0.944	0.945	0.921	0.948	0.901

12. Mesh Visualizations

We visualize the canonical mesh and present the number of faces of each mesh in Figure 20 and Figure 12. Note that the number of faces for each mesh is quite small. Though increasing the resolution of tetrahedra grids may improve the details of both geometry and materials, we do not conduct this experiment for it is orthogonal to our technical contributions.

13. Limitations and Further Discussions

Our method is biased for shape-material ambiguity [88, 45, 81, 63, 69, 87, 98]. Taking subject 315 from ZJU-MoCap as an example, the strips in the T-shirt are modeled as ravines on the surface. The high contrast color in the cloth surface makes our model biased for shape modeling. That might be resolved by introducing additional surface regularizers or pre-defined parameters for the materials.

Need foreground mask to enable the mesh optimization, akin to shape-from-silhouette. One future direction might be equipping our method with the ability to separate foreground and background automatically [31, 24]. It is also promising to model the background simultaneously during foreground subject optimization [31, 24], which eliminates the requirement of foreground mask processing.

Our method can digitize humans from visual footage, which may involve avatar misuse without the permission of the owners. Methods like implicit adversarial watermarks [11, 47] that disable the neural nets inference could assist the video creation to protect their portrait rights. Another concern is the deep fake misuse [62], which corrupts the identity in the visual footage rendered by our model. Methods like deep fake detection [64] could help to discover and prevent deep fake creations. Besides, our method involves training with GPUs, which leads to carbon emissions and increasing global warming [67].

Table 10: **Quantitative results of unseen pose novel view synthesis of ZJU-MoCap dataset**.

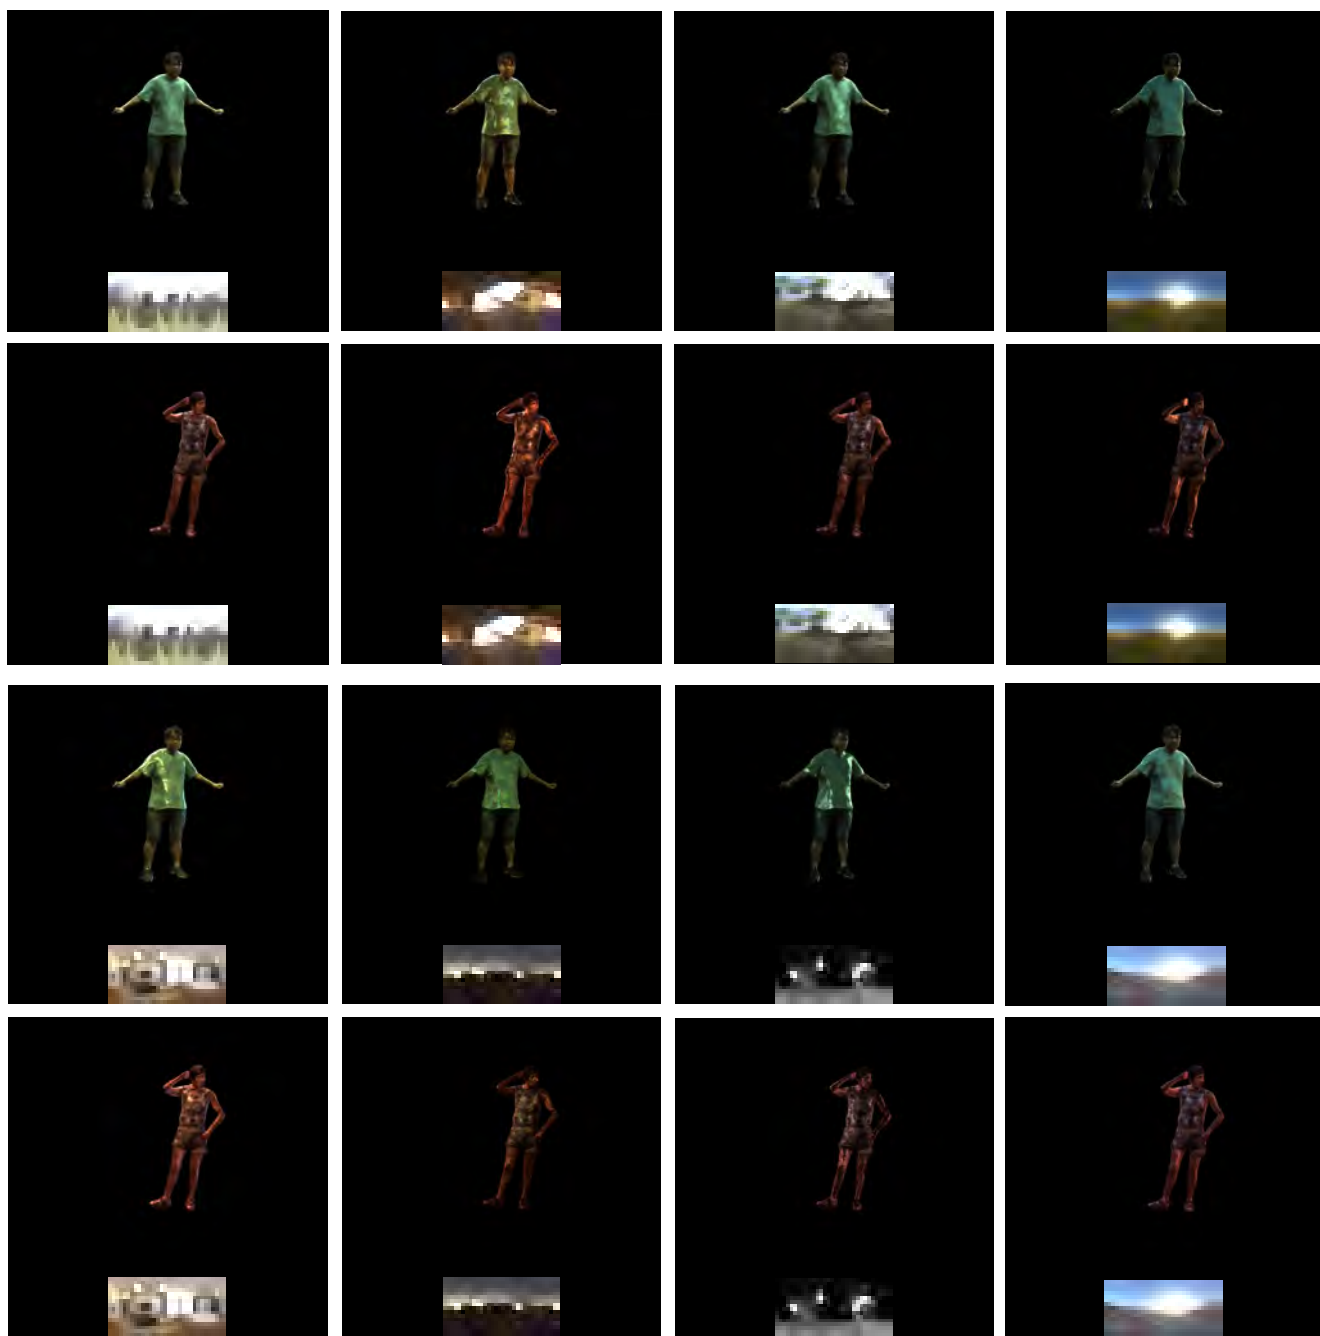| | Unseen pose | | | | | | | | | |
| | PSNR | | | | | SSIM | | | | |
	NB	SA-NeRF	Ani-NeRF	ARAH	Ours	NB	SA-NeRF	Ani-NeRF	ARAH	Ours
Twirl(313)	23.95	24.33	22.80	24.40	23.63	0.905	0.908	0.863	0.914	0.878
Taichi(315)	19.56	19.87	18.47	20.00	20.42	0.852	0.863	0.795	0.881	0.850
Swing1(392)	25.76	26.27	18.44	26.20	25.49	0.909	0.927	0.670	0.927	0.883
Swing2(393)	23.80	24.96	21.87	24.40	24.31	0.878	0.900	0.836	0.915	0.883
Swing3(394)	23.25	24.24	17.69	25.20	24.72	0.893	0.908	0.792	0.908	0.870
Warmup(377)	23.91	25.34	23.28	25.50	24.80	0.909	0.928	0.901	0.933	0.894
Punch1(386)	25.68	27.30	25.55	27.00	26.24	0.881	0.905	0.872	0.910	0.853
Punch2(387)	21.60	23.08	21.92	24.20	24.06	0.870	0.890	0.838	0.917	0.889
Kick(390)	23.90	24.43	23.90	24.80	25.79	0.870	0.889	0.887	0.896	0.873
Average	23.49	24.42	21.55	24.63	24.38	0.885	0.902	0.828	0.911	0.875

Figure 17: **Relighting visualization.** Zoom in for a better view. We strongly encourage our readers to view the supplemental video for a more comprehensive visual perception.
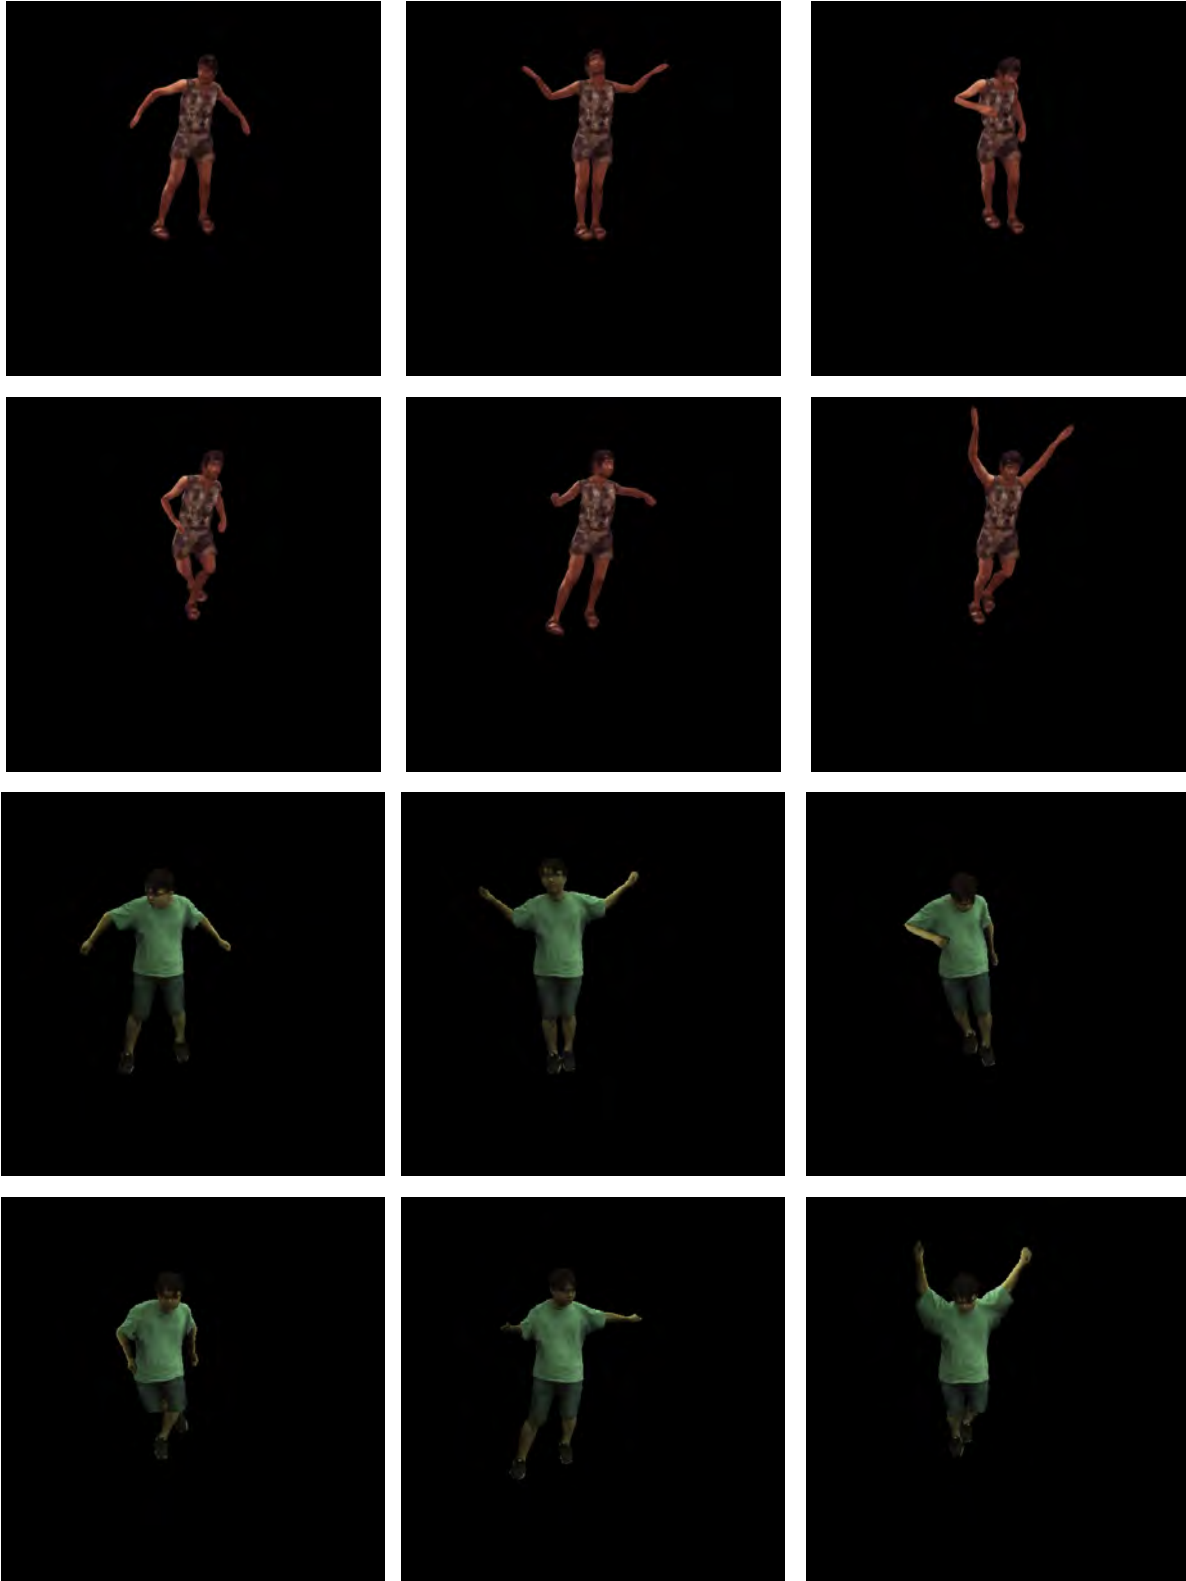
Figure 18: **Texture editing visualization.** Zoom in for a better view. We strongly encourage our readers to view the supplemental video for a more comprehensive visual perception.
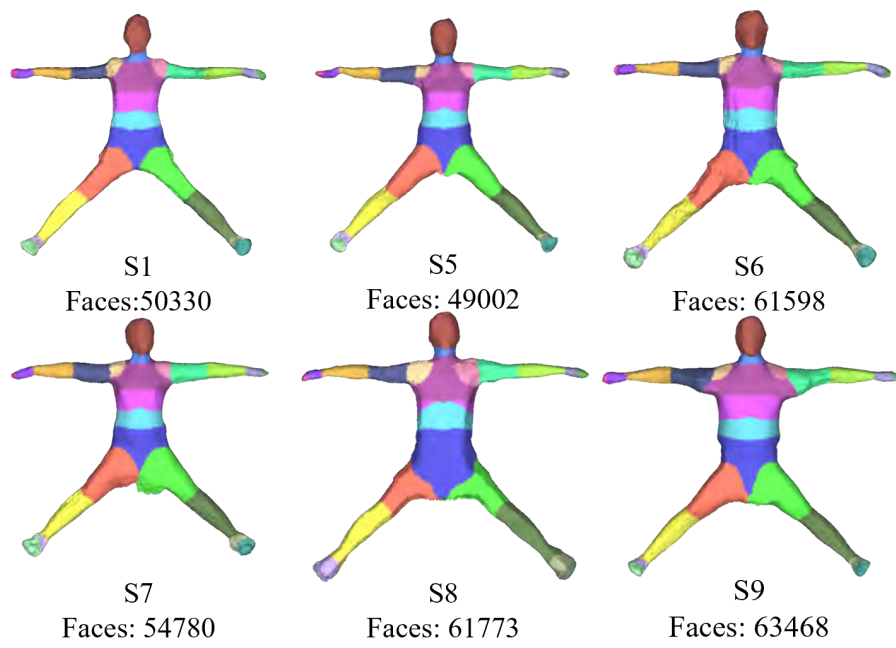
Figure 19: **Extreme pose visualization.** Zoom in for a better view. We strongly encourage our readers to view the supplemental video for a more comprehensive visual perception.

Figure 20: **Mesh visualization on the H36M dataset.** Zoom in for a better view.

313
Faces: 66250

315
Faces: 59426

377
Faces: 57514

386
Faces: 74474

387
Faces: 70176

390
Faces: 69646

392
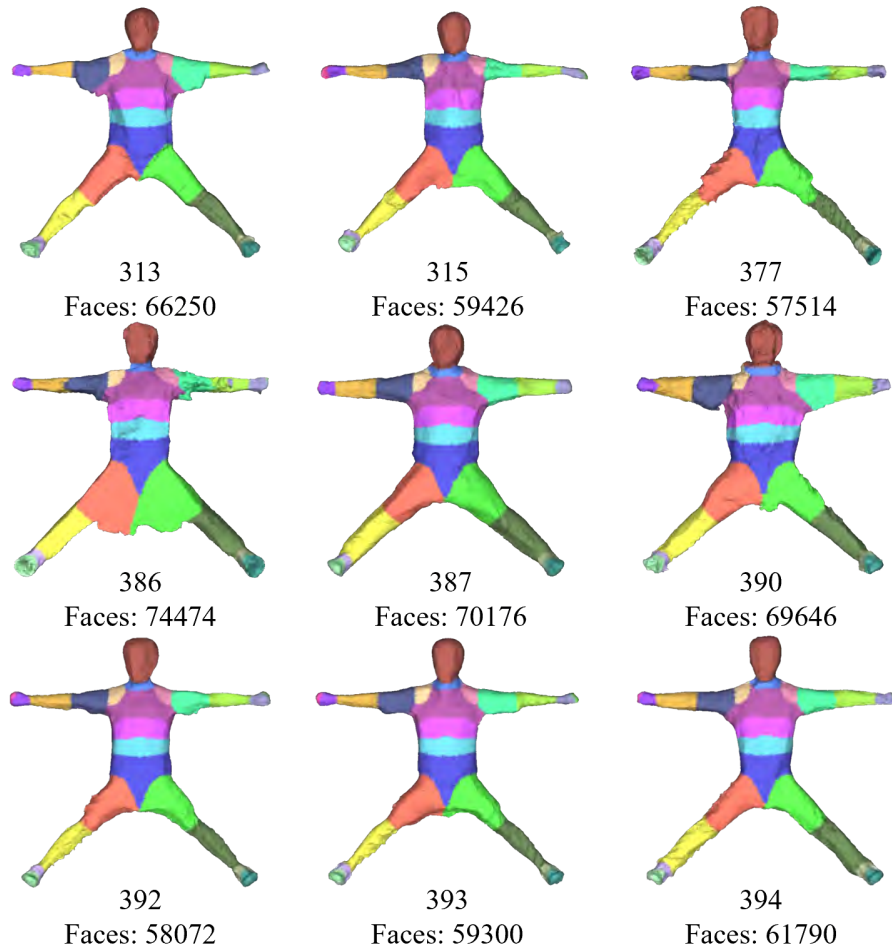Faces: 58072

393
Faces: 59300

394
Faces: 61790

Figure 21: **Mesh visualization on the ZJU-MoCap dataset.** Zoom in for a better view.