

TeHOR: Text-Guided 3D Human and Object Reconstruction with Textures

Hyeongjin Nam¹ Daniel Sungho Jung² Kyoung Mu Lee^{1,2}

¹Dept. of ECE&ASRI, ²IPAI, Seoul National University

{namhjsnu28, dqj5182, kyoungmu}@snu.ac.kr

<https://hygenie1228.github.io/TeHOR/>

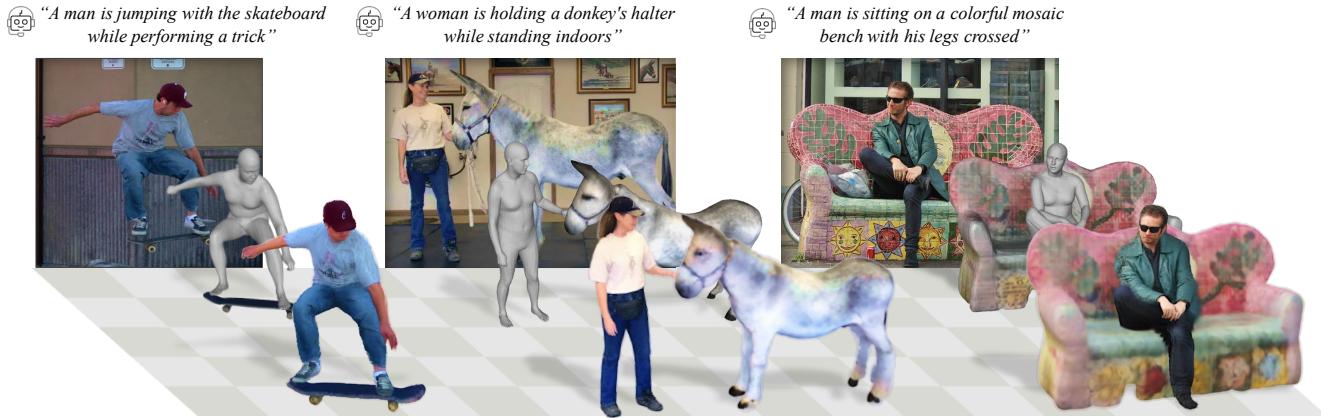


Figure 1. **TeHOR**. Given a single image, our framework jointly reconstructs textured 3D human and object by capturing their holistic and semantic interactions using text descriptions.

Abstract

*Joint reconstruction of 3D human and object from a single image is an active research area, with pivotal applications in robotics and digital content creation. Despite recent advances, existing approaches suffer from two fundamental limitations. First, their reconstructions rely heavily on physical contact information, which inherently cannot capture non-contact human–object interactions, such as gazing at or pointing toward an object. Second, the reconstruction process is primarily driven by local geometric proximity, neglecting the human and object appearances that provide global context crucial for understanding holistic interactions. To address these issues, we introduce **TeHOR**, a framework built upon two core designs. First, beyond contact information, our framework leverages text descriptions of human–object interactions to enforce semantic alignment between the 3D reconstruction and its textual cues, enabling reasoning over a wider spectrum of interactions, including non-contact cases. Second, we incorporate appearance cues of the 3D human and object into the alignment process to capture holistic contextual information, thereby*

ensuring visually plausible reconstructions. As a result, our framework produces accurate and semantically coherent reconstructions, achieving state-of-the-art performance.

1. Introduction

Joint reconstruction of 3D human and object from a single image is a cornerstone of human behavior understanding, as it provides insight into how humans interact physically and semantically with their surroundings, enabling broad applications in robotics, AR/VR, and digital content creation. To achieve physically and semantically accurate reconstruction, it is essential to faithfully capture human–object interaction (HOI), ranging from explicit physical contacts, such as grasping a cup, to implicit non-contact relations, such as reaching toward a door handle or gazing at a monitor.

Most existing methods for 3D human and object reconstruction from a single image [9, 13, 37, 50, 73, 76, 77, 81] have primarily utilized human–object contact information as a major cue for interaction reasoning. These approaches first predict contact regions on the 3D human and object surfaces, and subsequently enforce local geometric proximity

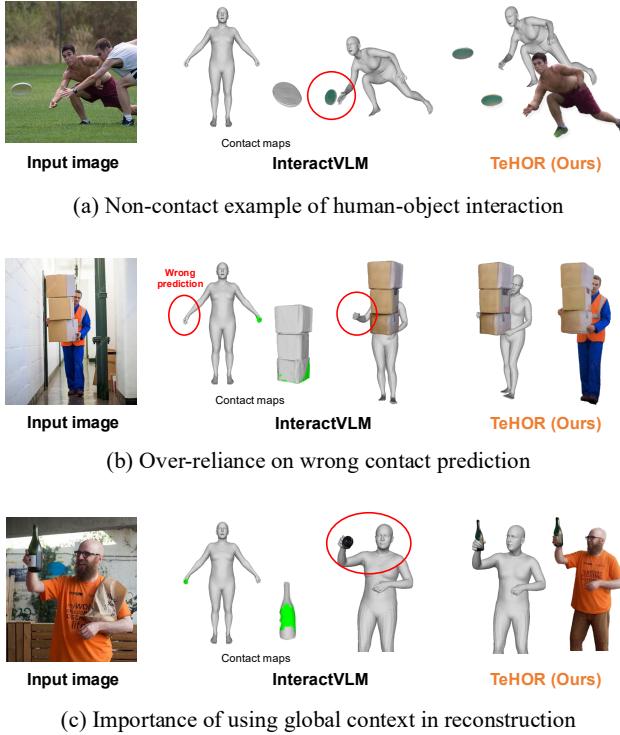


Figure 2. **Limitations of existing reconstruction methods.** Previous methods suffer from over-reliance on contact information and neglect the global interaction context, leading to implausible reconstructions. In contrast, TeHOR produces accurate and plausible 3D reconstructions by leveraging holistic and semantic guidance from text descriptions.

at the contact regions through iterative fitting. While these methods have shown notable progress, they suffer from two fundamental limitations. First, previous methods rely heavily on contact information, which provides limited physical cues and cannot capture non-contact interactions common in the real world, such as gazing at or pointing toward an object. As illustrated in Fig. 2 (a), when a person prepares to catch a frisbee, the absence of physical contact provides no cues for the reconstruction system to leverage, resulting in incorrect 3D reconstructions. Even in contact scenarios, over-reliance on contact information also makes these methods vulnerable to contact prediction errors, as shown in Fig. 2 (b), where inaccurate contact prediction directly leads to erroneous reconstruction results. Second, the fitting processes of the existing methods are primarily driven by local geometric proximity, ignoring global appearance cues of the human and object. The appearance cues, such as color and shading, provide rich contextual information for understanding the holistic structure of human–object interactions. However, the previous reconstruction systems primarily fit the geometric proximity between human and object without considering the holistic appearance, result-

ing in globally implausible results, such as an incorrectly oriented bottle and a misaligned human gaze in Fig. 2 (c).

To address these limitations, we propose **TeHOR** (Text-guided 3D Human and Object Reconstruction), which leverages rich text descriptions of human–object interaction as strong guidance for reconstruction. Given a single image, our framework first extracts text descriptions that specify the human–object interaction depicted in the image, using a vision–language model (*e.g.*, GPT-4 [1]). Based on the text descriptions, we jointly optimize the geometry and texture of the 3D human and object by supervising their rendered 2D appearances to align them with the semantic cues in the text. Specifically, this supervision is implemented by utilizing a pre-trained diffusion network (*e.g.*, StableDiffusion [62]) conditioned on the text descriptions. As the diffusion network encompasses strong prior knowledge of the association between visual appearance and text descriptions, it serves as a bridge between textual and visual domains. During the optimization, the diffusion network computes text-conditioned score gradients that drive the rendered appearances toward the visual distribution conditioned on the text, progressively refining the 3D human and object to better reflect the described interaction. This optimization is performed across multiple viewpoints, encouraging the reconstructed 3D human and object to exhibit consistent semantics across different viewpoints.

Built upon this design, our text-guided optimization offers two key advantages. First, as text descriptions contain semantic cues that extend beyond human–object contact, our framework enables reasoning about a wide range of interactions, including non-contact scenarios, such as approaching to catch a frisbee. Second, by supervising the 2D appearances of the human and object with a pre-trained diffusion network, our framework takes into account the holistic visual plausibility of the interaction, unlike previous methods that rely primarily on local geometric proximity. As a result, we show that our framework significantly improves the accuracy and plausibility of 3D human and object reconstruction and achieves state-of-the-art performance in both general and non-contact scenarios. Moreover, to the best of our knowledge, our framework is the first to jointly reconstruct full 3D textures of the human and the interacting object, which enables the creation of immersive and realistic 3D digital assets.

Our contributions can be summarized as follows.

- We propose **TeHOR**, which jointly reconstructs 3D human and object by leveraging text descriptions as semantic guidance, extending beyond physical contact cues.
- To capture the holistic context of human–object interaction, we incorporate appearance cues from the human and object by globally aligning their rendered appearances with the textual descriptions.
- Extensive experiments demonstrate that our proposed

framework significantly outperforms previous reconstruction methods across diverse interaction scenarios.

2. Related works

3D human and object reconstruction. Most of the recent works on 3D human and object reconstruction [7, 16, 37, 50, 76–78, 81, 85] primarily focus on capturing local interaction priors, particularly contact, to refine the spatial relationship between 3D human and object. PHOSA [85] optimizes the 3D spatial arrangement of humans and objects based on pre-defined contact labels, encouraging the contact regions on both human and object to ensure geometric proximity. CONTHO [50] is a framework that predicts contact maps directly from images and leverages a contact-based Transformer to enhance reconstruction using the predicted contact maps. More recently, several methods have been proposed for open-vocabulary reconstruction, enabling strong scalability and generalization to object categories unseen during training. InteractVLM [13] infers 2D contact maps in multiple views through a fine-tuned vision–language model [41] and uses the contact information to guide reconstruction. HOI-Gaussian [74] integrates a contact loss with an ordinal depth loss to more accurately capture relative depth between humans and objects.

Despite these advances, existing methods still rely heavily on physical cues, such as contact, causing them to struggle with complex interactions that require semantic reasoning. InteractVLM depends on accurate contact estimation, and failures in contact prediction directly degrade the quality of the final 3D reconstruction. While HOI-Gaussian introduces an ordinal depth constraint to mitigate contact failure, it remains vulnerable to complex interaction scenarios, such as severe human–object occlusion, as it relies on front-view cues without semantic reasoning. Unlike these methods, our framework is capable of holistic semantic reasoning about human–object interactions beyond contact information by leveraging global contextual information from text descriptions. Additionally, our framework goes beyond prior methods by reconstructing textured 3D humans and objects that provide richer representations than the non-textured surfaces, benefiting immersive downstream applications such as AR/VR.

Human-object interaction. With the emergence of large-scale 3D datasets [5, 25, 74], numerous studies [8, 17, 18, 24, 33, 66, 70, 82] have focused on capturing human–object interactions. Human–object contact has been a common cue in this line of research due to its intuitive physical representation observable in images. BSTRO [24] is a Transformer-based [11] framework that predicts dense body contact from a single image. DECO [70] utilizes a cross-attention network that jointly leverages human body parts and 2D scene context for contact estimation. LEMON [82] extends contact estimation approaches by modeling relations among

human contact, object affordance, and spatial configuration. More recently, ComA [32] introduces a probabilistic affordance representation that extends binary contact to incorporate relative orientation and proximity. Our framework goes beyond this line of work by leveraging textual descriptions that capture the semantic context of human–object interaction, enabling plausible 3D reconstructions.

3D human reconstruction. Most of the 3D human reconstruction methods are built upon parametric human models (*e.g.*, SMPL [46] and SMPL-X [56]) to estimate human body [3, 12, 14, 30, 36, 38, 49, 55, 72, 84] or clothed humans [2, 39, 47, 51–53, 60, 64, 67, 79, 87]. ARCH [19, 20] reconstructs animatable clothed humans using an implicit function that encodes occupancy, normals, and colors for detailed geometry and appearance. TeCH [26] employs a text-to-image diffusion model [63] to match the reconstructed appearance of the human with the textual description of the input image. LHM [59] represents human texture via 3D Gaussians in canonical space, enabling high-quality reconstruction and pose-controlled animation. In our framework, we utilize LHM for the initial 3D human reconstruction due to its strong generalization to in-the-wild scenarios.

3D object reconstruction. Recent works on 3D object reconstruction [15, 27, 42, 43, 45, 58, 68, 69, 75, 80, 86] have increasingly focused on open-vocabulary settings to generalize beyond fixed object categories. Zero-1-to-3 [43] introduces a viewpoint-conditioned image generative model that learns novel-view image synthesis, enabling zero-shot 3D reconstruction. DreamGaussian [69] optimizes a 3D Gaussian representation [31] for efficient and high-quality reconstruction. InstantMesh [80] is a feed-forward Transformer [71] that exploits a multi-view diffusion prior for efficient textured 3D object reconstruction. In all our experiments, we use the 3D object meshes produced by InstantMesh as the initial object reconstruction to ensure a fair and consistent baseline, given that most previous reconstruction methods are mesh-based pipelines.

3. TeHOR

Fig. 3 illustrates the overall pipeline of TeHOR. In the following sections, we first describe the 3D representation of the human and object (Sec. 3.1). Subsequently, we provide detailed descriptions of two stages: the reconstruction stage (Sec. 3.2) and the HOI optimization stage (Sec. 3.3).

3.1. 3D representation

We represent the 3D human and object, each as a set of 3D Gaussians, denoted by Φ_h and Φ_o , respectively. Following 3DGS [31], each 3D Gaussian is defined by its 3D position centroid and a set of Gaussian attributes (*e.g.*, scale, rotation, opacity, and appearance features) that jointly encode geometric and appearance properties. To project 3D Gaussians onto the 2D image space, we adopt the differentiable

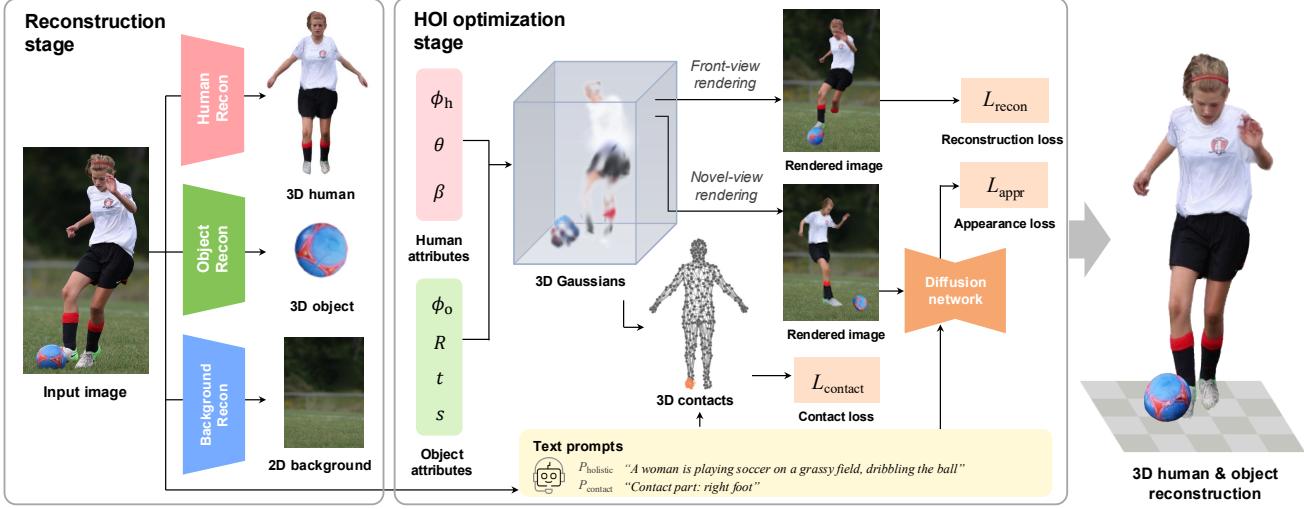


Figure 3. **Overall pipeline of TeHOR.** Given an input image, our framework initially reconstructs a 3D human, a 3D object, and a 2D background. Then, the initially reconstructed 3D human and object are jointly optimized using three core loss functions: reconstruction loss, appearance loss, and contact loss, to ensure accurate and semantically plausible human-object interaction.

rendering formulation introduced in Mip-Splatting [83].

Human Gaussians. 3D human Gaussians Φ_h are parameterized by Gaussian attributes ϕ_h along with SMPL-X human body pose θ and shape β parameters. The Gaussian attributes ϕ_h are defined in the canonical human pose, where each Gaussian is anchored to a surface point of the rest-posed SMPL-X [56] mesh. We then animate the Gaussians using linear blend skinning (LBS) driven by the human pose θ to obtain the final 3D human Gaussians. In this animation process, we follow ExAvatar [48] to make the Gaussians of the hands and face follow the original SMPL-X skinning weights, while those of the other body parts adopt averaged skinning weights from their neighboring SMPL-X vertices.

Object Gaussians. 3D object Gaussians Φ_o are parameterized by Gaussian attributes ϕ_o , rotation R , translation t , and scale s , where ϕ_o is defined in a canonical space. The final 3D Gaussians are obtained by applying an affine transformation with rotation R , translation t , and scale s to the canonical representation.

3.2. Reconstruction stage

In this stage, we acquire the text prompts, the initial 3D human and object, and 2D background, which serve as necessary components for the HOI optimization stage.

Text captioning. From the input image, we extract two text prompts, P_{holistic} and P_{contact} . P_{holistic} captures the global context of the human–object interaction depicted in the image, while P_{contact} specifies the human body parts (e.g., head and hands) involved in physical contact with the object. These prompts are acquired using a vision-language model, GPT-4 [1], which has powerful visual understanding capabilities learned from large-scale multimodal datasets.

Human reconstruction. To obtain an initial 3D human reconstruction, we first remove the interacting object by using SmartEraser [28]. From the object-removed image, we segment the human region and mask the background to obtain a clean human image. Based on the human segmented image, we derive the initial 3D Gaussian attributes ϕ_h with LHM [59]. Separately, we estimate the initial human pose θ and shape β parameters using Multi-HMR [3].

Object reconstruction. To obtain an initial 3D object reconstruction, we also isolate the object from the input image with SmartEraser [28] and SAM [35], acquiring a clean object image. Based on the object image, we reconstruct its 3D shape and texture as a mesh with InstantMesh [80] and subsequently convert the mesh into 3D Gaussian attributes ϕ_o using 3DGS [31]. The object pose parameters (R , t , and s) are estimated by aligning the reconstructed 3D object surface with the depth map predicted by ZoeDepth [4] from the original image.

Background reconstruction. To build a 2D background, we utilize SmartEraser [28] and remove the human and object from the input image. The 2D background image, along with 3D human Gaussians Φ_h and 3D object Gaussians Φ_o , are used for constructing realistic front- and novel-view rendering, which is crucial for the appearance loss in the HOI optimization stage (Sec. 3.3).

3.3. HOI optimization stage

In this stage, we jointly refine the initial 3D human and object reconstructions using text prompts to ensure a holistic and semantically aligned reconstruction with the text. Specifically, we optimize the 3D human Gaussians Φ_h and the 3D object Gaussians Φ_o over $N = 200$ optimization

steps, driven by the following overall loss function:

$$\mathcal{L} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{appr}} + \mathcal{L}_{\text{contact}} + \mathcal{L}_{\text{collision}}, \quad (1)$$

where $\mathcal{L}_{\text{collision}}$ is a penalization term that discourages interpenetration between the 3D human and object, following Jiang *et al.* [29]. The remaining terms are detailed below.

Reconstruction loss. The reconstruction loss $\mathcal{L}_{\text{recon}}$ is defined as the discrepancy between the input image and the front-view rendering composed of 3D human Gaussians, 3D object Gaussians, and 2D background. It consists of two mean squared error (MSE) terms: (1) loss between the rendered RGB image and the input image, and (2) loss between the rendered silhouettes and the corresponding segmentation masks of the human and object in the input image.

Appearance loss. The appearance loss $\mathcal{L}_{\text{appr}}$ is defined as the semantic distance between the holistic text prompt P_{holistic} and the novel-view rendering composed of 3D human Gaussians, 3D object Gaussians, and the 2D background. Specifically, we sample a random viewpoint uniformly over a sphere, render the 3D human and object Gaussians from that viewpoint, and composite the result onto the 2D background image. Based on the rendered image, we apply the score distillation sampling strategy [57] by leveraging the rich visual prior knowledge of a pre-trained diffusion network (*e.g.*, StableDiffusion [62]) to align the rendered appearance with the text prompt. The appearance loss is computed as

$$\nabla_{\Phi} \mathcal{L}_{\text{appr}} = \mathbb{E}[w_t(\hat{\epsilon}_t(\mathbf{x}_t; P_{\text{holistic}}) - \epsilon_t) \frac{\partial \mathbf{x}_t}{\partial \Phi}], \quad (2)$$

where t denotes the noise level, \mathbf{x}_t is the rendered image perturbed by Gaussian noise ϵ_t , and w_t is a weighting factor determined by the noise level t . The loss function minimizes the discrepancy between the predicted noise $\hat{\epsilon}_t(\cdot)$ and the true noise ϵ_t . This encourages the optimization of the 3D Gaussians, through the rendered images, to align with the distribution of plausible human-object interaction appearances learned by the pre-trained diffusion network.

The appearance loss effectively utilizes the visual prior knowledge from the pre-trained diffusion network to guide the reconstruction toward plausible human-object interactions, addressing the absence of interaction reasoning in the initial 3D human and object reconstruction. It captures holistic contextual cues beyond physical contact, enabling reasoning about non-contact interactions and object orientation related to human intent. By applying such holistic guidance, it complements the contact information and significantly improves the accuracy and plausibility of 3D human and object reconstruction.

Contact loss. The contact loss $\mathcal{L}_{\text{contact}}$ is defined as the proximity between the object surface and human body parts predicted to be in contact. We identify the set of Gaussian center points $V_{h,c}$ corresponding to the contacting body

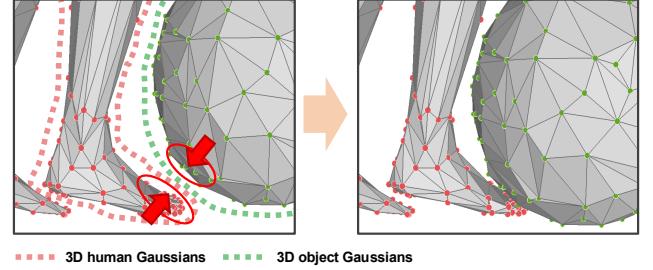


Figure 4. Gaussians-to-mesh conversion process.

parts specified in the SMPL-X body model, based on the contact text prompt P_{contact} (*e.g.*, right foot). Then, we compute the loss that minimizes the distance between the contacting 3D human points $V_{h,c}$ and their nearest 3D object points V_o , within a threshold $\tau = 10$ cm, calculated as:

$$\mathcal{L}_{\text{contact}} = \frac{1}{|V_{h,c}|} \sum_{v_h \in V_{h,c}} d(v_h, V_o) \cdot \mathbb{1}[d(v_h, V_o) < \tau]. \quad (3)$$

This loss enforces local physical plausibility between the estimated contacting regions, enhancing reconstruction accuracy in conjunction with the appearance loss $\mathcal{L}_{\text{appr}}$.

3.4. Gaussians-to-mesh conversion

Fig. 4 illustrates the process of converting the 3D Gaussians of the final human and object reconstruction into 3D human and object meshes. Naturally, the 3D Gaussians deviate from their underlying base meshes, which are the SMPL-X body model [56] for the human and the 3D object mesh from InstantMesh [80]. For contact regions, direct conversion from 3D Gaussians to a mesh can result in inconsistencies between the contacts defined by the 3D Gaussians and those defined by the corresponding mesh surfaces. To ensure consistent contact, we apply a local shift that moves the object mesh toward the human surface to resolve such inconsistencies. To this end, we identify contact regions where the distance between the 3D human and object Gaussians is less than 5 cm. Then, we select the mesh vertices on both the human and object surfaces that correspond to these contact regions and minimize the distance between these vertices to zero. We use the 3D human and object meshes after this conversion in all experiments described in Sec. 4, for comparison with existing mesh-based reconstruction methods.

4. Experiments

4.1. Datasets

Open3DHOI [74] and BEHAVE [5] datasets are used for our experiments. Open3DHOI is an open-vocabulary, in-the-wild 3D HOI dataset that we use only for evaluation. It

contains over 2.5K images and 133 object categories. BEHAVE is an indoor 3D HOI dataset that captures the interactions of 8 human subjects and 20 objects in a controlled setting. We use its official test set, which consists of 4.5K images, for the evaluation.

4.2. Evaluation metrics

For all evaluation metrics, we align the root position between the 3D human reconstruction and ground-truth (GT).

Chamfer distance (CD_{human}, CD_{object}). We evaluate the 3D human and object reconstruction using the Chamfer distance between the predicted 3D surface and the corresponding GT. The Chamfer distance is separately computed for the human (CD_{human}) and object (CD_{object}) in centimeters.

Contact score. We evaluate contact fidelity between reconstructed human and object using the F1 score of contact regions derived from their 3D surfaces. We obtain a contact map by extracting human vertices of the SMPL-X mesh surface [56] within 5 cm of the object mesh. Then, the contact score (Contact) is computed as the harmonic mean of precision and recall, defined as $2 \times \frac{PR}{(P+R)}$, where P and R denote precision and recall. Note that contact evaluation is only performed on the SMPL-X mesh surface, since GT contacts are defined on the SMPL-X topology and are incompatible with the topology-free Gaussian representation.

Collision. To evaluate physical plausibility, we evaluate collision by measuring the interpenetration between the reconstructed 3D human and object. We compute the percentage of human vertices located within the object mesh.

4.3. Ablation study

We carry out all ablation studies on Open3DHOI [74].

Effectiveness of text-guided optimization. Fig. 5 and Tab. 1 demonstrate that our optimization stage effectively refines the 3D human and object by capturing semantic interactions described in text descriptions. When we remove the text prompt condition from the appearance loss $\mathcal{L}_{\text{appr}}$, the reconstruction fails to capture the global context of the human-object interaction. As shown in Fig. 6, optimization without text conditioning fails to orient the human’s gaze toward the right hand, whereas incorporating text information corrects this implicit context, resulting in accurate 3D human and object reconstruction. This is mainly because the text descriptions provide the holistic context about the interaction beyond the physical contact information. Moreover, the texts provide the semantic prior knowledge for inferring non-contact interactions, which cannot be derived from the contact information. Thus, leveraging the comprehensive textual information enables our framework to produce accurate and plausible reconstructions that capture the global interaction context.

Ablation on loss configurations. Fig. 7 and Tab. 2 demonstrate the effectiveness of our loss design, comprising ap-

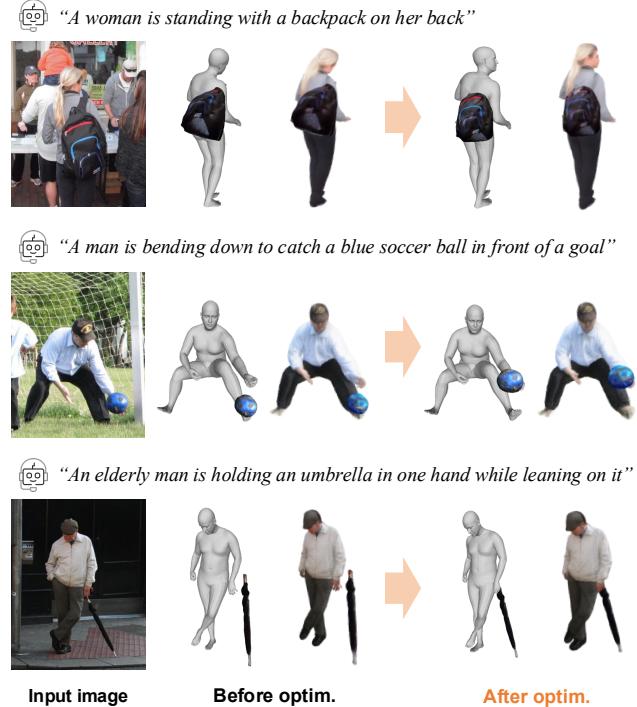


Figure 5. **Optimization results of TeHOR.** Our text-guided optimization accurately refines the 3D human and object by utilizing their corresponding text descriptions.

pearance loss $\mathcal{L}_{\text{appr}}$ and contact loss $\mathcal{L}_{\text{contact}}$. The first row of the table shows that removing the appearance loss fails to capture holistic interaction cues, leading to implausible reconstructions. As shown in Fig. 7, the teddy bear held by the person is displaced from the chest and even penetrates the arm. With the inclusion of appearance loss, contextual cues from the input image are effectively captured and reflected in the 3D reconstruction, resulting in a consistent human-object interaction. To further examine the effectiveness of our appearance loss, we replace the appearance loss $\mathcal{L}_{\text{appr}}$ with CLIP loss $\mathcal{L}_{\text{CLIP}}$, which also provides appearance-level supervision using the textual information. Specifically, the CLIP loss function minimizes the cosine distance between the feature embeddings of the rendered image and the corresponding text prompt within the pre-trained CLIP embedding space [61]. As shown in the third and fourth rows of Tab. 2, our appearance loss achieves significantly better object Chamfer distance (CD_{object}) and contact F1-score. CLIP encodes text features into a single 1D embedding vector, which limits its ability to model dense spatial human-object relationships. In contrast, our appearance loss $\mathcal{L}_{\text{appr}}$ applies dense, pixel-level gradients in rendered 2D appearances, distilling rich priors from the pre-trained diffusion network to enable accurate, fine-grained 3D reconstruction. Accordingly, our appearance loss effectively optimizes the 3D human and object to follow the ap-



Figure 6. Effectiveness of text descriptions in optimization.

	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
Before optim.	5.252	31.268	0.305	0.040
After optim. w/o text prompts	5.028	20.348	0.374	0.052
After optim. (Ours)	4.941	16.701	0.412	0.047

Table 1. Effectiveness of text-guided optimization.

pearance cues under the diffusion prior, ensuring both plausible and accurate 3D reconstruction.

Ablation on appearance rendering. Tab. 3 investigates the impact of key components in our framework that enable realistic appearance rendering. We first evaluate the advantage of the 3D Gaussian representation over mesh, which is a widely used representation in 3D human and object reconstruction [9, 13, 85]. As shown in the first row, 3D Gaussians significantly outperform mesh representation for both the human and object, due to two main factors. First, 3D Gaussians excel at modeling high-fidelity visual appearances, providing richer signals that enable the appearance loss $\mathcal{L}_{\text{appr}}$ to more accurately align the reconstruction with the textual semantics. Second, their flexible and topology-free structure allows for more effective optimization of the 3D spatial relationships between the human and object. We also conduct an ablation study on the use of a 2D background for realistic rendering. The second row of Tab. 3 shows a significant degradation in performance when the 2D background is removed. This demonstrates that the 2D background plays a crucial role in providing complete scene context, which allows the appearance loss to fully exploit the image prior knowledge from the diffusion network, leading to more precise optimization.

4.4. Comparison with state-of-the-art methods

We compare our framework with state-of-the-art 3D human and object reconstruction methods. Since PICO [9] requires human–object contact as input, we run the method using contact estimation results from LEMON [82]. For a fair comparison, we use the same 3D human and object initializations for all methods. Specifically, the initial 3D human pose is obtained from Multi-HMR [3], and the initial 3D object pose is estimated by aligning the reconstructed 3D object with the depth map predicted by ZoeDepth [4]. All



Figure 7. Effectiveness of each loss function in our framework.

L_{appr}	L_{contact}	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
✗	✓	5.191	22.094	0.330	0.049
✓	✗	5.311	19.849	0.374	0.054
L_{CLIP}	✓	5.018	18.504	0.366	0.047
✓ (Ours)	✓	4.941	16.701	0.412	0.047

Table 2. Ablation studies for loss configurations.

	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
3D Gaussians → Mesh	5.153	25.162	0.308	0.054
w/o 2D background	5.002	18.196	0.389	0.049
TeHOR (Ours)	4.941	16.701	0.412	0.047

Table 3. Ablation studies for appearance rendering.

object shapes are reconstructed using InstantMesh [80]. As TeHOR is a Gaussian-based framework, we convert the initial object mesh into 3D Gaussians via 3DGS [31]. The † symbol denotes the evaluation that directly compares the reconstructed 3D Gaussians against the GT mesh surfaces, without Gaussian-to-mesh conversion.

Fig. 8 and Tab. 4 show that our TeHOR largely outperforms all state-of-the-art methods both qualitatively and quantitatively. As previous methods rely heavily on local geometric cues, such as human–object contact information, they often fail to handle interactions that require global context, including object orientation. Although HOI-Gaussian [74] uses depth map from the input image as additional information for the reconstruction, it is unable to capture the semantic context of the interaction and is vulnerable to severe human-object occlusion. Unlike the previous methods, our TeHOR leverages rich text descriptions of human–object interactions as a key prior, providing holistic and semantic context that contact information alone cannot capture. Additionally, Tab. 5 demonstrates our superiority in non-contact scenarios, where we evaluate the methods on a subset of Open3DHOI that excludes samples with physical contact between the ground-truth 3D human and object. Such cases are particularly challenging because contact-based cues vanish entirely, forcing the reconstruction system to reason about the interaction from global contextual signals such as object orientation, gaze direction, and body posture. While previous methods fail to reason about

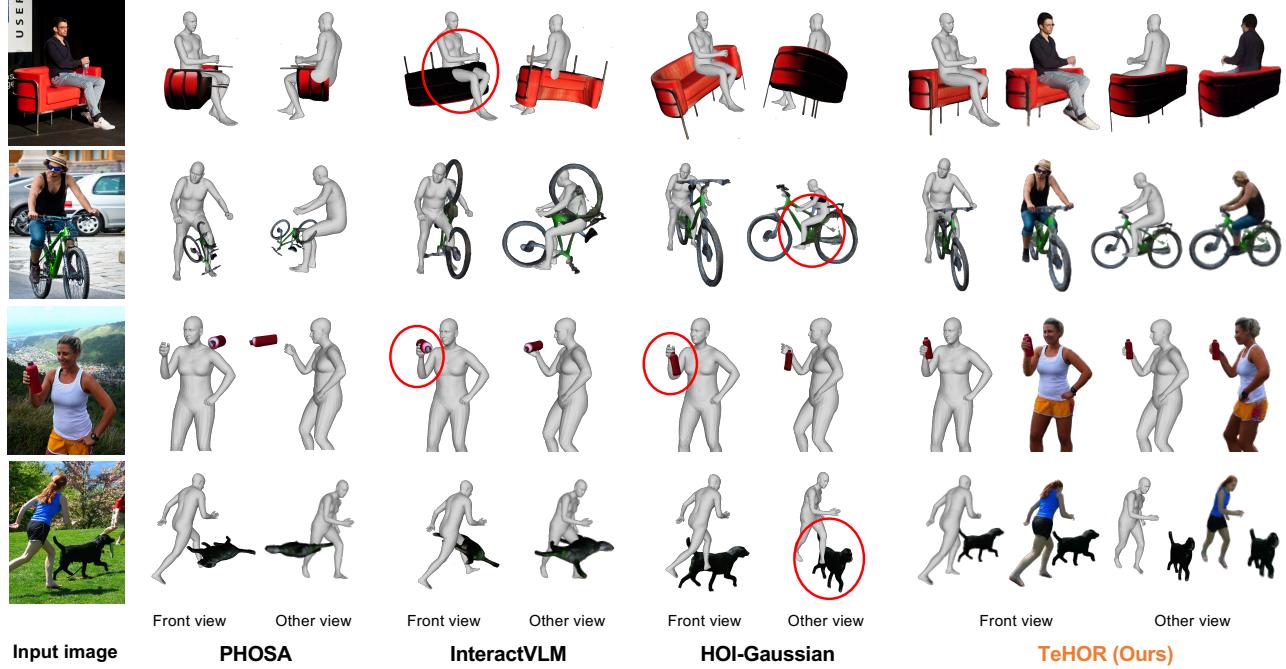


Figure 8. **Qualitative comparison with state-of-the-art methods.** We highlight their representative failure cases with red circles.

Methods	Open3DHOI				BEHAVE			
	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
PHOSA [85]	5.342	49.180	0.243	0.044	5.758	46.003	0.257	0.010
LEMON [82] + PICO [9]	5.948	25.889	0.335	0.078	6.159	22.585	0.082	0.045
InteractVLM [13]	5.252	24.238	<u>0.392</u>	0.054	5.770	<u>19.197</u>	<u>0.379</u>	0.021
HOI-Gaussian [74]	<u>5.111</u>	<u>19.363</u>	0.348	0.070	<u>5.748</u>	21.774	0.371	0.019
TeHOR (Ours)	4.941	16.701	0.412	<u>0.047</u>	5.615	17.339	0.412	<u>0.016</u>
TeHOR[†] (Ours)	4.403	16.697	–	0.045	5.241	17.341	–	0.012

Table 4. **Quantitative comparison with state-of-the-art methods on Open3DHOI [74] and BEHAVE [5].** Bold and underlined values indicate the best and second-best scores, respectively. [†] denotes the evaluation results directly computed based on the centers of reconstructed 3D Gaussians instead of 3D mesh vertices.

Methods	CD _{human} ↓	CD _{object} ↓	Collision↓
PHOSA [85]	5.401	65.537	0.028
LEMON [82] + PICO [9]	5.635	33.073	0.029
InteractVLM [13]	5.390	46.819	0.011
HOI-Gaussian [74]	5.244	25.374	0.037
TeHOR (Ours)	4.958	17.546	0.005

Table 5. **Quantitative comparison with state-of-the-art methods for non-contact scenarios on Open3DHOI [74].**

non-contact interaction, our framework benefits from text descriptions, achieving superior performance in the non-contact scenarios. Overall, our framework, which effectively optimizes the 3D human and object using these comprehensive text descriptions, provides more accurate and robust reconstructions than previous methods.

5. Conclusion

We propose **TeHOR**, a text-guided framework for joint reconstruction of the 3D human and object from a single image. Our framework leverages text descriptions to enforce semantic alignment, enabling reasoning over a wide spectrum of interactions, including non-contact cases. It further incorporates appearance cues from the 3D human and object to capture holistic contextual information, ensuring visual plausibility of the reconstruction. Extensive experiments demonstrate that our proposed framework produces accurate and semantically plausible reconstructions, achieving state-of-the-art performance.

TeHOR: Text-Guided 3D Human and Object Reconstruction with Textures

Supplementary Material

In this supplementary material, we present additional technical details and more experimental results that could not be included in the main manuscript due to the lack of pages. The contents are summarized below:

- **S1.** Visualization in video format
- **S2.** Evaluation on semantic alignment
- **S3.** Impact of contact estimation accuracy
- **S4.** Impact of Gaussian attributes optimization
- **S5.** Impact of Gaussians-to-mesh conversion
- **S6.** Details of text captioning
- **S7.** Implementation details
- **S8.** More qualitative results
- **S9.** Limitations and future work

S1. Visualization in video format

The supplementary video "**TeHOR_Supp1.mp4**" consists of three parts. The first part visualizes the optimization process of our framework, demonstrating how text-derived semantic cues guide the reconstruction. Representative examples include 1) adjusting the arms to hug a teddy bear, 2) following the human's gaze toward the water bottle, and 3) positioning a leg on the bench. In particular, the gaze-following case highlights that our framework reasons about global interaction context beyond explicit contact. The second part compares our proposed TeHOR with state-of-the-art reconstruction methods: InteractVLM [13] and HOI-Gaussian [74]. The last part shows 360-degree renderings of the final reconstructions.

S2. Evaluation on semantic alignment

In this section, we introduce additional evaluation to compare TeHOR with state-of-the-art reconstruction methods on semantic alignment between 3D reconstructions and text descriptions. Since direct comparison between 3D reconstructions and text is not feasible, we instead evaluate appearance-text alignment metrics on 2D renderings of the reconstructions, following the process shown in Fig. S1. To ensure a fair comparison, we unify the underlying 3D representation across all methods, since existing methods predominantly use mesh-based representations, whereas our framework is Gaussian-based. Specifically, we use the same initial 3D human and object Gaussians, including both shape and texture attributes, for all methods. We then extract each method's human (θ and β) and object (R, t , and s) pose parameters and apply them to transform the 3D Gaus-

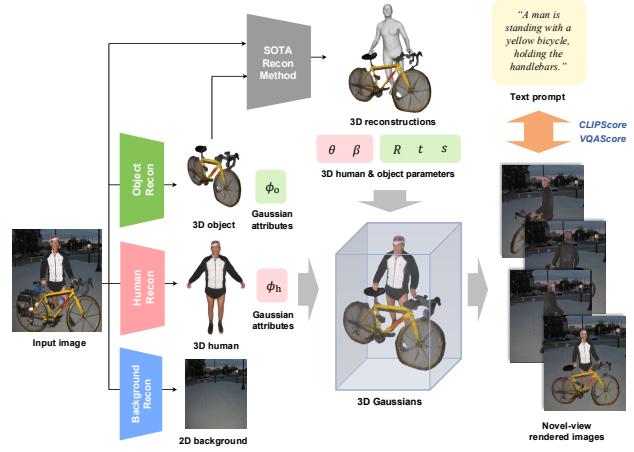


Figure S1. **Process of evaluating text alignment for state-of-the-art reconstruction methods.**

Methods	CLIPScore↑	VQAScore↑
PHOSA [85]	0.689	0.631
LEMON [82] + PICO [9]	0.696	0.642
InteractVLM [13]	0.694	0.647
HOI-Gaussian [74]	0.698	0.648
TeHOR (Ours)	0.706	0.652

Table S1. **Quantitative evaluation of appearance-text alignment on Open3DHOI [74].**

sians. This setup ensures that the only variable in the experiments is the set of 3D pose parameters provided by each method. For evaluation, we render the transformed Gaussians on the 2D background from pre-defined viewpoints 0° , 90° , 180° , and 270° . Each rendered image is then paired with its corresponding text description to compute two image-text alignment metrics: 1) CLIPScore [21] and 2) VQAScore [40]. CLIPScore computes the cosine similarity between the embeddings of the rendered image and the text description. VQAScore utilizes a powerful visual-question-answering (VQA) model to compute the alignment score by converting the text description into a simple query and measuring the generative likelihood of a desired response. Here, we use InstructBLIP-FlanT5-XL [10] as the underlying VQA model to compute VQAScore. Tab. S1 shows that our framework outperforms other state-of-the-art methods in text alignment by effectively capturing the holistic and semantic context of human-object interaction.

Contact estimation methods	Contact estimation			3D reconstruction			
	Contact _p ↑	Contact _r ↑	Contact _{f1} ↑	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
w/o contact	—	—	—	5.311	19.849	0.374	0.054
P_{Contact}	0.282	0.342	0.309	4.941	16.701	0.412	0.047
DECO [70]	0.200	0.264	0.228	5.115	17.229	0.353	0.051
LEMON [82]	0.426	0.225	0.295	5.084	17.060	0.389	0.050
InteractVLM [13]	0.422	0.458	0.439	4.988	16.009	0.408	0.052

Table S2. Impact of contact estimation accuracy on TeHOR’s reconstruction performance.



Figure S2. Enhancement of Gaussian via optimization process.

S3. Impact of contact estimation accuracy

Tab. S2 illustrates the impact of contact estimation accuracy on the 3D reconstruction performance of our framework. We compare our framework under different contact estimation settings, including specialized contact prediction models such as LEMON [82] and InteractVLM [13]. While these models improve the precision of contact localization on human and object surfaces, their contribution to final 3D reconstruction quality remains marginal. Moreover, as shown in Tab. 4, even with less accurate contact estimation, our framework still achieves superior 3D reconstruction performance compared to other reconstruction methods that rely on more accurate contact prediction. The contact estimation methods primarily focus on accurately predicting the boundaries of the contact region at a fine-grained level. However, regardless of how precise these contact boundaries are, the contact information alone cannot capture the holistic and semantic context of human-object interaction. This observation suggests that capturing the holistic interaction context is far more important for the joint reconstruction of 3D human and object than precisely delineating fine-grained contact boundaries. Accordingly, the core strength of our framework is determined by holistic contact reasoning supported by text-guided optimization rather than by accurate contact prediction. This validates our use of the contact text prompt P_{contact} as a lightweight yet effective alternative to external contact prediction models.

S4. Impact of Gaussian attributes optimization

Fig. S2 demonstrates the importance of optimizing the 3D human and object Gaussian attributes (ϕ_h and ϕ_o) within our framework. Since the initial Gaussian attributes can occasionally be incomplete due to occlusions in the input image,

	CD _{human} ↓	CD _{object} ↓	Contact↑	Collision↓
Before conversion	5.020	16.987	0.394	0.052
After conversion (Ours)	4.941	16.701	0.412	0.047

Table S3. Impact of Gaussians-to-mesh conversion

we further refine them through the appearance loss L_{appr} . This optimization process enhances the visual plausibility and overall coherence of the reconstructed human–object interactions. As there is no existing 3D HOI dataset that provides both geometry and texture annotations, quantitative evaluation of this optimization remains challenging; thus, we primarily present qualitative results.

S5. Impact of Gaussians-to-mesh conversion

Tab. S3 shows that our Gaussians-to-mesh conversion process, detailed in Sec. 3.4, is a crucial step for accurate mesh reconstruction. Direct conversion of 3D Gaussians to mesh surfaces often produces inconsistencies near contact regions. Accordingly, we use the conversion procedure that enforces geometric consistency between Gaussian-defined contact regions and the corresponding mesh vertices. As a result, it improves overall geometric accuracy and yields substantial gains in the contact evaluation score (Contact).

S6. Details of text captioning

Fig. S3 and Fig. S4 illustrate the two captioning instructions used to generate the holistic text prompt P_{holistic} and contact text prompt P_{contact} , with the GPT-4 [1] vision-language model (VLM). First, we generate the holistic prompt P_{holistic} , which describes the interaction between the person closest to the image center and the object most directly involved with that person. Then, we generate the contact prompt P_{contact} by providing both the input image and the holistic description P_{holistic} as inference cues, enabling it to infer which human body parts are in direct physical contact with the object. This two-stage captioning strategy allows us to separately infer the global interaction semantics and the localized contact information.

Fig. S5 highlights the strong capability of the text captioning process. As shown in the examples, it successfully captures key contextual cues essential for reasoning

```

### TASK ###
Your goal is to provide a detailed description of the
given image, which depicts the interaction between a
person and an object.

- Focus only on the person whose body center is closest
  to the image center.
- Identify the object most directly interacted with and
  state the action.
- Output must be one sentence, no explanations, labels,
  or reasoning.
- Additionally, explicitly output the single object that
  is most directly interacted with.

### OUTPUT FORMAT ###
Output: {{interacting object}}, {{description}}

### OUTPUT EXAMPLE ###
Example 1 - Output: soccer ball, A woman is playing
soccer on a grassy field, dribbling the ball.
Example 2 - Output: small box, A man is seated on a
small box with legs crossed.
Example 3 - Output: chair, A woman is moving a chair
with one hand.

```

Figure S3. Captioning instruction for the VLM [1] to acquire holistic text prompt P_{holistic} .

about human–object interactions, including the human’s action (*e.g.*, sitting, riding, and performing) and the surrounding environment (*e.g.*, pathway, mid-air, and grassy field). Even when the same object appears in different interaction scenarios, the VLM provides accurate and semantically appropriate descriptions. This demonstrates the richness of holistic contextual information, in contrast to contact cues that convey only local geometric proximity. Such comprehensive interaction cues play a crucial role in guiding our reconstruction framework toward more accurate and globally coherent 3D human and object reconstructions.

S7. Implementation details

We explain the implementation details of two stages: the reconstruction stage (Sec. 3.2) and HOI optimization stage (Sec. 3.3), below. PyTorch [54] is used for implementation.

S7.1. Reconstruction stage

Human reconstruction. When using SmartEraser [28], the object regions to be removed are inpainted using classifier-free guidance with a guidance scale of 1.5 in its generative diffusion network. To segment human region from the inpainted image, we use Grounded-SAM [35, 44] with a text prompt corresponding to the object category name obtained from the text captioning (Sec. S6). From the segmented human image, LHM operates on a canonical set of 40,000 Gaussian anchors uniformly sampled over the SMPL-X surface. For each anchor, LHM predicts the Gaussian attributes ϕ_h , including canonical offsets, opacity, scale, and appearance features, through a single feed-forward inference.

Object reconstruction. When using SmartEraser [28], we adopt the same settings as in the human reconstruc-

```

### TASK ###
Your goal is to list the body parts of the person in the
given image that are in direct physical contact
with the object.

- Choose ONLY from this pre-defined list (multi-select
  allowed): head, hips, ...
- The interacting object and reference description are
  provided as follows: "{object}", "{description}".
- Focus only on the person whose body center is closest
  to the image center.
- Identify the object most directly interacted with and
  state the action.
- LEFT/RIGHT must be relative to the person (egocentric)
  , not the viewer/camera.
- If no clear physical contact is visible, output none.

### OUTPUT FORMAT ###
Output: {{comma-separated list}}

### OUTPUT EXAMPLE ###
Example 1 - Output: left hand, right hand
Example 2 - Output: right foot
Example 3 - Output: none

```

Figure S4. Captioning instruction for the VLM [1] to acquire contact text prompt P_{contact} .

tion stage. To segment object region from the inpainted image, we use Grounded-SAM [35, 44] with the text prompt “human”. From the segmented object image, InstantMesh [80] first synthesizes six multi-view images using Zero123++ [65] with 75 diffusion steps, and then reconstructs a textured mesh through its triplane-based reconstruction network. The resulting textured mesh is subsequently converted into 3D object Gaussians ϕ_o , where the Gaussian centroids are placed at the mesh vertex positions and their initial appearance features are assigned from the mesh vertex colors. The Gaussian attributes are further optimized to match the 2D images rendered from the reconstructed textured mesh at 360 uniformly sampled viewpoints, following the optimization procedure of 3DGS [31].

S7.2. HOI optimization stage

We use the Adam [34] optimizer with an exponentially decaying learning rate. The initial learning rate is set to 1×10^{-2} for the object pose parameters (R , t , and s), 1×10^{-4} for the human pose parameters (θ and β), and 1×10^{-4} for the human and object Gaussian attributes (ϕ_h and ϕ_o). We optimize the network for $N = 200$ steps on a single NVIDIA RTX 8000 GPU. Under this setting, the average optimization time per sample is 134 seconds.

Appearance rendering. During optimization, we render the 3D human Gaussians Φ_h and object Gaussians Φ_o using a spherical coordinate system (r, v, ψ) , where r denotes the distance to the spherical origin, v the elevation angle, and ψ the azimuth angle. We uniformly sample viewpoints with $r \in [1.0, 2.5]$, $v \in [-30^\circ, 30^\circ]$, and $\psi \in [-180^\circ, 180^\circ]$. Since human-object interaction primarily involves the upper body, such as the head and hands, we additionally



Figure S5. Text captioning results on Open3DHOI [74]. Our text captioning produces accurate and rich text descriptions for a wide range of interaction scenarios.

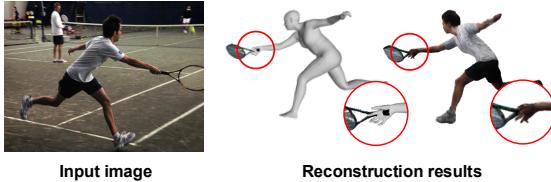


Figure S6. Failure case of reconstructing local details.

use zoomed-in camera views focused on this region. For these upper-body views, we set the spherical origin to the 3D position of the SMPL-X spine keypoint and sample $r \in [0.7, 1.5]$, $v \in [-30^\circ, 30^\circ]$, and $\psi \in [-180^\circ, 180^\circ]$.

Appearance loss. We compute the appearance loss $\mathcal{L}_{\text{appr}}$ of Eq. (2) using StableDiffusion-v2.1 [62] and apply classifier-free guidance [22] with a guidance scale of 15.0 for noise estimation. The noise levels are defined at randomly sampled timesteps within $[0.02, 0.98]$. To ensure stable optimization, we clip the loss gradients to a maximum norm of 1.0.

S8. More qualitative results

We provide additional qualitative comparison results of our TeHOR in Figs. S7 to S10. These examples further demonstrate the effectiveness of our method in reconstructing realistic and semantically coherent human–object interactions. Please note that the left-side results of TeHOR are mesh-

based renderings, while the right-side results are Gaussian-based renderings. Due to the inherent characteristics of 3D Gaussian representations, Gaussian renderings can appear slightly larger and exhibit blurred boundaries.

S9. Limitations and future work

Reconstruction of local details. While our framework captures holistic human–object interactions effectively, it may overlook fine-grained local details such as small accessories or subtle surface deformations, as shown in Fig. S6. This limitation occurs because the appearance loss of our framework primarily offers global guidance and lacks fine-grained supervision that specifically addresses local regions. A promising future direction is to design localized, text-driven supervision that specializes in local regions to further enhance fine-detail reconstruction.

Video as input. Our framework aims to jointly reconstruct 3D human and object from a single image. When extending the method to video input, additional considerations become essential, such as maintaining temporal consistency across frames and ensuring consistent geometry and texture over time. With the recent emergence of text-to-video generative models [6, 23], future work could leverage these advances by using text descriptions as a key guidance, enabling more stable and temporally coherent 3D HOI reconstruction.

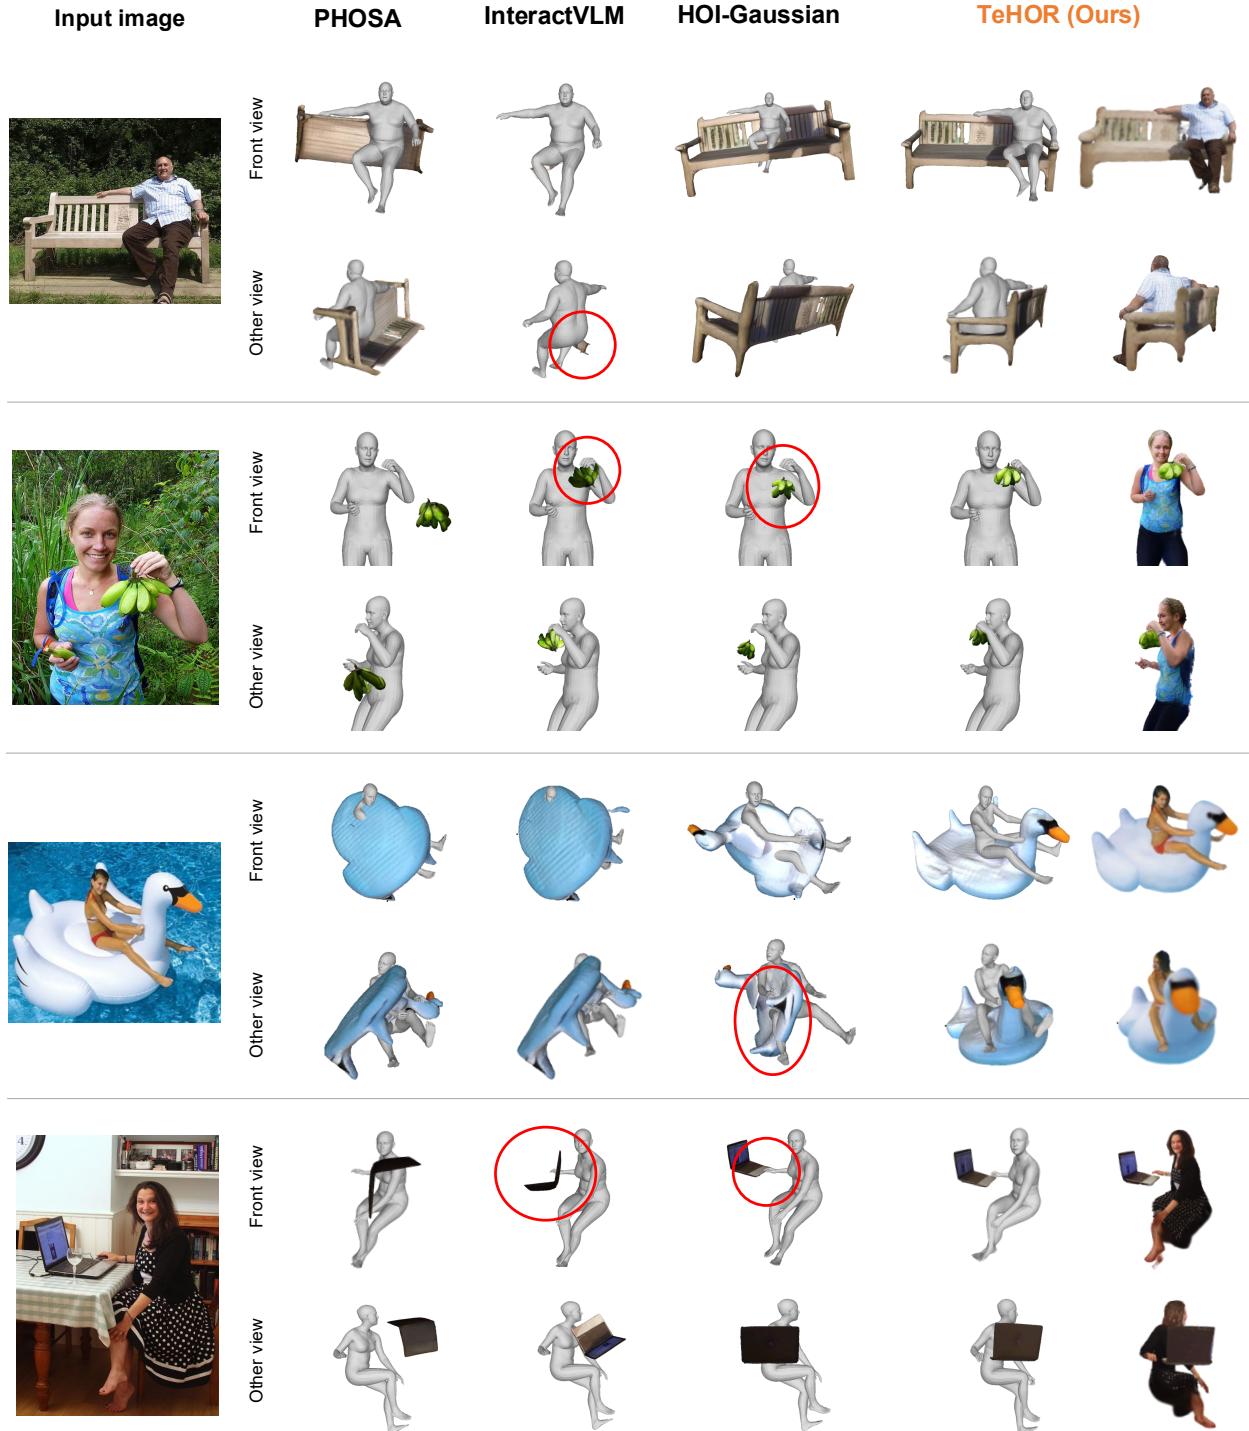


Figure S7. **More qualitative comparison of 3D human and object reconstruction with PHOSA [85], InteractVLM [13], and HOI-Gaussian [74], on Open3DHOI [74].** We highlight their representative failure cases with red circles.

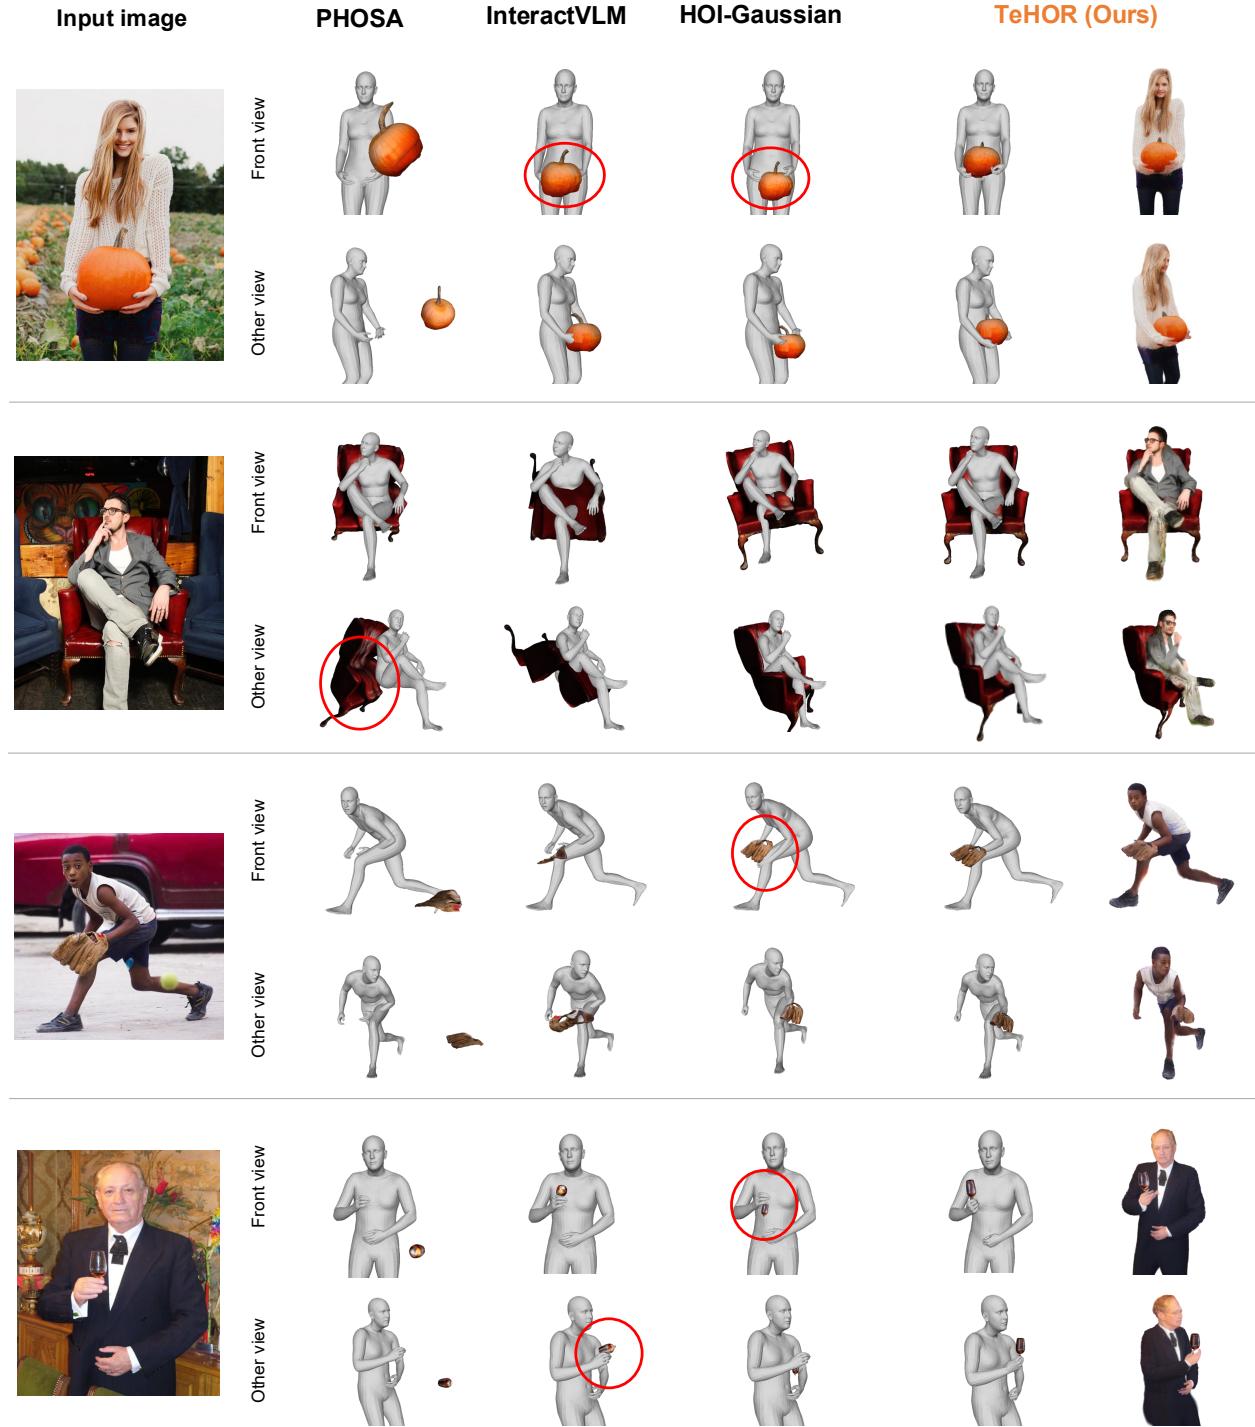


Figure S8. More qualitative comparison of 3D human and object reconstruction with PHOSA [85], InteractVLM [13], and HOI-Gaussian [74], on Open3DHOI [74]. We highlight their representative failure cases with red circles.

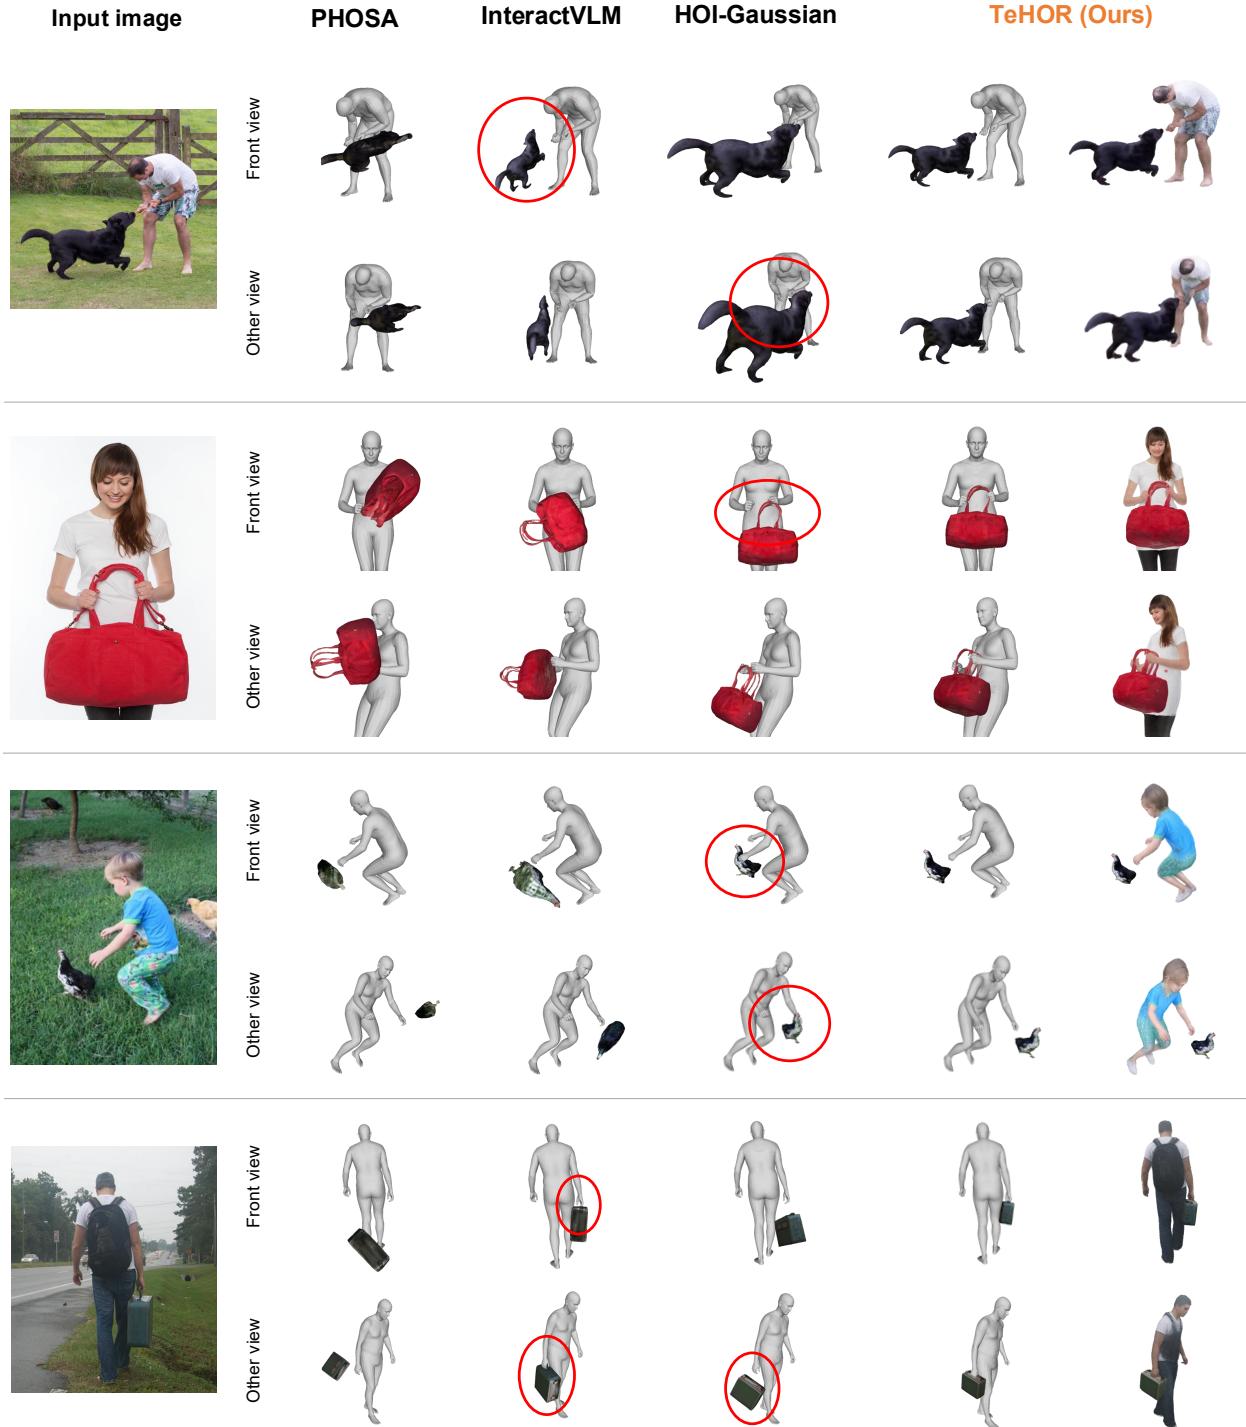


Figure S9. **More qualitative comparison of 3D human and object reconstruction with PHOSA [85], InteractVLM [13], and HOI-Gaussian [74], on Open3DHOI [74].** We highlight their representative failure cases with red circles.

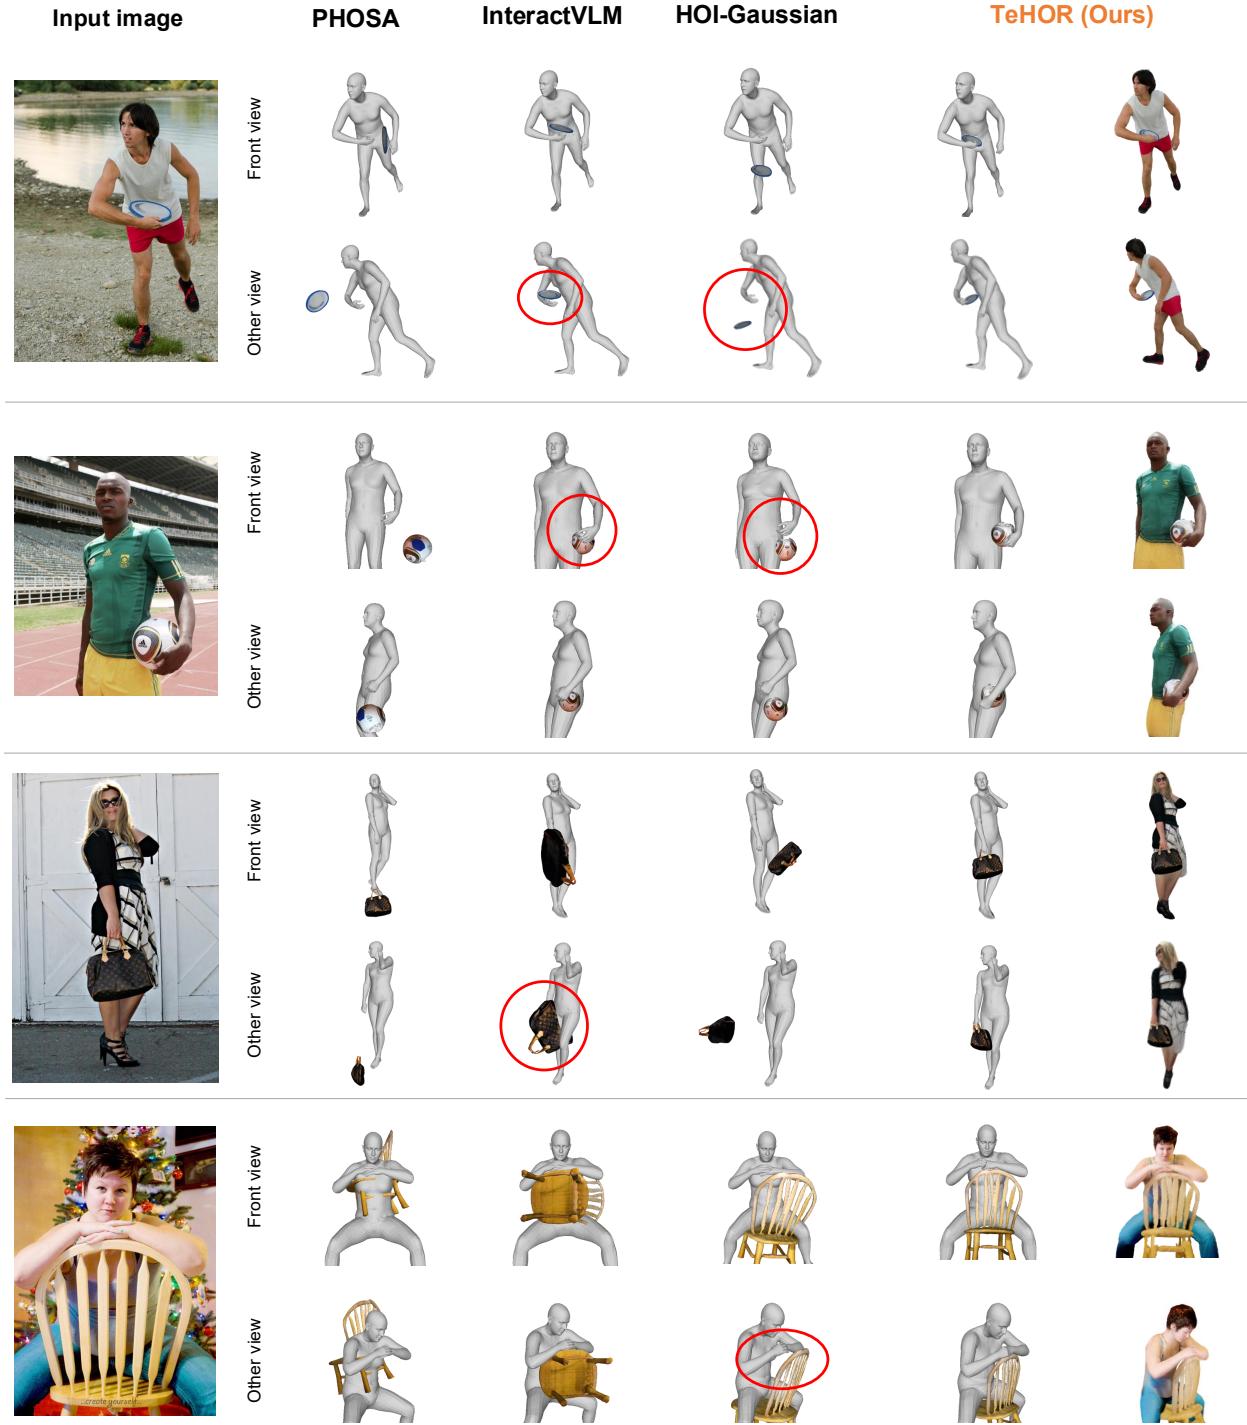


Figure S10. More qualitative comparison of 3D human and object reconstruction with PHOSA [85], InteractVLM [13], and HOI-Gaussian [74], on Open3DHOI [74]. We highlight their representative failure cases with red circles.

References

[1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2, 4, 3

[2] Thiendo Alldieck, Mihai Zanfir, and Cristian Sminchisescu. Photorealistic monocular 3D reconstruction of humans wearing clothing. In *CVPR*, 2022. 3

[3] Fabien Baradel, Matthieu Armando, Salma Galaaoui, Romain Brégier, Philippe Weinzaepfel, Grégory Rogez, and Thomas Lucas. Multi-HMR: Multi-person whole-body human mesh recovery in a single shot. In *ECCV*, 2024. 3, 4, 7

[4] Shariq Farooq Bhat, Reiner Birk, Diana Wofk, Peter Wonka, and Matthias Müller. ZoeDepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 4, 7

[5] Bharat Lal Bhatnagar, Xianghui Xie, Ilya A Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *CVPR*, 2022. 3, 5, 8

[6] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023. 4

[7] Yixin Chen, Siyuan Huang, Tao Yuan, Siyuan Qi, Yixin Zhu, and Song-Chun Zhu. Holistic++ scene understanding: Single-view 3D holistic scene parsing and human pose estimation with human-object interaction and physical commonsense. In *ICCV*, 2019. 3

[8] Yixin Chen, Sai Kumar Dwivedi, Michael J Black, and Dimitrios Tzionas. Detecting human-object contact in images. In *CVPR*, 2023. 3

[9] Alpár Cseke, Shashank Tripathi, Sai Kumar Dwivedi, Arjun Lakshminath, Agniv Chatterjee, Michael J. Black, and Dimitrios Tzionas. PICO: Reconstructing 3D people in contact with objects. In *CVPR*, 2025. 1, 7, 8

[10] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. InstantBLIP: Towards general-purpose vision-language models with instruction tuning. *NeurIPS*, 2023. 1

[11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional Transformers for language understanding. In *NAACL*, 2019. 3

[12] Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, Yao Feng, and Michael J Black. TokenHMR: Advancing human mesh recovery with a tokenized pose representation. In *CVPR*, 2024. 3

[13] Sai Kumar Dwivedi, Dimitrije Antić, Shashank Tripathi, Omid Taheri, Cordelia Schmid, Michael J Black, and Dimitrios Tzionas. InteractVLM: 3D interaction reasoning from 2D foundational models. In *CVPR*, 2025. 1, 3, 7, 8, 2, 5, 6

[14] Yao Feng, Vasileios Choutas, Timo Bolkart, Dimitrios Tzionas, and Michael J Black. Collaborative regression of expressive bodies using moderation. In *3DV*, 2021. 3

[15] Ruiqi Gao, Aleksander Holynski, Philipp Henzler, Arthur Brussee, Ricardo Martin-Brualla, Pratul Srinivasan, Jonathan T Barron, and Ben Poole. CAT3D: Create anything in 3D with multi-view diffusion models. In *NeurIPS*, 2024. 3

[16] Alexey Gavryushin, Yifei Liu, Daoji Huang, Yen-Ling Kuo, Julien Valentin, Luc van Gool, Otmar Hilliges, and Xi Wang. ROMEO: Revisiting optimization methods for reconstructing 3D human-object interaction models from images. In *ECCVW*, 2024. 3

[17] Sookwan Han and Hanbyul Joo. CHORUS: Learning canonicalized 3D human-object spatial relations from unbounded synthesized images. In *ICCV*, 2023. 3

[18] Mohamed Hassan, Partha Ghosh, Joachim Tesch, Dimitrios Tzionas, and Michael J Black. Populating 3D scenes by learning human-scene interaction. In *CVPR*, 2021. 3

[19] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH: Animatable reconstruction of clothed humans. In *CVPR*, 2020. 3

[20] Tong He, Yuanlu Xu, Shunsuke Saito, Stefano Soatto, and Tony Tung. ARCH++: Animation-ready clothed human reconstruction revisited. In *ICCV*, 2021. 3

[21] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. CLIPScore: A reference-free evaluation metric for image captioning. *EMNLP*, 2021. 1

[22] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS workshop*, 2021. 4

[23] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *NeurIPS*, 2022. 4

[24] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovsky, Daniel Scharstein, and Michael J Black. Capturing and inferring dense full-body human-scene contact. In *CVPR*, 2022. 3

[25] Yinghao Huang, Omid Taheri, Michael J Black, and Dimitrios Tzionas. InterCap: Joint markerless 3D tracking of humans and objects in interaction. In *GCPR*, 2022. 3

[26] Yangyi Huang, Hongwei Yi, Yuliang Xiu, Tingting Liao, Jiaxiang Tang, Deng Cai, and Justus Thies. TeCH: Text-guided reconstruction of lifelike clothed humans. In *3DV*, 2024. 3

[27] Zixuan Huang, Stefan Stojanov, Anh Thai, Varun Jampani, and James M Rehg. ZeroShape: Regression-based zero-shot shape reconstruction. In *CVPR*, 2024. 3

[28] Longtao Jiang, Zhendong Wang, Jianmin Bao, Wengang Zhou, Dongdong Chen, Lei Shi, Dong Chen, and Houqiang Li. SmartEraser: Remove anything from images using masked-region guidance. In *CVPR*, 2025. 4, 3

[29] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *CVPR*, 2020. 5

[30] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. 3

[31] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM TOG*, 2023. 3, 4, 7

[32] Hyeonwoo Kim, Sookwan Han, Patrick Kwon, and Hanbyul Joo. Beyond the contact: Discovering comprehensive affordance for 3D objects from pre-trained 2D diffusion models. In *ECCV*, 2024. 3

[33] Hyeonwoo Kim, Sangwon Baik, and Hanbyul Joo. DAViD: Modeling dynamic affordance of 3D objects using pre-trained video diffusion models. In *ICCV*, 2025. 3

[34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 3

[35] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, 2023. 4, 3

[36] Muhammed Kocabas, Chun-Hao P Huang, Otmar Hilliges, and Michael J Black. PARE: Part attention regressor for 3D human body estimation. In *CVPR*, 2021. 3

[37] Ao Li, Jinpeng Liu, Yixuan Zhu, and Yansong Tang. Score-HOI: Physically plausible reconstruction of human-object interaction via score-guided diffusion. In *ICCV*, 2025. 1, 3

[38] Zhihao Li, Jianzhuang Liu, Zhensong Zhang, Songcen Xu, and Youliang Yan. CLIFF: Carrying location information in full frames into human pose and shape estimation. In *ECCV*, 2022. 3

[39] Tingting Liao, Xiaomei Zhang, Yuliang Xiu, Hongwei Yi, Xudong Liu, Guo-Jun Qi, Yong Zhang, Xuan Wang, Xianyu Zhu, and Zhen Lei. High-fidelity clothed avatar reconstruction from a single image. In *CVPR*, 2023. 3

[40] Zhiqiu Lin, Deepak Pathak, Baiqi Li, Jiayao Li, Xide Xia, Graham Neubig, Pengchuan Zhang, and Deva Ramanan. Evaluating text-to-visual generation with image-to-text generation. In *ECCV*, 2024. 1

[41] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 3

[42] Minghua Liu, Chao Xu, Haian Jin, Linghao Chen, Mukund Varma T, Zexiang Xu, and Hao Su. One-2-3-45: Any single image to 3D mesh in 45 seconds without per-shape optimization. In *NeurIPS*, 2023. 3

[43] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3D object. In *ICCV*, 2023. 3

[44] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding DINO: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 2024. 3

[45] Xiaoxiao Long, Yuan-Chen Guo, Cheng Lin, Yuan Liu, Zhiyang Dou, Lingjie Liu, Yuexin Ma, Song-Hai Zhang, Marc Habermann, Christian Theobalt, et al. Wonder3D: Single image to 3D using cross-domain diffusion. In *CVPR*, 2024. 3

[46] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. SMPL: A skinned multi-person linear model. *ACM TOG*, 2015. 3

[47] Gyeongsik Moon, Hyeongjin Nam, Takaaki Shiratori, and Kyoung Mu Lee. 3D clothed human reconstruction in the wild. In *ECCV*, 2022. 3

[48] Gyeongsik Moon, Takaaki Shiratori, and Shunsuke Saito. Expressive whole-body 3D gaussian avatar. In *ECCV*, 2024. 4

[49] Hyeongjin Nam, Daniel Sungho Jung, Yeonguk Oh, and Kyoung Mu Lee. Cyclic test-time adaptation on monocular video for 3D human mesh reconstruction. In *ICCV*, 2023. 3

[50] Hyeongjin Nam, Daniel Sungho Jung, Gyeongsik Moon, and Kyoung Mu Lee. Joint reconstruction of 3D human and object via contact-based refinement transformer. In *CVPR*, 2024. 1, 3

[51] Hyeongjin Nam, Donghwan Kim, Gyeongsik Moon, and Kyoung Mu Lee. PARTE: Part-guided texturing for 3D human reconstruction from a single image. In *ICCV*, 2025. 3

[52] Panwang Pan, Zhuo Su, Chenguo Lin, Zhen Fan, Yongjie Zhang, Zeming Li, Tingting Shen, Yadong Mu, and Yebin Liu. HumanSplat: Generalizable single-image human gaussian splatting with structure priors. In *NeurIPS*, 2024.

[53] Panwang Pan, Zhuo Su, Chenguo Lin, Zhen Fan, Yongjie Zhang, Zeming Li, Tingting Shen, Yadong Mu, and Yebin Liu. HumanSplat: Generalizable single-image human gaussian splatting with structure priors. *NeurIPS*, 2025. 3

[54] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. 2017. 3

[55] Priyanka Patel and Michael J. Black. CameraHMR: Aligning people with perspective. In *3DV*, 2025. 3

[56] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019. 3, 4, 5, 6

[57] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. DreamFusion: Text-to-3D using 2D diffusion. In *ICLR*, 2023. 5

[58] Guocheng Qian, Jinjie Mai, Abdullah Hamdi, Jian Ren, Aliaksandr Siarohin, Bing Li, Hsin-Ying Lee, Ivan Skorokhodov, Peter Wonka, Sergey Tulyakov, et al. Magic123: One image to high-quality 3D object generation using both 2D and 3D diffusion priors. In *ICLR*, 2024. 3

[59] Lingteng Qiu, Xiaodong Gu, Peihao Li, Qi Zuo, Weichao Shen, Junfei Zhang, Kejie Qiu, Weihao Yuan, Guanying Chen, Zilong Dong, et al. LHM: Large animatable human reconstruction model from a single image in seconds. In *ICCV*, 2025. 3, 4

[60] Lingteng Qiu, Shenhao Zhu, Qi Zuo, Xiaodong Gu, Yuan Dong, Junfei Zhang, Chao Xu, Zhe Li, Weihao Yuan, Liefeng Bo, et al. AniGS: Animatable gaussian avatar from a single image with inconsistent gaussian reconstruction. In *CVPR*, 2025. 3

[61] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 6

[62] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 5, 4

[63] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. DreamBooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 3

[64] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. PIFuHD: Multi-level pixel-aligned implicit function for high-resolution 3D human digitization. In *CVPR*, 2020. 3

[65] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: A single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. 3

[66] Soshi Shimada, Vladislav Golyanik, Zhi Li, Patrick Pérez, Weipeng Xu, and Christian Theobalt. HULC: 3D human motion capture with pose manifold sampling and dense contact guidance. In *ECCV*, 2022. 3

[67] Geonhee Sim and Gyeongsik Moon. PERSONA: Personalized whole-body 3D avatar with pose-driven deformations from a single image. In *ICCV*, 2025. 3

[68] Jiaxiang Tang, Zhaoxi Chen, Xiaokang Chen, Tengfei Wang, Gang Zeng, and Ziwei Liu. LGM: Large multi-view gaussian model for high-resolution 3D content creation. In *ECCV*, 2024. 3

[69] Jiaxiang Tang, Jiawei Ren, Hang Zhou, Ziwei Liu, and Gang Zeng. DreamGaussian: Generative gaussian splatting for efficient 3D content creation. In *ICLR*, 2024. 3

[70] Shashank Tripathi, Agniv Chatterjee, Jean-Claude Passy, Hongwei Yi, Dimitrios Tzionas, and Michael J Black. DECO: Dense estimation of 3D human-scene contact in the wild. In *ICCV*, 2023. 3, 2

[71] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 3

[72] Yufu Wang, Yu Sun, Priyanka Patel, Kostas Daniilidis, Michael J Black, and Muhammed Kocabas. PromptHMR: Promptable human mesh recovery. In *CVPR*, 2025. 3

[73] Zhenrong Wang, Qi Zheng, Sihan Ma, Maosheng Ye, Yibing Zhan, and Dongjiang Li. End-to-end HOI reconstruction Transformer with graph-based encoding. In *CVPR*, 2025. 1

[74] Boran Wen, Dingbang Huang, Zichen Zhang, Jiahong Zhou, Jianbin Deng, Jingyu Gong, Yulong Chen, Lizhuang Ma, and Yong-Lu Li. Reconstructing in-the-wild open-vocabulary human-object interactions. In *CVPR*, 2025. 3, 5, 6, 7, 8, 1, 4

[75] Jianfeng Xiang, Zelong Lv, Sicheng Xu, Yu Deng, Ruicheng Wang, Bowen Zhang, Dong Chen, Xin Tong, and Jiaolong Yang. Structured 3d latents for scalable and versatile 3d generation. In *CVPR*, 2025. 3

[76] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. CHORE: Contact, human and object reconstruction from a single RGB image. In *ECCV*, 2022. 1, 3

[77] Xianghui Xie, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Template free reconstruction of human-object interaction with procedural interaction generation. In *CVPR*, 2024. 1

[78] Xianghui Xie, Xi Wang, Nikos Athanasiou, Bharat Lal Bhatnagar, Chun-Hao P Huang, Kaichun Mo, Hao Chen, Xia Jia, Zerui Zhang, Liangxian Cui, et al. RHOBIN Challenge: Reconstruction of human object interaction. *arXiv preprint arXiv:2401.04143*, 2024. 3

[79] Yuliang Xiu, Jinlong Yang, Xu Cao, Dimitrios Tzionas, and Michael J Black. ECON: Explicit clothed humans optimized via normal integration. In *CVPR*, 2023. 3

[80] Jiale Xu, Weihao Cheng, Yiming Gao, Xintao Wang, Shenghua Gao, and Ying Shan. InstantMesh: Efficient 3D mesh generation from a single image with sparse-view large reconstruction models. *arXiv preprint arXiv:2404.07191*, 2024. 3, 4, 5, 7

[81] Xiang Xu, Hanbyul Joo, Greg Mori, and Manolis Savva. D3D-HOI: Dynamic 3D human-object interactions from videos. *arXiv preprint arXiv:2108.08420*, 2021. 1, 3

[82] Yuhang Yang, Wei Zhai, Hongchen Luo, Yang Cao, and Zheng-Jun Zha. LEMON: Learning 3D human-object interaction relation from 2D images. In *CVPR*, 2024. 3, 7, 8, 1, 2

[83] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-Splatting: Alias-free 3D gaussian splatting. In *CVPR*, 2024. 4

[84] Hongwen Zhang, Yating Tian, Xinch 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*, 2021. 3

[85] 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 *ECCV*, 2020. 3, 7, 8, 1, 5, 6

[86] Zibo Zhao, Wen Liu, Xin Chen, Xianfang Zeng, Rui Wang, Pei Cheng, Bin Fu, Tao Chen, Gang Yu, and Shenghua Gao. Michelangelo: Conditional 3D shape generation based on shape-image-text aligned latent representation. In *NeurIPS*, 2023. 3

[87] Zerong Zheng, Tao Yu, Yebin Liu, and Qionghai Dai. PaMIR: Parametric model-conditioned implicit representation for image-based human reconstruction. *TPAMI*, 2021. 3