

NeuralHOFusion: Neural Volumetric Rendering under Human-object Interactions

Yuheng Jiang¹ Suyi Jiang¹ Guoxing Sun¹ Zhuo Su² Kaiwen Guo³
 Minye Wu⁴ Jingyi Yu^{1,5} Lan Xu^{1,5}

¹ShanghaiTech University ²Tencent ³Meta Research Lab ⁴KU leuven
⁵Shanghai Engineering Research Center of Intelligent Vision and Imaging

Abstract

4D modeling of human-object interactions is critical for numerous applications. However, efficient volumetric capture and rendering of complex interaction scenarios, especially from sparse inputs, remain challenging. In this paper, we propose NeuralHOFusion, a neural approach for volumetric human-object capture and rendering using sparse consumer RGBD sensors. It marries traditional non-rigid fusion with recent neural implicit modeling and blending advances, where the captured humans and objects are layer-wise disentangled. For geometry modeling, we propose a neural implicit inference scheme with non-rigid key-volume fusion, as well as a template-aid robust object tracking pipeline. Our scheme enables detailed and complete geometry generation under complex interactions and occlusions. Moreover, we introduce a layer-wise human-object texture rendering scheme, which combines volumetric and image-based rendering in both spatial and temporal domains to obtain photo-realistic results. Extensive experiments demonstrate the effectiveness and efficiency of our approach in synthesizing photo-realistic free-view results under complex human-object interactions.

1. Introduction

Human-centric 4D content generation enables numerous applications for VR/AR, telepresence and education. However, conveniently reconstructing and rendering human activities under human-object interactions remain unsolved.

Early high-end solutions [5, 7, 12, 13, 20, 29] require dense cameras and custom-designed lighting conditions for high-fidelity reconstruction. But such a complicated and expensive system setup is undesirable for consumer-level usage. Light-weight volumetric performance capture is more practical and attractive. Early solutions [14, 21, 22, 51] rely on pre-scanned templates which are unsuitable for on-the-fly human-object interaction modeling. The volumetric fu-

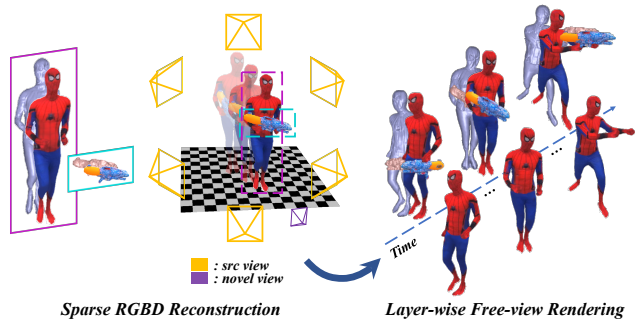


Figure 1. Our NeuralHOFusion achieves layer-wise and photo-realistic reconstruction results, using only 6 RGBD cameras.

sion approaches Fusion4D [9] and Motion2Fusion [8] further reconstruct complex human-object interaction scenes with topology changes in real-time. But they heavily rely on high-quality depth sensors and up to 9 high-end GPUs, which are infeasible for consumer usage. Besides, the low-end fusion approaches [34, 43, 44, 57, 63] adopt the most handy monocular setup with a temporal fusion pipeline [35], but suffer from the inherent self-occlusion constraint. Moreover, the appearance results of the fusion methods are restricted by the limited geometry resolution.

Recent learning-based techniques enable robust human modeling from only light-weight inputs. In particular, various approaches [40, 41, 48] utilize implicit function to model human geometry, which is also widely adopted in the volumetric capture pipeline [25, 26, 44, 62]. But these methods are restricted to only human without modeling human-object interactions, let alone generating compelling photo-realistic texture. Similarly, despite the progress for realistic human rendering [30, 32, 33, 48, 56], few researchers explore the neural rendering strategies for human-object interactions, especially under the volumetric capture framework. On the other hand, various researchers [16, 17, 38, 50, 64–67] model the interactions between humans and the surrounding objects or environments. But they only recover the parametric human model rather than reconstructing and rendering the interaction scenes. Only recently, a few methods [45, 47] explicitly model human and object simultaneously in the

volumetric capture framework. But they still cannot handle the interaction scenes, which highly limit the practicality.

In this paper, we present *NeuralHOFusion* – a neural volumetric human-object capture and rendering system using light-weight consumer RGBD sensors (see Fig. 1 for overview). In stark contrast with existing systems, our approach handles various complex human-object interaction scenarios and even multi-person interactions. It achieves photo-realistic layer-wise geometry and texture rendering in novel views for both the performers and interacted objects.

Generating such a human-object free-viewpoint video with the layer-wise visual effect whilst maintaining light-weight and efficient setting is non-trivial. Our key idea is to organically combine traditional volumetric non-rigid fusion pipeline with recent neural implicit modeling and blending advances, besides embracing a layer-wise scene decoupling strategy. To this end, we first utilize off-the-shelf instance segmentation approach to distinguish the human and object from the six RGBD streams. For human reconstruction, we propose a fusion-based neural implicit scheme to reason about the human-only geometry details in novel views. Specifically, it combines pixel-aligned features with an occlusion-aware truncated projective SDF (TSDF) feature [62], by utilizing a traditional key-volume non-rigid fusion pipeline [9, 57] in a human-only manner. Such a key-volume fusion-based implicit scheme handles occlusions effectively. For object reconstruction, inspired by the recent work [45], we adopt a template-aid robust object tracking pipeline with a specific initialization process for the following neural blending. Finally, based on the human-object geometry proxy above, we propose a layer-wise neural blending scheme to disentangle human and object for photo-realistic performance rendering. For the human phase, we combine the image-based rendering with the traditional per-vertex texturing using albedo volume [59], through occlusion-aware blending weight learning. It enables accurate human appearance rendering in the target view with the level of texture detail in the spatially adjacent input. For the object rendering, we extend the spatial neural blending into the temporal domain, which learns the blending weight from both the spatial and temporal candidate input views for photo-realistic rendering. To summarize, our main contributions include:

- We present the first neural volumetric capture and rendering system for human-object interaction scenarios using light-weight consumer RGBD sensors.
- We propose a fusion-based neural implicit inference scheme for detail-preserved human-object reconstruction in an occlusion-aware manner.
- We introduce a layer-wise neural rendering scheme, which combines volumetric and image-based rendering in both spatial and temporal domains.

2. Related Work

Human-Object Capture. Markerless human-object performance capture techniques have been widely investigated to achieve free-viewpoint video or immersive telepresence. Early high-end works [7, 13] use dense cameras for reconstruction and rendering of human and objects through mesh reconstruction and motion tracking, but it is expensive to build the synchronized and calibrated multi-camera systems. The recent low-end approaches enable light-weight performance capture under the single-RGB setup [23, 60, 64], single-RGBD setup [6, 45] or sparse RGBs setup [47, 48]. In another line, [16, 17, 38, 50, 65–67] model the interaction between humans and the objects or the surrounding environments. PHOSA [64] runs human-object capture without any 3D supervision, considering the relationship between human and objects to eliminate ambiguity. But they only recover the naked human template and produce a visually reasonable spatial arrangement. RobustFusion [45] captures human and objects by volumetric fusion, as well as tracks object by Iterative Closest Point. But they cannot handle topology changes and their texture suffers from blur artifacts. HOI-FVV [47] utilizes a decoupling strategy to process human and object respectively under six RGB cameras. Though they show impressive rendering results for human-object interactions, they only process simple pose of human and the inference speed is very slow. Comparably, our approach achieves high-fidelity capture and rendering for various human-object interactions with complex human pose and severe occlusions at fast speed.

Human Volumetric Capture. Volumetric fusion based methods [34, 58, 61, 63, 69] allow free-form dynamic reconstruction in a template-free, single-view, real-time way, through updating depth into the canonical model and performing non-rigid deformation. A series of works are proposed to make volumetric fusion more robust with SIFT features [18], human articulated skeleton prior [61, 63], extra IMU sensors [69], data-driven prior [44], learned correspondences [4] or neural deformation graph [3]. Since these single-view setups suffer from tracking error in the occluded parts, multi-view setups are introduced to mitigate this problem with improved fusion methods. Fusion4D [9] proposes a key volume updating strategy. Motion2fusion [8] incorporates learning-based surface matching into pipeline. UnstructuredFusion [59] achieves an unstructured multi-view setup. Function4D [62] combines temporal volumetric fusion and implicit functions to generate complete geometry. However, these methods either cannot handle modeling human-object interactions or generate photo-realistic rendering results. Comparably, our approach realizes the abilities of high-fidelity capture and rendering of human-object interactions.

Neural Rendering and Blending. In the area of photo-realistic novel view synthesis and 3D scene reconstruc-

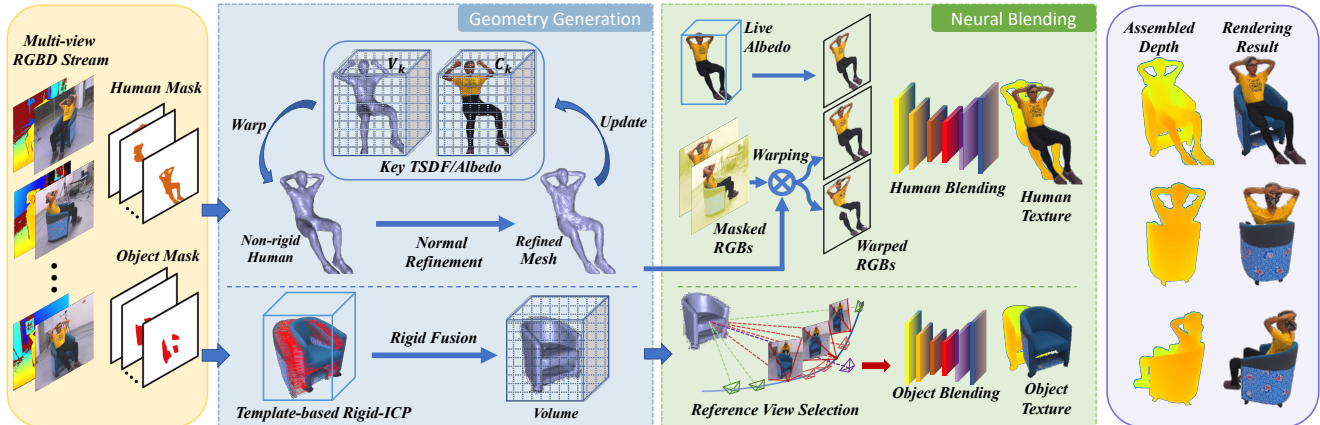


Figure 2. Our approach consists of two stages. The geometry module includes neural human reconstruction (Sec. 4.1) and template-aided object fusion (Sec. 4.2), and the blending module includes neural human blending (Sec. 4.3) and temporal neural object blending (Sec. 4.4).

tion, neural rendering shows great power and huge potential. Various data representations are adopted to obtain better performance and characteristics, such as point-clouds [1, 49, 56], voxels [30], texture meshes [28, 52] or implicit functions [33, 36] and hybrid neural blending [47, 48]. NHR [56] embeds spatial features into sparse dynamic point-clouds, Neural Volumes [30] transforms input images into a 3D volume representation by a VAE network. More recently, [24, 37, 39, 53, 68] extend neural radiance field [33] into the dynamic setting. However, for all approaches above, dense spatial views or full temporal frames are required in training for high fidelity novel view rendering. Blending based methods learn blending weight for adjacent views and synthesize photo-realistic novel views in a light-weight way. [48] uses the occlusion map as guidance for blending weight estimation. [47] incorporates the direction information to reduce artifacts in wide baseline. However, they cannot handle the occlusion region. Comparably, our blending with spatial-temporal information, enables recovery of photo-realistic texture of human and objects even under the extreme occlusion region.

3. Overview

Given human-object interaction videos under sparse RGBDs setting, NeuralHOFusion can reconstruct high-quality geometries and synthesize layer-wise photo-realistic free-viewpoint videos even in challenging scenarios with extreme poses, occlusions. As illustrated in Fig. 2, NeuralHOFusion includes two streams for human and objects separately, and each stream includes two steps: Geometry Generation and Neural Blending.

Geometry Generation. To achieve high-quality human-object geometry for neural rendering under sparse RGBD cameras setting, NeuralHOFusion incorporates global temporal information into key volumes. For humans, we dynamically maintain a key TSDF volume V_k and its albedo

volume C_k . NeuralHOFusion generates a complete geometry of the non-rigid human via a proposed pixel-aligned approach that accompanies either fused TSDF volume V_t of current frame or V_k to assist global reconstruction. Besides, normal refinement helps to restore more geometry details. We then utilize complete geometry for accurate neural texture blending and key TSDF volume V_k updating. For object, we adopt rigid tracking and volumetric fusion to reconstruct the geometry with the aid of the template generated by an occupancy regression network.

Neural Blending. To produce photo-realistic textures based on the above geometries, we propose neural blending schemes to extract features from input textures and predict their blending weights. For human, the blending network takes a projected image from albedo volume and adjacent warped RGB images as input, and then predict the blending weight to blend the final texture. As for object, we reserve non-occluded spatial-temporal observations into an observation-angle group. We then retrieve “adjacent views” from this group to perform temporal blending. After assembling blended human and object textures, NeuralHOFusion outputs final rendering results.

4. Method

4.1. Neural Human Reconstruction

To reconstruct complete and fine-detailed human geometry, we sequentially perform fusion-based implicit reconstruction and key volume update.

Fusion-based Implicit Reconstruction. Neural networks based on implicit functions are good at complete reconstruction but lack geometry details and temporal consistency, while traditional volumetric capture methods [9, 34, 59] have enabled the temporal-consistent reconstruction results. From the insight of this complementarity, we follow Function4d [62] to combine the non-rigid fusion with implicit functions by extracting features from TSDF volume.

However, when facing severe occlusions in human-object interactions, [62] will fail since their non-rigid fusion in a sliding way could not provide complete TSDF features. Therefore, we non-rigidly fuse a key volume V_k as a reference model, in which we track the motion field from the live frame to it and integrate each depth into this TSDF volume. Note that the motion field is represented by embedded deformation graph(ED-graph) [46] and SMPL model [31], please refer to [44, 59, 63] for details about this non-rigid fusion process. Then we combine the key volume V_k with current volume V_t which is extracted from current depth and RGBD images to our implicit reconstruction network f for inferring detailed and complete geometry, as shown in Fig. 3. We follow the network architecture of PIFu [40] to regress an implicit function f to predict occupancy of every 3D point X in the space and formulate it as:

$$f(\phi(X), \alpha(X), z(X)) = s : s \in [0.0, 1.0],$$

$$\phi(X) = \frac{1}{n} \sum_i^n F_i(\pi_i(X)), \quad (1)$$

where $\pi_i(\cdot)$ denotes projection matrix of i -th camera; $z(\cdot)$ is the depth value of X . $F_i(\pi_i(X)) = g(I_i(\pi_i(X)))$ is the image feature of X on RGBD images, $g(\cdot)$ is feature extraction network, I_i is input image. $\alpha(\cdot)$ represents the queried TSDF value from V_k or V_t of the 3D point X , in which V_t gives the detailed geometry information of current frame, while V_k reserves global information in the occluded region. To make the most of them, we introduce a dynamic selection strategy: when the point X is close to the visible human body, we choose V_t 's value, $\alpha(X) = V_t(X)$, otherwise, we choose the value in warped V_k . This operation ensures our method to reconstruct high fidelity geometry in the visible area and reasonable geometry in the occluded part.

Key volume update. Our V_k provides a global prior to support the network to infer a complete and temporal-consistent human reconstructions in occlusion scenarios. In practice, we reset V_k periodically to reduce misalignment caused by its difference to live frame and enable to handle the topology-changing issues. More specifically, we not only fuse the TSDF volume V_k with current depth and the final geometry output through the estimated motion fields [63] for each frame, we also reset key volume and update ED-graph through re-sampling nodes on the output mesh and SMPL model to re-initialize the motion field at a fixed frequency(40 frames in the paper).

4.2. Template-aid Object Fusion.

Although naive volumetric fusion [45] provides a generally correct object geometry, deteriorated surface caused by the partial depth loss and limited overlap between the matched point-clouds may also occur, as shown in Fig. 7 (b). To achieve stable and accurate object reconstruction,




Figure 3. Illustration of human geometry reconstruction.(Sec. 4.1)

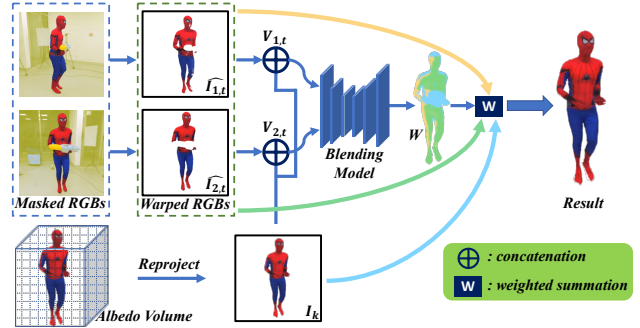


Figure 4. Network structure of neural human blending.(Sec. 4.3)

we also incorporate an object template generation module. Specifically, this module provides a global cue to regulate object tracking, which greatly improves the robustness of tracking, thus leading to more accurate geometry results.

Firstly, we utilize a data-driven occupancy regression network to generate a complete object template from multi-view RGBD images, in which the formulation resembles Eqn. 1 without TSDF features. We then optimize the rigid motions T of the corresponding object point-clouds under the ICP framework as:

$$\begin{aligned} E_{\text{object}}(T) = & \lambda_{\text{geo}} \sum_{(\mathbf{p}, \mathbf{q}) \in \mathcal{R}_{\text{geo}}} (\mathbf{n}_{\mathbf{p}}^T (\mathbf{p} - T\mathbf{q}))^2 + \\ & \lambda_{\text{tem}} \sum_{(\mathbf{p}_t, \mathbf{q}) \in \mathcal{R}_{\text{tem}}} (\mathbf{n}_{\mathbf{p}_t}^T (\mathbf{p}_t - T\mathbf{q}))^2 + \\ & \lambda_{\text{sp.o}} E_{\text{sp.o}}, \end{aligned} \quad (2)$$

where \mathcal{R}_{geo} is the correspondence pair sets between source point-clouds and fused point-clouds, \mathcal{R}_{tem} is the correspondence pair sets between source point-clouds and template point-clouds. \mathbf{q} is the point from source, \mathbf{p} is the point from fused model and \mathbf{p}_t from template. $E_{\text{sp.o}}$ is a term to punish mesh interpenetration, please refer to [45] for more details. Finally, with the estimated T , source point-clouds are fused into a TSDF volume to update the object geometry.

4.3. Neural Human Blending

To enable fast novel view synthesis, we adopt a neural blending pipeline to generate photo-realistic and non-

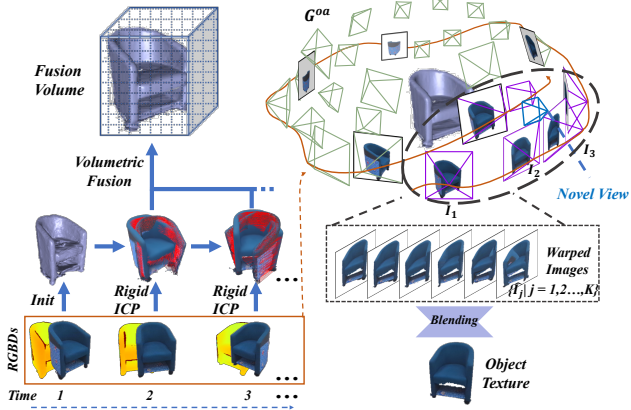


Figure 5. Illustration of our object capture and rendering. The left side is the object rigid tracking and fusion(Sec. 4.2). The right side is the temporal neural object blending(Sec.4.4).

occluded human textures under novel view, which incorporates information from the albedo volume maintained in the key frame [15], adjacent input views and local fine-detailed geometry, as illustrated in Fig. 4.

We expand the work [48] to define the variance and occlusion maps to learn our blending network. Specifically, we obtain the albedo image and depth map in target view (I_k and D_k) using the key albedo volume and occupancy field respectively in 4.1. The albedo image I_k lacks details but reserves complete texture. It has rich information for a blending pipeline in the occluded part through providing color candidate and help finding occluded part. We then warp adjacent RGBD maps into target view denoted by $\hat{I}_{1,t}, \hat{I}_{2,t}, \hat{D}_{1,t}, \hat{D}_{2,t}$. Subsequently, we calculate the occlusion maps as $O_i = D_k - \hat{D}_{i,t}, i = 1, 2$, and compute a per-element variance maps $V_{i,t} = (\hat{I}_{i,t} - I_k)^2, i = 1, 2$. Our blending network Θ_{HBN} utilizes the color information from albedo image and adjacent images, the variance information $V_{i,t}$ and occlusion information O_i to predict pixel-wise blending map W , which can be formulated as:

$$W = \Theta_{HBN}(I_{1,t}, V_{1,t}, O_1, I_{2,t}, V_{2,t}, O_2). \quad (3)$$

Our novel view result of human I_n can be formulated as:

$$I_n = \hat{W}_1 \cdot \hat{I}_{1,t} + \hat{W}_2 \cdot \hat{I}_{2,t} + \hat{W}_3 \cdot I_k, \quad (4)$$

where \hat{W}_i denotes blending weights with the sum of 1.0. **Normal Refinement.** To further improve the quality of geometry in novel view, we follow [48] to perform a normal refinement to infer the displacement of the target depth by a normal refinement network Θ_{HRN} . We apply [55] to source view RGBD images $I_{i,t}, I_k$ with novel depth D_k , and I_n with novel depth D_k respectively, we get normal map of source $N_{i,t}$, albedo novel view N_k , and novel view N_n . We blend adjacent $N_{i,t}$ and N_k with \hat{W}_i to get the blended normal map N_b . Θ_{HRN} takes N_n, N_b and D_k as input, and finally predicts the depth displacement.

4.4. Temporal Neural Object Blending.

With the temporal observations from object fusion and subsequent tracking, we introduce a strategy to combine these observations into our object blending pipeline smoothly and effectively. In the initialization stage of object capture, we intentionally show non-occluded images to the cameras and the system works on-line with rigid tracking to collect a group G^{oa} of observation-angle pairs. In the tracking stage, non-occluded observation-angle pairs are also added into G^{oa} . For novel view generation, we interpolate ‘‘nearby views’’. More specifically, we follow [54] to identify a pool of 18 ‘‘nearby views’’ from the pool. In this manner, our blending can collect more information in a wider baseline. We then introduce a temporal neural object blending pipeline to predict novel view object textures as illustrated in Fig. 5. This blending network can be formulated as:

$$W = \Theta_{TBN}(\{\hat{I}_j, O_j | j = 1, 2, \dots, 6\}). \quad (5)$$

\hat{I}_j is the warped image of nearby source view I_j from G^{oa} . O_j denotes the occlusion map. Novel view images will be generated like Eqn. 4. Although the albedo volume of object can also be used for object blending, we find our blending strategy with G^{oa} observations is sufficient to get good and complete rendering results.

4.5. Implementation Details

For human-object segmentation, we first use [27] for background separation then train [2] to get the initial object coarse mask. Subsequently, We follow [45] to refine the mask. For image encoders g , we follow [48] to use a U-Net, which outputs 64 channels feature maps. For implicit decoders f , we use MLP with skip connections as [62], in which the hidden neurons are (128,128,128,128,128). Furthermore, the loss function of MLP f minimizes the average of mean squared error like [40]. $\Theta_{HBN}, \Theta_{HRN}$ and Θ_{TBN} adopt the U-Net structure. For geometry training, we firstly collect 100 human sequences and 40 objects in a dome and utilize existing object meshes from 3D-FUTURE [11]. Then, we place a human in the center of our camera setting and add objects with predefined trajectories to simulate human-object interactions. We then render RGBD images under our camera parameters and add synthesized noises according to [10]. For training blending networks, we render the RGBD images, normals and masks under 180 novel views as ground truth. Besides, both L1 loss and perceptual loss [19] are used. For template-aid object fusion, we use the following empirically determined parameters: $\lambda_{geo} = 0.8, \lambda_{tem} = 0.2, \lambda_{sp-o} = 1.0$.



Figure 6. The geometry and texture results of our NeuralHOFusion on various interaction sequences, including “floral dress”, “nesting in a sofa”, “undressing coat” and “shaking hands”.

5. Experimental Results

In this section, we evaluate our NeuralHOFusion in various extremely challenging human-object interactions. All the experiments are run on a PC with an Nvidia GeForce RTX3090 GPU and an Intel i7-8700k CPU. Capturing six RGBD streams from a synchronized-Kinect system, our method produces high-quality geometry and texture results. In order to achieve fast performance, we follow [62] and [48] to implement our entire pipeline on GPU. For each part, the human geometry generation takes 129 ms and 7 GB, the following neural human blending pipeline costs 20 ms and 2.3 GB. The object geometry initialization takes 10 s where the robust object rigid fusion takes around 33 ms. Finally, the neural object blending costs 42 ms. The whole object branch takes around 7 GB of memory consumption. Various geometry and texture results of our NeuralHOFusion are shown in Fig. 6, including different type interactions, and even severe occlusion and topology changes, such as nesting in a sofa, shaking hands, and removing clothes.

5.1. Comparison

We compare our NeuralHOFusion against the start-of-art methods UnstructuredFusion [59], RobustFusion [45] and

NeuralHumanFVV [48] both in geometry and texture. As illustrated in Fig. 7, UnstructuredFusion [59] fails to handle the human-object interactions, RobustFusion [45] does not support topology changes, and NeuralHumanFVV [48] cannot reconstruct correct geometry facing different object types. While our NeuralHOFusion achieves more detailed, complete, and isolated geometry results and significantly more photo-realistic rendering results, even under the challenging interactions and extreme human pose. Please note that our approach can also enable layer-wise rendering which is not supported by UnstructuredFusion and NeuralHumanFVV. The quantitative results in Tab. 1 and Tab. 2 also demonstrate that our approach can achieve consistent better results on all the metrics.

5.2. Ablation Study

Neural Human Geometry Generation. As shown in Fig. 8 (b), without the TSDF feature, the generation model is unable to recover the occluded human part, and lacks of details, while after encoding the TSDF feature in Fig. 8 (c), the model can generate a complete human body with mid-level geometry details such as the clothing wrinkles but still suffers from over-smooth results, especially on the facial




Figure 7. Qualitative comparison. (a) Input images. (b-e) are the geometry and texture results in a novel view from UnstructuredFusion [59], RobustFusion [45], NeuralHumanFVV [48] and ours, respectively.

Table 1. Quantitative comparison of rendering results.

Method	PSNR \uparrow	SSIM \uparrow	MAE \downarrow
UnstructuredFusion [59]	22.456	0.937	3.058
RobustFusion [45]	26.537	0.941	1.868
NeuralhumanFVV [48]	27.526	0.979	1.131
Ours	33.59	0.984	0.627

Table 2. Quantitative comparison of geometry reconstructions.

Method	P2S $\times 10^{-4}\downarrow$	Chamfer $\times 10^{-4}\downarrow$
Multi-PIFu [40]	14.475	10.564
RobustFusion [44]	5.770	6.375
Ours	2.692	2.853

area. In contrast, our full approach with normal refinement achieves detailed human geometry reconstruction as shown in Fig. 8 (d). For further quantitative analysis, we evaluate each component using CD and P2S as shown in Tab. 3, which highlights the contributions of our each component.

Template-aid Object Fusion. As shown in Fig. 9 (b), without the template produced by MLP, simple rigid ICP is prone to slight mismatches, which affects the generation of fused mesh. Furthermore, due to the limitation of volume resolution, template-aid rigid ICP still suffers from the over-smooth issues as shown in Fig. 9 (c). In contrast, our full pipeline with refinement can recover more detailed object geometry. Further quantitative analysis in Tab. 3 demonstrates that our method achieves higher accuracy.

Neural Human Blending. In Fig. 10, we evaluate different variants of texturing schemes with the same geometry proxy. The texture extracted from albedo volume as shown in Fig. 10 (b) is blurred, while the naive neural blending results in Fig. 10 (c) suffer from severe block artifacts, which blends the object texture to the human. In contrast, our full neural human blending scheme achieves both photo-realistic and complete texture results as shown in Fig. 10 (d). Besides, we also make a comparison on a synthetic sequence with 400 frames and generate 180 different target views to evaluate. The Tab. 4 demonstrates that our method

Table 3. Quantitative evaluation of reconstruction schemes.

Method	P2S $\times 10^{-4}\downarrow$	Chamfer $\times 10^{-4}\downarrow$
w/o TSDF (human)	7.7407	7.969
w/o normal refinement (human)	3.1425	3.4086
neural human geometry generation	2.855	3.239
w/o template (object)	35.138	19.383
w/o normal refinement (object)	11.529	9.256
template-aid object fusion	11.480	9.166

Table 4. Quantitative evaluation of texturing schemes.

Method	PSNR \uparrow	SSIM \uparrow	MAE \downarrow
albedo volume (human)	26.758	0.925	2.167
naive neural blending (human)	25.983	0.962	1.735
neural human blending	30.040	0.968	0.945
albedo volume (object)	33.455	0.950	0.564
naive neural blending (object)	30.760	0.968	1.255
temporal neural object blending	37.901	0.971	0.376

achieves higher accuracy.

Temporal Neural Object Blending. As for evaluation of object texturing, Fig. 11 (b) demonstrates that texture fusion scheme leads to blur, Fig. 11 (c) shows that the naive neural texturing blending scheme wrongly recovers the texture which belongs to the human part. In contrast, our temporal neural object blending makes full use of both the previous non-occluded frames and the current frame. Therefore, we can faithfully recover the accurate texture even when some parts are severe occluded in Fig. 11 (d). The quantitative experiments on synthetic object sequences can refer to Tab. 4.

Camera Number. We evaluate the influence of input view number in our multi-view setting, where the cameras are placed around a circle uniformly and numbered from 0 to 5. We compare the results of two-camera system (0, 3), four-camera system (0, 1, 3, 4) and six-camera system. As shown in Fig. 12, without sufficient camera views, the reconstructed geometry is a little downgrading, while the textured results significantly get worse. Empirically, we find the system with six-camera produces good results in a compromise of camera number and quality.

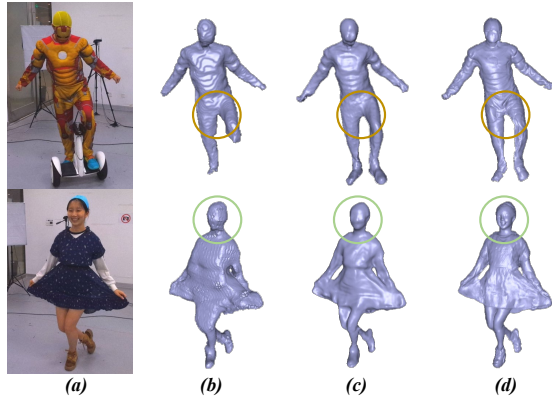


Figure 8. Qualitative evaluation of human geometry generation. (a) Input images. (b) Geometry without TSMDF feature; (c) Geometry without normal refinement; (d) Ours.

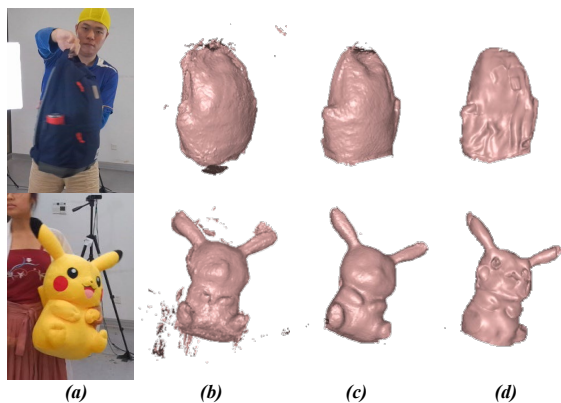


Figure 9. Qualitative evaluation of object geometry generation. (a) Input images. (b) Geometry without template; (c) Geometry without normal refinement; (d) Ours.

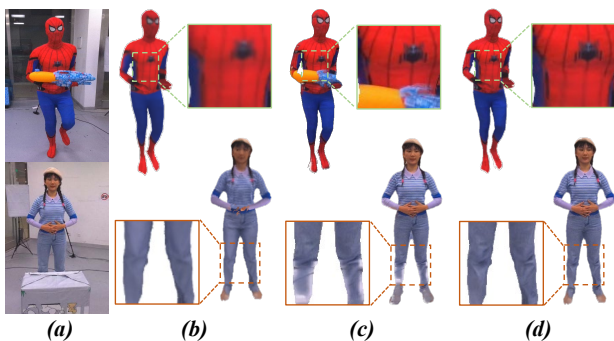


Figure 10. Qualitative evaluation of neural blending. (a) Input image; (b) Per-vertex texture; (c) Naive neural blending; (d) Ours.

5.3. Limitation

Although NeuralHOFusion can perform detailed and complete reconstruction and layer-wise photo-realistic rendering under complex human-object interactions by fully utilizing temporal observations, it still owns some limitations. First, our method heavily relies on instance segmentation method, bad segmentation will lead to awful geometry and interlaced texture. Besides, since depth sensors lack observations from specific materials such as shaggy hair and




Figure 11. Qualitative evaluation of neural blending scheme on objects. (a) Input images; (b) Per-vertex texture; (c) Naive neural texture blending; (d) Ours.




Figure 12. Evaluation of the number input camera views. (a) Reference image. (b, c, d) Our reconstructed texture results using two, four and six cameras, respectively. (e) Cumulative distribution function of the mean absolute error.

yarn clothes, our method cannot get good geometry to these areas. A semantic-aware implicit function on different human parts will be critical for such problem. Furthermore, our approach will produce texture-copy artifacts after normal refinement. Our current pipeline models human and objects separately, and it is an interesting direction to build a physical framework such as [42].

6. Conclusion

We have presented a practical neural volumetric capture and rendering approach for complex human-object interaction scenes, using sparse RGBD cameras. By combining traditional non-rigid fusion with neural implicit modeling and blending, our system achieves detailed and realistic results with the unique layer-wise viewing experience. Our fusion-based neural implicit inference and template-aid object tracking enable detailed and complete geometry generation under occlusions, while our texturing scheme combines volumetric and image-based rendering in both spatial and temporal domains to synthesize photo-realistic texture. Our experimental results demonstrate the effectiveness of NeuralHOFusion in complex interaction scenarios with various poses and clothing types. We believe that our approach is a critical step to virtually but realistic teleport human performances under complex interactions, with many potential applications like consumer-level telepresence, active object scanning and human behavior analysis.

Acknowledgements. This work was supported by Shanghai YangFan Program (21YF1429500), Shanghai Local college capacity building program (22010502800).

References

- [1] Kara-Ali Aliev, Artem Sevastopolsky, Maria Kolos, Dmitry Ulyanov, and Victor Lempitsky. Neural point-based graphics. *arXiv preprint arXiv:1906.08240*, 2019. 3
- [2] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact++: Better real-time instance segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 5
- [3] Aljaz Bozic, Pablo Palafox, Michael Zollhofer, Justus Thies, Angela Dai, and Matthias Nießner. Neural deformation graphs for globally-consistent non-rigid reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1450–1459, 2021. 2
- [4] Aljaz Bozic, Michael Zollhofer, Christian Theobalt, and Matthias Nießner. Deepdeform: Learning non-rigid rgb-d reconstruction with semi-supervised data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7002–7012, 2020. 2
- [5] Derek Bradley, Tiberiu Popa, Alla Sheffer, Wolfgang Heidrich, and Tamy Boubekeur. Markerless garment capture. In *ACM SIGGRAPH 2008 papers*, pages 1–9. 2008. 1
- [6] Andrei Burov, Matthias Nießner, and Justus Thies. Dynamic surface function networks for clothed human bodies. *arXiv preprint arXiv:2104.03978*, 2021. 2
- [7] Alvaro Collet, Ming Chuang, Pat Sweeney, Don Gillett, Dennis Evseev, David Calabrese, Hugues Hoppe, Adam Kirk, and Steve Sullivan. High-quality streamable free-viewpoint video. *ACM Transactions on Graphics (TOG)*, 34(4):69, 2015. 1, 2
- [8] Mingsong Dou, Philip Davidson, Sean Ryan Fanello, Sameh Khamis, Adarsh Kowdle, Christoph Rhemann, Vladimir Tankovich, and Shahram Izadi. Motion2fusion: Real-time volumetric performance capture. *ACM Transactions on Graphics (TOG)*, 36(6):1–16, 2017. 1, 2
- [9] Mingsong Dou, Sameh Khamis, Yury Degtyarev, Philip Davidson, Sean Ryan Fanello, Adarsh Kowdle, Sergio Orts Escolano, Christoph Rhemann, David Kim, Jonathan Taylor, et al. Fusion4d: Real-time performance capture of challenging scenes. *ACM Transactions on Graphics (ToG)*, 35(4):1–13, 2016. 1, 2, 3
- [10] Péter Fankhauser, Michael Bloesch, Diego Rodriguez, Ralf Kaestner, Marco Hutter, and Roland Siegwart. Kinect v2 for mobile robot navigation: Evaluation and modeling. In *2015 International Conference on Advanced Robotics (ICAR)*, pages 388–394. IEEE, 2015. 5
- [11] Huan Fu, Rongfei Jia, Lin Gao, Mingming Gong, Binqiang Zhao, Steve Maybank, and Dacheng Tao. 3d-future: 3d furniture shape with texture. *arXiv preprint arXiv:2009.09633*, 2020. 5
- [12] Juergen Gall, Carsten Stoll, Edilson De Aguiar, Christian Theobalt, Bodo Rosenhahn, and Hans-Peter Seidel. Motion capture using joint skeleton tracking and surface estimation. In *CVPR*, 2009. 1
- [13] Kaiwen Guo, Peter Lincoln, Philip Davidson, Jay Busch, Xueming Yu, Matt Whalen, Geoff Harvey, Sergio Orts-Escolano, Rohit Pandey, Jason Dourgarian, et al. The relightables: Volumetric performance capture of humans with realistic relighting. *ACM Transactions on Graphics (TOG)*, 38(6):1–19, 2019. 1, 2
- [14] Kaiwen Guo, Feng Xu, Yangang Wang, Yebin Liu, and Qionghai Dai. Robust Non-Rigid Motion Tracking and Surface Reconstruction Using L0 Regularization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3083–3091, 2015. 1
- [15] Kaiwen Guo, Feng Xu, Tao Yu, Xiaoyang Liu, Qionghai Dai, and Yebin Liu. Real-time geometry, albedo and motion reconstruction using a single rgbd camera. *ACM Transactions on Graphics (TOG)*, 2017. 5
- [16] Vladimir Guzov, Aymen Mir, Torsten Sattler, and Gerard Pons-Moll. Human positioning system (hps): 3d human pose estimation and self-localization in large scenes from body-mounted sensors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, jun 2021. 1, 2
- [17] Mohamed Hassan, Partha Ghosh, Joachim Tesch, Dimitrios Tzionas, and Michael J. Black. Populating 3D scenes by learning human-scene interaction. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 1, 2
- [18] Matthias Innmann, Michael Zollhöfer, Matthias Nießner, Christian Theobalt, and Marc Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *European Conference on Computer Vision*, pages 362–379. Springer, 2016. 2
- [19] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016. 5
- [20] Hanbyul Joo, Tomas Simon, and Yaser Sheikh. Total capture: A 3d deformation model for tracking faces, hands, and bodies. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [21] Hao Li, Bart Adams, Leonidas J Guibas, and Mark Pauly. Robust single-view geometry and motion reconstruction. *ACM Transactions on Graphics (ToG)*, 28(5):1–10, 2009. 1
- [22] Hao Li, Linjie Luo, Daniel Vlasic, Pieter Peers, Jovan Popović, Mark Pauly, and Szymon Rusinkiewicz. Temporally coherent completion of dynamic shapes. *ACM Transactions on Graphics (TOG)*, 31(1):1–11, 2012. 1
- [23] Ruilong Li, Yuliang Xiu, Shunsuke Saito, Zeng Huang, Kyle Olszewski, and Hao Li. Monocular real-time volumetric performance capture. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 49–67, Cham, 2020. Springer International Publishing. 2
- [24] Zhengqi Li, Simon Niklaus, Noah Snavely, and Oliver Wang. Neural scene flow fields for space-time view synthesis of dynamic scenes. *arXiv preprint arXiv:2011.13084*, 2020. 3
- [25] Zhe Li, Tao Yu, Chuanyu Pan, Zerong Zheng, and Yebin Liu. Robust 3d self-portraits in seconds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1
- [26] Zhe Li, Tao Yu, Zerong Zheng, Kaiwen Guo, and Yebin Liu. Posefusion: Pose-guided selective fusion for single-view hu-

- man volumetric capture. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2021. 1
- [27] Shanchuan Lin, Andrey Ryabtsev, Soumyadip Sengupta, Brian Curless, Steve Seitz, and Ira Kemelmacher-Shlizerman. Real-time high-resolution background matting. *arXiv*, pages arXiv-2012, 2020. 5
- [28] Lingjie Liu, Weipeng Xu, Michael Zollhoefer, Hyeongwoo Kim, Florian Bernard, Marc Habermann, Wenping Wang, and Christian Theobalt. Neural rendering and reenactment of human actor videos. *ACM Transactions on Graphics (TOG)*, 38(5):1–14, 2019. 3
- [29] Yebin Liu, Juergen Gall, Carsten Stoll, Qionghai Dai, Hans-Peter Seidel, and Christian Theobalt. Markerless motion capture of multiple characters using multiview image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(11):2720–2735, 2013. 1
- [30] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *arXiv preprint arXiv:1906.07751*, 2019. 1, 3
- [31] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. In *SIGGRAPH Asia*, volume 34, pages 248:1–248:16, 2015. 4
- [32] Ricardo Martin-Brualla, Rohit Pandey, Shuoran Yang, Pavel Pidlypenskyi, Jonathan Taylor, Julien Valentin, Sameh Khamis, Philip Davidson, Anastasia Tkach, Peter Lincoln, and et al. Lookingood: Enhancing performance capture with real-time neural re-rendering. *ACM Trans. Graph.*, 37(6), Dec. 2018. 1
- [33] 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 Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 405–421, Cham, 2020. Springer International Publishing. 1, 3
- [34] Richard A Newcombe, Dieter Fox, and Steven M Seitz. Dynamicfusion: Reconstruction and tracking of non-rigid scenes in real-time. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 343–352, 2015. 1, 2, 3
- [35] Richard A. Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J. Davison, Pushmeet Kohli, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-Time Dense Surface Mapping and Tracking. In *Proc. of ISMAR*, pages 127–136, 2011. 1
- [36] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 165–174, 2019. 3
- [37] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo-Martin Brualla. Deformable neural radiance fields. *arXiv preprint arXiv:2011.12948*, 2020. 3
- [38] Priyanka Patel, Chun-Hao P. Huang, Joachim Tesch, David T. Hoffmann, Shashank Tripathi, and Michael J. Black. AGORA: Avatars in geography optimized for regression analysis. In *Proceedings IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, June 2021. 1, 2
- [39] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 3
- [40] 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 *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 1, 4, 5, 7
- [41] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1
- [42] Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Transactions on Graphics (TOG)*, 39(6):1–16, 2020. 8
- [43] Miroslava Slavcheva, Maximilian Baust, Daniel Cremers, and Slobodan Ilic. Killingfusion: Non-rigid 3d reconstruction without correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1386–1395, 2017. 1
- [44] Zhuo Su, Lan Xu, Zerong Zheng, Tao Yu, Yebin Liu, and Lu Fang. Robustfusion: Human volumetric capture with data-driven visual cues using a rgbd camera. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 246–264, Cham, 2020. Springer International Publishing. 1, 2, 4, 7
- [45] Zhuo Su, Lan Xu, Dawei Zhong, Zhong Li, Fan Deng, Shuxue Quan, and Lu Fang. Robustfusion: Robust volumetric performance reconstruction under human-object interactions from monocular rgbd stream. *arXiv preprint arXiv:2104.14837*, 2021. 1, 2, 4, 5, 6, 7
- [46] Robert W Sumner, Johannes Schmid, and Mark Pauly. Embedded deformation for shape manipulation. *ACM Transactions on Graphics (TOG)*, 26(3):80, 2007. 4
- [47] Guoxing Sun, Xin Chen, Yizhang Chen, Anqi Pang, Pei Lin, Yuheng Jiang, Lan Xu, Jingya Wang, and Jingyi Yu. Neural free-viewpoint performance rendering under complex human-object interactions. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 1, 2, 3
- [48] Xin Suo, Yuheng Jiang, Pei Lin, Yingliang Zhang, Minye Wu, Kaiwen Guo, and Lan Xu. Neuralhumanfvv: Real-time neural volumetric human performance rendering using rgb cameras. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6226–6237, 2021. 1, 2, 3, 5, 6, 7
- [49] Xin Suo, Minye Wu, Yanshun Zhang, Yingliang Zhang, Lan Xu, Qiang Hu, and Jingyi Yu. Neural3d: Light-weight neural portrait scanning via context-aware correspondence learning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3651–3660, 2020. 3

- [50] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *Computer Vision – ECCV 2020*, volume LNCS 12355, pages 581–600, Cham, Aug. 2020. Springer International Publishing. 1, 2
- [51] Jonathan Taylor, Jamie Shotton, Toby Sharp, and Andrew Fitzgibbon. The vitruvian manifold: Inferring dense correspondences for one-shot human pose estimation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 103–110. IEEE, 2012. 1
- [52] Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *ACM Transactions on Graphics (TOG)*, 38(4):1–12, 2019. 3
- [53] Liao Wang, Jiakai Zhang, Xinhang Liu, Fuqiang Zhao, Yanshun Zhang, Yingliang Zhang, Minye Wu, Lan Xu, and Jingyi Yu. Fourier plenotrees for dynamic radiance field rendering in real-time. *arXiv preprint arXiv:2202.08614*, 2022. 3
- [54] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 5
- [55] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 5
- [56] Minye Wu, Yuehao Wang, Qiang Hu, and Jingyi Yu. Multi-view neural human rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1, 3
- [57] Lan Xu, Wei Cheng, Kaiwen Guo, Lei Han, Yebin Liu, and Lu Fang. Flyfusion: Realtime dynamic scene reconstruction using a flying depth camera. *IEEE transactions on visualization and computer graphics*, 27(1):68–82, 2019. 1, 2
- [58] Lan Xu, Yebin Liu, Wei Cheng, Kaiwen Guo, Guyue Zhou, Qionghai Dai, and Lu Fang. Flycap: Markerless motion capture using multiple autonomous flying cameras. *IEEE Transactions on Visualization and Computer Graphics*, 24(8):2284–2297, Aug 2018. 2
- [59] Lan Xu, Zhuo Su, Lei Han, Tao Yu, Yebin Liu, and Lu Fang. Unstructuredfusion: realtime 4d geometry and texture reconstruction using commercial rgbd cameras. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2508–2522, 2019. 2, 3, 4, 6, 7
- [60] Weipeng Xu, Avishek Chatterjee, Michael Zollhöfer, Helge Rhodin, Dushyant Mehta, Hans-Peter Seidel, and Christian Theobalt. Monoperfcap: Human performance capture from monocular video. *ACM Transactions on Graphics (TOG)*, 37(2):27:1–27:15, 2018. 2
- [61] Tao Yu, Kaiwen Guo, Feng Xu, Yuan Dong, Zhaoqi Su, Jianhui Zhao, Jianguo Li, Qionghai Dai, and Yebin Liu. Bodyfusion: Real-time capture of human motion and surface geometry using a single depth camera. In *The IEEE International Conference on Computer Vision (ICCV)*. ACM, October 2017. 2
- [62] Tao Yu, Zerong Zheng, Kaiwen Guo, Pengpeng Liu, Qionghai Dai, and Yebin Liu. Function4d: Real-time human volumetric capture from very sparse consumer rgbd sensors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5746–5756, 2021. 1, 2, 3, 4, 5, 6
- [63] Tao Yu, Zerong Zheng, Kaiwen Guo, Jianhui Zhao, Qionghai Dai, Hao Li, Gerard Pons-Moll, and Yebin Liu. Doublefusion: Real-time capture of human performances with inner body shapes from a single depth sensor. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019. 1, 2, 4
- [64] Jason Y. Zhang, Sam Pepose, Hanbyul Joo, Deva Ramanan, Jitendra Malik, and Angjoo Kanazawa. Perceiving 3d human-object spatial arrangements from a single image in the wild. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2
- [65] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J. Black, and Siyu Tang. PLACE: Proximity learning of articulation and contact in 3D environments. In *International Conference on 3D Vision (3DV)*, Nov. 2020. 1, 2
- [66] Tianshu Zhang, Buzhen Huang, and Yangang Wang. Object-occluded human shape and pose estimation from a single color image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7376–7385, 2020. 1, 2
- [67] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J. Black, and Siyu Tang. Generating 3d people in scenes without people. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6194–6204, June 2020. 1, 2
- [68] Fuqiang Zhao, Wei Yang, Jiakai Zhang, Pei Lin, Yingliang Zhang, Jingyi Yu, and Lan Xu. Humannerf: Generalizable neural human radiance field from sparse inputs. *arXiv preprint arXiv:2112.02789*, 2021. 3
- [69] Zerong Zheng, Tao Yu, Hao Li, Kaiwen Guo, Qionghai Dai, Lu Fang, and Yebin Liu. Hybridfusion: Real-time performance capture using a single depth sensor and sparse imus. In *European Conference on Computer Vision (ECCV)*, Sept 2018. 2