

Hoi3DGen: Generating High-Quality Human-Object-Interactions in 3D

Agniv Sharma^{1,4,†} Xianghui Xie^{1,2,3,†} Tom Fischer⁴ Eddy Ilg⁴ Gerard Pons-Moll^{1,2,3}

¹University of Tübingen ²Tübingen AI Center ³Max Planck Institute for Informatics ⁴Technische Universität Nürnberg

<https://virtualhumans.mpi-inf.mpg.de/hoi3dgen/>




Figure 1. Given detailed text descriptions of human, object and their interactions, Hoi3DGen generates high quality textured human and object meshes that follow precisely the contact semantics, together with an aligned animatable SMPL model.

Abstract

Modeling and generating 3D human-object interactions from text is crucial for applications in AR, XR, and gaming. Existing approaches often rely on score distillation from text-to-image models, but their results suffer from the Janus problem and do not follow text prompts faithfully due to the scarcity of high-quality interaction data. We introduce Hoi3DGen, a framework that generates high-quality

textured meshes of human-object interaction that follow the input interaction descriptions precisely. We first curate realistic and high-quality interaction data leveraging multi-modal large language models, and then create a full text-to-3D pipeline, which achieves orders-of-magnitude improvements in interaction fidelity. Our method surpasses baselines by 4–15× in text consistency and 3–7× in 3D model quality, exhibiting strong generalization to diverse categories and interaction types, while maintaining high-quality 3D generation.

[†]Equal contribution.

1. Introduction

Modeling human-object-interaction is highly important for gaming, virtual and augmented reality. Generating such interactions as 3D models from text prompts is particularly interesting, since the manual creation of interacting humans and objects is laborious. Despite its significance, this problem has been overlooked, and most methods focus on either human-only [5, 23] or object-only [32, 56] generation. Only a few methods [8, 11, 77, 80] have pioneered this field with score distillation sampling (SDS). While showing promising results, SDS generations are unreliable at inference and often lead to unnatural poses and low-quality 3D interactions with the Janus problem.

A major challenge for text-to-3D interaction generation is the lack of paired text captions and 3D interaction data. Existing 3D interaction datasets [3, 30, 49, 60] feature interactions with natural poses and high-quality contacts, but typically cover only a limited set of object categories or lack fine-grained textual descriptions. Consequently, current approaches often adopt training-free SDS-based pipelines that leverage powerful image diffusion models [11, 77, 80]. While the underlying models were trained on billions of images, these models do not generalize well for interactions and produce implausible results.

In this paper, we introduce **Hoi3DGen**, a framework that generates high-quality 3D-human-object-interaction, while also maintaining strong generalization ability. Our key idea is to automatically create high-quality text descriptions for 3D interactions and fine-tune a model to generate the compositional interaction without losing its own capability for generating diverse humans and objects. Specifically, we propose an automatic interaction labelling pipeline that leverages multimodal large language models to generate high-quality and detailed descriptions for 3D human object interaction. We decompose the complex interaction captioning task into simpler subtasks that describe appearance, action, and body parts in contact, respectively, and fuse the results together to generate the final caption. Based on automatically generated data, we design a text-to-3D generation method that equips the text-to-image model with view conditioning, lifts the 2D image to a 3D mesh, and segments and registers the 3D mesh with the SMPL model to obtain semantic contacts. In summary, our main contributions are:

- We present an automatic data annotation pipeline that generates high-quality detailed text captions for 3D-human-object-interactions by decomposing the complex caption generation into subtasks, solvable by open-source multimodal large language models.
- We introduce a text-to-3D pipeline that generates high-quality interacting segmented human-object meshes with an aligned animatable SMPL model.
- We demonstrate that Hoi3DGen surpasses baselines by 4–15× in text to 3D consistency and 3–7× in 3D model

quality.

Our code, data, and pretrained models will be released.

2. Related work

Text-to-3D generation. For general objects, recent methods can be classified as Score Distillation Sampling (SDS)-based or learning-based. SDS-based approaches like DreamFusion [42], ProlificDreamer [53] and more [7, 33, 39, 50] distill 3D objects from pre-trained 2D image diffusion models [45]. Learning-based approaches either fine-tune image diffusion models to generate novel views conditioned on camera poses [25, 47, 67, 68], or adopt native 3D representations like Triplane [6, 21], SDF [51], occupancy [78], or hybrid ones [56] to directly obtain 3D. Trained on large-scale datasets [12, 13], these methods can generate high-quality 3D from text. Orthogonal to these, human avatar creation relies more heavily on SDS due to the lack of high-quality and diverse human data [5, 23, 27, 35, 52, 75]. Recent works [69, 81] propose to distill 2D generation models to create large-scale human data. While showing promising results, these methods consider only humans or only objects and cannot model the complex relationship during interaction.

Human-object interaction. Most existing methods focus on the accurate capture of hand-object [20, 72], full-body-object [58, 62] or human-scene [19, 70, 73] interactions from images [60] or videos [9, 16, 59, 61, 63]. Based on the captured interaction motion data [3, 30, 49], several works learn to generate HOI motion sequences conditioned on signals like human or object position [30, 54], past motion sequence [64], and text [15, 31, 41, 55, 65]. Despite impressive performance, these approaches are typically limited to non-clothed SMPL body meshes and object templates from the training datasets, restricting generalization. Another line of work synthesizes interactions from provided human and object meshes and optimizes human poses via SDS [77, 80], with [77] further improving contacts using open-set affordances and LLMs. However, they require user-supplied meshes and cannot generate fully-textured HOIs from text. The most related work [11], employs SDS to optimize separate human and object NeRFs and fuses them into a combined scene, but suffers from the Janus problem, noisy textures, and severe interpenetrations, making segmentation difficult and lacking contact control. In contrast, our method produces realistic human-object interactions with high-quality textures, accurate contacts, and also enables accurate human-object segmentation.

3D datasets with text annotation. With the great progress in large language models (LLM), automated text annotation for various 2D [46, 69] and 3D data [29, 79] has become a viable option. For objects, CAP3D [37] proposed a scalable pipeline using vision language models and annotated the Objaverse dataset [13]. Follow-up work [38] further im-

proves the accuracy by robustly selecting better views for annotation. These works output only one description of the object, while Marvel-40M+ [48] introduces multi-level annotations, which allow control at different complexity levels. Orthogonal to these, PoseScript [14] develops a procedural pipeline to assign text labels based on axis angles and further train a model for automatic human-pose-to-text annotation. However, these methods handle objects and humans separately. Some works [41, 54, 66] annotated interaction motion with text manually, which is not scalable. Another work [71] utilizes GPT-4V for detailed HOI labeling, but its dependence on proprietary models introduces cost constraints, and its lengthy descriptions deviate from natural human phrasing. In contrast, we propose a fully automatic and scalable pipeline based on open-source models to annotate complex interactions through natural language.

3. Method

We present Hoi3DGen, a framework to obtain high-quality 3D models of humans interacting with objects. An overview of our method is provided in Fig. 2.

First, we introduce an automatic data annotation pipeline that produces detailed text captions for 3D-human-object-interaction (Sec. 3.1). We then adopt a text-to-image and 2D to 3D lifting paradigm for 3D generation. We propose a view-conditioned image generator that synthesizes interaction images from text (Sec. 3.2) and lifts them into high-quality textured 3D meshes, which are then semantically segmented into human and object components, and aligned with the SMPL model for consistent semantics and animation (Sec. 3.3).

3.1. Data Curation

Given a 3D human-object-interaction mesh, we aim to automatically generate captions that describe appearance, actions, and contacts. For scalable automation, we decompose the task into subtasks that multimodal LLMs can solve efficiently and reliably. As is common in many interaction datasets [3, 60, 76] and also the case in ProciGen, we assume that the 3D interaction is represented as a textured SMPL [36] model for the human and a separate textured mesh for the object, see Fig. 2 top left. We then divide the interaction annotation into 1.) appearance labeling, 2.) interaction labeling and 3.) caption generation.

Appearance Labeling. To label humans and object, our method begins by rendering them separately into four orthogonal views: front, back, and either side. We then use these views as queries for the multimodal LLM InternVL [10] and prompt it to describe the attributes shown in the image. For human labeling, we focus on clothing, hairstyle, and footwear. For object labeling, we annotate attributes such as color, texture, and overall structural form.

Interaction Labeling. A detailed interaction description should define the *what* and *how* of the interaction. We query the InternVL with the four views of human-object interaction renderings and ask it to describe the type of interaction by picking a label from a predefined list of possible actions of the shown object category. Alternatively, one could let InternVL propose actions, but we found that this leads to poor annotation quality due to hallucinations or ambiguous wording for semantically similar actions. Addressing the *how* is also crucial, since the same action can produce very different interactions depending on the contact points (i.e. *holding* a basket in the *right* or *left* hand). To integrate this information, we analyse contact points of the SMPL mesh with the object and filter out body parts whose distances are smaller than 4cm as the contacting parts.

Combined Caption Generation. To generate the final caption, we integrate human, object, action, and interaction labels, along with the object category, using LLaMA 3.1 (70B) [18]. The strong generative capability of the model helps us in obtaining high-quality, natural, and detailed text captions.

Filtering. We apply our data annotation method to the entire ProciGen dataset [60], yielding over 750k pairs of 3D models and captions. Although ProciGen features large object shape diversity, the human appearances and interaction types are limited (100 subjects and 18 categories, respectively). More importantly, many ProciGen interactions involve multiple simultaneous contacts, which leads to a less rich training signal when finetuning, since the model has to distinguish contacts such as left and right part labels. To avoid forgetting during the finetuning, we aim to minimize the training iterations and therefore curate a small subset of high-quality and diverse samples.

We separate interactions according to $K = 8$ frequent but distinct contact configurations: on back, right hand, left hand, right leg, left leg, both hands, no contact, and others. The goal is to construct disjoint subsets $\mathcal{D}^k = \{x_i^k\}$, where each sample $x_i^k = (t, \mathcal{I})$ consists of text description t and renderings \mathcal{I} that *only* have contact configuration k . For instance, *right hand* includes only those with contact limited to the right hand and forearm. Then, we filter each subset by removing data where the object is either significantly overlapping with the human or is far away from both the human and the ground. We then group the remaining data based on action-contact pairs and remove samples where the action clearly mismatches the contact (i.e. a sample with contact point *right hand* and action *kicking*). Finally, we arbitrarily select 50 samples from each subset \mathcal{D}^k , leading to 400 remaining high-quality 3D HOI instances.

3.2. View-Conditioned 2D Interaction Generation

For text-to-image generation, we build on top of the latent diffusion model SANA [57], which can already gen-

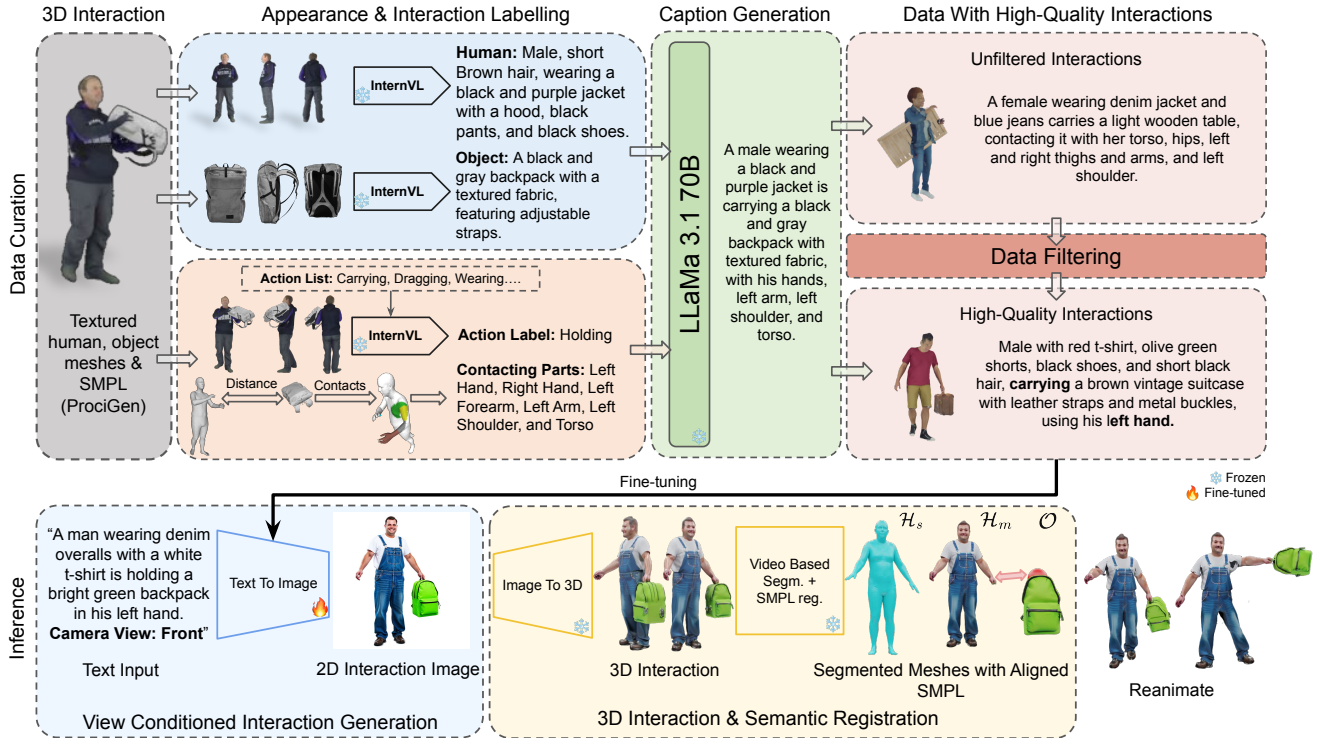


Figure 2. **HOI3D framework overview.** **Top:** We first leverage the existing multimodal foundation model InternVL [10] to perform decomposed annotation of human, object, and human-object-interaction of samples from the ProciGen [60] dataset. We then use LLaMa [18] to create a final detailed caption for the sample. **Bottom:** We leverage our data consisting of high-quality and diverse human-object-interactions to fine-tune an existing text-to-image model. Subsequently, we establish a pipeline to reconstruct high-fidelity textured 3D meshes. The output of our final text-to-3D inference pipeline consists of segmented meshes for the human and object, as well as an animatable SMPL model.

erate high quality human-only or object-only images but cannot follow precisely the interaction prompts. A key discovery of our work is that fine-tuning on the 400 high-quality and diverse interaction samples is sufficient to adjust the learned representation, enabling the model to generate high-quality human-object interactions while maintaining its original power of generating diverse humans and objects.

Another key innovation to enable lifting to high-quality 3D models is to introduce a conditioning of the text-to-image model on the camera view angle. Specifically, we append a view description t_v to the interaction prompt t to generate the 2D interaction image I using our latent diffusion model ϵ_θ . We allow control of three distinct view angles $t_v \in \{\text{front, left diagonal, right diagonal}\}$ that avoid occlusion for most interaction types. These correspond to rendering 3D interaction meshes from azimuth angles of $0^\circ, -45^\circ, +45^\circ$ respectively. We then finetune the model using the standard diffusion loss [22]:

$$\mathcal{L} = \mathbb{E}_{x_0, \epsilon, \sigma} \left[\|\epsilon - \epsilon_\theta(z_t, \sigma, \text{cond})\|_2^2 \right], \quad (1)$$

where $\epsilon \sim \mathcal{N}(0, I)$, $\sigma \sim \mathcal{U}(0, 1)$, and $\text{cond} = (t, t_v)$. Fi-

nally, since training on the highly correlated finetuning data will inadvertently bias the model and reduce output variety, we improve overall texture quality and fidelity in the generated output by retexturing with the Flux model [4].

We show in Tab. 4 that our view-conditioned sampling produces more accurate 3D contacts and our fine-tuning significantly improves the existing text-to-image model for interaction image generation in Tab. 2.

3.3. 3D Interaction and Semantic Registration

3D Interaction Generation. Given the high-quality image of our finetuned diffusion model, we can obtain a 3D mesh using a large-scale image-to-3D model. In our experiments, we use Hunyuan3D [51] due to the observation that it can generate diverse shapes and interactions if the 2D image is of high quality. Our view control allows us to sample three images of the same interaction type, where at least one view has the full interaction visible from which Hunyuan3D can easily generate high quality and complete 3D mesh. Hence, we directly apply Hunyuan3D to all generated images and obtain three different textured 3D inter-

actions. The user can either pick the best one manually or we perform automatic selection based on segmentation and contact semantics, which we describe next.

Interaction Segmentation. Despite high quality, the 3D generation from Hunyuan3D is a single combined mesh in a normalized space. For practical applications, one also needs to understand the semantics, such as which part the human is and where contacts happen. To this end, we propose to segment the mesh and register a SMPL [36] body model that provides semantics and allows further animation.

Given the combined 3D mesh \mathcal{M} produced by Hunyuan3D, our objective is to separate it into two semantically meaningful components: the human mesh \mathcal{H}_m and the object mesh \mathcal{O} . To this end, we render \mathcal{M} into a video sequence along a smooth camera trajectory that spans elevations in $[-60^\circ, 60^\circ]$ and a full 360° azimuth. We then apply the open-vocabulary video segmentation model Grounded-Segment Anything 2 (GSAM2) [26, 34, 44] to obtain temporally consistent binary mask sequences $\{\mathbf{M}_i^h, \mathbf{M}_i^o\}_{i=1}^N$ of human and object, for each rendered view i . We input *person* and the target object category to prompt GSAM2 for human and object segmentation respectively.

We aggregate the 2D masks to segment 3D mesh vertices based on vertex visibility and majority voting. For each rendered view i , we have access to the object mask \mathbf{M}_i^o of GSAM2 and the depth map \mathbf{D}_i and camera parameters of the rendering process. A vertex \mathbf{v} is considered *visible* in view i if it projects inside the image and passes a z-buffer consistency check:

$$\text{vis}_i(\mathbf{v}) = 1[-\delta \leq z_i(\mathbf{v}) - \mathbf{D}_i(\pi_i(\mathbf{v})) \leq \delta], \quad (2)$$

where δ is a depth tolerance, $\pi_i(\cdot)$ is the projection function of frame i , and $z_i(\mathbf{v})$ denotes the projected depth of \mathbf{v} to view i . For the set of views where \mathbf{v} is visible, we compute the fraction in which the vertex lies inside the object mask. Let $\mathcal{V}(\mathbf{v})$ denote the set of visible views of \mathbf{v} . A vertex is then assigned the binary object label if this fraction exceeds a certain threshold τ (empirically, 0.5):

$$l(\mathbf{v}) = \begin{cases} 1, & \text{if } \frac{1}{|\mathcal{V}(\mathbf{v})|} \sum_i \mathbf{M}_i^o[\pi_i(\mathbf{v})] > \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

Finally, the mesh \mathcal{M} is split into the human and object components ($\mathcal{H}_m, \mathcal{O}$) according to these per-vertex labels.

SMPL Registration. To obtain semantics for the generated interaction, we register a SMPL [36] body model \mathcal{H}_s to our segmented human mesh \mathcal{H}_m . Since the human mesh is often incomplete, off-the-shelf human registration methods do not work well for our setup, since they require complete scans [60] or were trained on limited poses [28]. To solve this issue, we introduce a simple yet effective approach to obtain an aligned SMPL mesh. First, we apply CameraHMR [40] to a arbitrary rendering with zero elevation to



Figure 3. **Analysis of the CLIP score.** While our model clearly generates images that follow input interaction descriptions more precisely than SANA [57], the CLIP score indicates the opposite, rendering it unusable as a metric for our task.

identify the front view of the mesh. Then, we render the mesh from the front to get a higher quality SMPL by reapplying CameraHMR. The SMPL model is then scaled and translated to match the center and scale of our mesh output. Finally, scale, global translation, and rotation (7DoF) are refined with a few rounds of Chamfer distance-based optimization.

4. Experiment

4.1. Experimental Setup

Implementation Details. For our text-to-image pipeline, we fine-tune the pretrained *Sana* model at a resolution of 1024×1024 . The fine-tuning process is carried out on four H100 GPUs over 24 hours, with each GPU handling a batch size of 4, yielding an effective batch size of 16. All inferences are performed on a single A100 GPU.

Evaluation Metrics. We evaluate model performance in two aspects: consistency to input text prompt and quality of the generated 3D interactions. For text consistency, we report GPT score, CLIP score, and contact accuracy. The **GPT score** is defined in the same way as in InterFusion [11], where GPT-4V is prompted to select the one that is most consistent with the input text among the generations from different methods. Each 3D generation is rendered into 4 orthogonal views as the input to GPT-4V. **CLIP score** averages CLIP similarity [43] across eight rendered views. While we report this standard metric for completeness, we

find that CLIP is poorly suited for fine-grained interaction recognition as we show in Fig. 3. We also report **contact accuracy** to evaluate if our 3D interactions faithfully follow the contacts defined in input prompts. Using our SMPL registration to the 3D interaction, we can segment out the body parts that are mentioned in the prompt and consider it in contact if its minimum distance to the object is smaller than 4cm, which is the same as Section 3.1. We then calculate the percentage of contacts that correctly follow the contacts defined in prompts as the contact accuracy.

For 3D quality, we also instruct GPT-4V to select the best one among generations from different methods. In this case, however, we only prompt the model to consider the visual quality alone without providing the prompt to generate the 3D. We additionally conduct user studies to evaluate the text consistency and 3D fidelity, detailed later.

Interaction Prompts. We ask ChatGPT to generate 100 prompts describing humans and objects in various interaction scenarios and use these prompts as input to produce our 3D interactions. These prompts cover the most general interactions and are used to report the GPT score in text consistency, quality, CLIP score, and the user study. For contact accuracy, we use ChatGPT to generate 60 prompts focusing specifically on the most important body parts for interaction: hands and feet.

4.2. Text-to-3D Interaction Generation

Our Hoi3DGen allows generation of 3D human-object interaction (HOI) with contacts controlled precisely by text descriptions. We compare our method against TRELIS [56], a state of the art method for text to general 3D object synthesis, and InterFusion [11], current state-of-the-art method for text-to-3D textured HOI generation.

Quantitative evaluation. We report the text consistency and 3D quality metrics in Table 1. It can be seen that our model significantly outperforms baselines in GPT scores. TRELIS generates 3D interaction as a full mesh, hence it cannot reason about the semantic contacts between human and object. InterFusion generates the interaction conditioned on a SMPL mesh, yet the interaction is not aligned with the original human mesh, hence it also cannot reason about the contacts. Our model generates 3D interaction together with a registered human body model, and the contact follows precisely the input prompts with an accuracy of 90%.

As shown in Table 1, to obtain most faithful assessments, we also conduct a user study to assess text consistency and 3D quality on 40 randomly selected examples, 20 for each criterion. Participants view 360° videos of TRELIS, InterFusion, and our results in random order. Among 33 participants, our method is preferred by 91.09% for text consistency and 85.56% for 3D quality, significantly outperform-

Method	Text Consistency				3D Quality	
	GPT ↑	CLIP ↑	Contact ↑	User ↑	GPT ↑	User ↑
TRELIS	0.04	0.32	N/A	3.44%	0.21	10.16%
InterFusion	0.15	0.35	N/A	5.47%	0.00	3.28%
Ours	0.81	0.42	90%	91.09%	0.79	85.56%

Table 1. **Quantitative comparison with text-to-3D models.** We compare our method against TRELIS [56], a general text-to-3D object generation model, and InterFusion [11], a text-to-3D interaction model. Our method outperforms all prior arts in both consistency to input text and quality of the generated 3D interactions by a very large margin.

ing other methods. We provide the details on the user study in the Supplemental material.

Notably, InterFusion achieves better text consistency than TRELIS, yet its 3D quality is worse than direct 3D generation from TRELIS. This is attributed to its reliance on score distillation sampling which suffers from low-resolution results and the Janus problem. Our method generates 3D HOI directly, while also maintaining high quality.

Qualitative comparison. In Fig. 4, we present representative examples comparing our method with baselines. As shown in rows 1 and 3, Interfusion suffers from Janus artifacts, resulting in multiple hands and missing faces. Even when the generated results appear plausible (row 2), the overall quality is low and the contacts are incorrect.


TRELIS, trained natively with 3D data, generates much higher-quality 3D results. However, it is not interaction-aware. The model either generates partial objects (Fig. 4, row 2) or completely omits object generation (Fig. 4, rows 1 and 3).

In contrast, after guiding Sana with our annotated high-quality interaction data, our model is able to precisely follow the text prompt and generate accurate interactions with coherent contacts. Note also how well our model generalizes to different characters, clothes, and hairstyles, while the training data ProciGen contains only 100 human subjects. This clearly shows the advantage of our high-quality data and the great potential of existing models: interaction capability is there, we just need to distill them from structured data.


4.3. 2D Interaction Generation

To obtain high-quality 3D generation with correct contacts, we fine-tune a 2D image generation model on our curated dataset. We evaluate this by comparing our model against the pretrained baseline using the CLIP score, the GPT score for text consistency, the GPT score for quality, and for contact accuracy. CLIP and GPT scores are computed directly on generated 2D images, while contact accuracy is mea-

“A short man with a thick beard, dressed in overalls and a white undershirt, carries a wooden box painted light blue with **both hands**.”



“An elderly man with glasses, wearing a tan blazer and slacks, carries a black leather briefcase-style toolbox in **right hand**.”



“Spiderman holding a blue backpack in his **left hand**.”




Figure 4. **Qualitative comparison for text to 3D generation.** InterFusion [11] is based on Score Distillation Sampling and hence is slow and produces low-quality 3D due to the well-known Janus problem. TRELLIS [56] is a learning based native 3D generation method, hence it can produce better 3D but is not interaction-aware. Our method faithfully follows the text prompts, especially the detailed body contact specifications. Our contacts are highlighted with spheres coloured based on contacting body parts.

sured by applying our 3D interaction generation and segmentation pipeline to extract 3D HOI and human registrations from the generated images. As shown in Table 2, our method substantially improves contact accuracy, reflecting enhanced interaction awareness in the 2D generation model. Our model also achieves significantly better GPT scores, demonstrating more faithful adherence to input prompts.

In Contact score, our model achieves an accuracy of 90% in comparison to 45.76% of the base model. Notably, removing prompts involving *Right Hand* and *Both Feet* causes the baseline score to drop to 23.07%, whereas our method only decreases slightly to 87.5%, indicating the base model’s bias toward a limited set of contact configurations. The detailed results are provided in the Supplemental

material.

Although our model attains a slightly lower CLIP score than the pretrained model, we attribute this to CLIP’s limited sensitivity to fine-grained physical interactions [24, 74], as illustrated in Fig. 3. In contrast, the gains in contact accuracy and text consistency reflect stronger interaction understanding beyond what CLIP can capture.

4.4. Ablation Studies

We ablate different design choices of our method.

High-Quality Interaction Data. A critical step in our data curation pipeline is filtering undesirable data with issues such as interpenetration, implausible actions, and

Method	Text Consistency			Generation Quality
	GPT \uparrow	CLIP \uparrow	Contact \uparrow	GPT \uparrow
SANA	0.31	0.42	45.76%	0.24
Ours	0.69	0.40	90%	0.76

Table 2. **Quantitative comparison with 2D baselines.** We compare our method against pretrained text-to-2D generation SANA [57] model. Our method achieves higher text consistency and 3D interaction quality. The CLIP score is however not very informative due to its limited sensitivity to fine-grained text [24, 74] as we also show in Fig. 3.

mismatched action-contact pairs. We compare our model against one trained on the entire ProciGen dataset in Tab. 3b. While the full ProciGen dataset provides diverse interaction examples, the aforementioned issues cause the model to learn incorrect and implausible actions. Consequently, our model trained on filtered data points achieves higher contact accuracy and GPT scores.

Retexturing. We also observe that removing retexturing causes a small drop in contact scores, but significantly reduces GPT-Score, as can be seen in Tab. 3c. This indicates that retexturing is essential for better texture quality.

Method	Text Consistency		
	GPT \uparrow	CLIP \uparrow	Contact accuracy \uparrow
a. Ours w/o data filter	0.20	0.413	0.80
b. Ours w/o retexturing	0.05	0.412	0.85
c. Ours	0.75	0.417	0.90

Table 3. **Ablation studies.** Our proposed data filtering improves contact accuracy and retexturing improves consistency to text input (GPT score). Combining both achieves the best result.

View conditioned sampling. To produce the final 3D model with correct contacts, we rely on a pretrained image-to-3D model. This image to 3D lifting is mostly accurate when both the human and object are visible. To make the 3D generation more stable, we propose to condition the 2D image sampling on a text that specifies which view to generate. To evaluate this, we compute the contact accuracy of generating one, two or three views, respectively. When more than one views are generated, the contact accuracy is simply the maximum of the generated views. We send the ground-truth image renderings from ProciGen and also the images generated by our view conditioned model to Hunyuan3D and report the contact accuracies in Tab. 4. It can be seen that Hunyuan3D can preserve the contacts over 79% of the time when only one view is given. However, it faithfully preserves the contacts with more than 94% accuracy when we sample three times with different views. This highlights the importance of our view conditioned sampling.

Method	3 views	2 view	1 view
ProciGen [60] images	94.7%	89.5%	79.6%
Our generated images	90.0%	86.7%	78.3%

Table 4. **Contact accuracy** of the lifted 3D given different number of 2D images. Our view-conditioned sampling allows generating three views conditioned on the view description, leading to more stable 3D results.

Human registration. We propose a simple yet effective pipeline to register SMPL model to the generated interaction mesh. We compare this against ETCH [28], a state-of-the-art human registration method. We compute the one-directional Chamfer distance from the segmented human mesh to the registered SMPL mesh, and the results are: 4.30cm (ETCH) vs. 1.60cm (ours). ETCH was trained on scans with mainly standing poses, hence it cannot handle complex poses like sitting and bending. Our method is training-free and overall more robust to different poses.

5. Limitation and Future Work

Despite that our method can follow detailed text prompts and generate coherent 3D interactions, one limitation that we noticed is that while our model can reason about the contacts very well, it has difficulties with following very complex human pose descriptions. Notably, such human poses can be very ambiguous when described in text. As opposed to body part contacts, it is also difficult to obtain unique examples demonstrating characteristic poses. Therefore, future work could introduce a dedicated model for text-to-human-pose generation to further improve the pose awareness, similar to InterFusion[11] or ChatPose[17].

6. Conclusion

In this paper, we address the challenging problem of generating 3D-human-object-interactions from detailed text prompts by leveraging existing foundation models for curating diverse and high-quality training data, and to establish an accurate text-to-3D pipeline. We show that little data for interaction is enough to adjust the rich representations learned in general text-to-image models, and that with view-conditioning, the text-to-image output can be lifted to high-quality 3D meshes with accurate contact. Our results outperform the baselines by nearly an order of magnitude, exhibiting strong generalization to diverse categories and interaction types.

Acknowledgements. We thank RVH members [2] and CVMP Members [1] for their helpful discussions. This work is funded by the Deutsche Forschungsgemeinschaft (German Research Foundation) - 409792180 (Emmy Noether Programme, project: Real Virtual Humans), and German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039A, and Amazon-MPI science hub. We gratefully ac-

knowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU) under the BayernKI project b266be. BayernKI funding is provided by Bavarian state authorities. Gerard Pons-Moll is a Professor at the University of Tübingen endowed by the Carl Zeiss Foundation, at the Department of Computer Science and a member of the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645.

References

- [1] <https://cvmplab.github.io/>. 8
- [2] <http://virtualhumans.mpi-inf.mpg.de/people.html>. 8
- [3] Bharat Lal Bhatnagar, Xianghui Xie, Ilya Petrov, Cristian Sminchisescu, Christian Theobalt, and Gerard Pons-Moll. Behave: Dataset and method for tracking human object interactions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3
- [4] BFL Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 4
- [5] Yukang Cao, Yan-Pei Cao, Kai Han, Ying Shan, and Kwan-Yee K. Wong. Dreamavatar: Text-and-shape guided 3d human avatar generation via diffusion models. *arXiv preprint arXiv:2304.00916*, 2023. 2
- [6] Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. In *CVPR*, 2022. 2
- [7] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. *arXiv preprint arXiv:2303.13873*, 2023. 2
- [8] Yongwei Chen, Tengfei Wang, Tong Wu, Xingang Pan, Kui Jia, and Ziwei Liu. Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. *arXiv preprint arXiv:2403.12409*, 2024. 2
- [9] Yue Chen, Xingyu Chen, Yuxuan Xue, Anpei Chen, Yuliang Xiu, and Pons-Moll Gerard. Human3r: Everyone everywhere all at once. *arXiv preprint arXiv:2510.06219*, 2025. 2
- [10] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024. 3, 4
- [11] Sisi Dai, Wenhao Li, Haowen Sun, Haibin Huang, Chongyang Ma, Hui Huang, Kai Xu, and Ruizhen Hu. Interfusion: Text-driven generation of 3d human-object interaction. In *ECCV*, 2024. 2, 5, 6, 7, 8
- [12] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 2
- [13] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. *arXiv preprint arXiv:2307.05663*, 2023. 2
- [14] Delmas, Ginger and Weinzaepfel, Philippe and Lucas, Thomas and Moreno-Noguer, Francesc and Rogez, Grégory. PoseScript: 3D Human Poses from Natural Language. In *ECCV*, 2022. 3
- [15] Christian Diller and Angela Dai. Cg-hoi: Contact-guided 3d human-object interaction generation. 2024. 2
- [16] Zicong Fan, Maria Parelli, Maria Eleni Kadoglou, Muhammed Kocabas, Xu Chen, Michael J Black, and Otmar Hilliges. HOLD: Category-agnostic 3d reconstruction of interacting hands and objects from video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 494–504, 2024. 2
- [17] Yao Feng, Jing Lin, Sai Kumar Dwivedi, Yu Sun, Priyanka Patel, and Michael J. Black. Chatpose: Chatting about 3d human pose. In *CVPR*, 2024. 8
- [18] Aaron Grattafiori, Abhimanyu Dubey, and et al. The llama 3 herd of models, 2024. 3, 4
- [19] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *International Conference on Computer Vision*, 2019. 2
- [20] Yana Hasson, Gül Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *CVPR*, 2019. 2
- [21] Zexin He and Tengfei Wang. Openlrm: Open-source large reconstruction models. <https://github.com/3DTopia/OpenLRM>, 2023. 2
- [22] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 4
- [23] Fangzhou Hong, Mingyuan Zhang, Liang Pan, Zhongang Cai, Lei Yang, and Ziwei Liu. Avatarclip: Zero-shot text-driven generation and animation of 3d avatars. *ACM Transactions on Graphics (TOG)*, 41(4):1–19, 2022. 2
- [24] Raphi Kang, Yue Song, Georgia Gkioxari, and Pietro Perona. Is clip ideal? no. can we fix it? yes!, 2025. 7, 8
- [25] Yash Kant, Ziyi Wu, Michael Vasilkovsky, Guocheng Qian, Jian Ren, Riza Alp Guler, Bernard Ghanem, Sergey Tulyakov, Igor Gilitschenski, and Aliaksandr Siarohin. Spad: Spatially aware multiview diffusers, 2024. 2
- [26] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023. 5
- [27] Nikos Kolotouros, Thiemo Alldieck, Andrei Zanfir, Eduard Gabriel Bazavan, Mihai Fieraru, and Cristian Sminchis-

- escu. Dreamhuman: Animatable 3d avatars from text. *arXiv preprint arXiv:2306.09329*, 2023. 2
- [28] Boqian Li, Haiwen Feng, Zeyu Cai, Michael J. Black, and Yuliang Xiu. ETCH: Generalizing Body Fitting to Clothed Humans via Equivariant Tightness. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 5, 8
- [29] Chuqiao Li, Xianghui Xie, Yong Cao, Andreas Geiger, and Gerard Pons-Moll. Frankenmotion: Part-level human motion generation and composition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2
- [30] Jiaman Li, Jiajun Wu, and C Karen Liu. Object motion guided human motion synthesis. *ACM Trans. Graph.*, 42(6), 2023. 2
- [31] Jiaman Li, Alexander Clegg, Roozbeh Mottaghi, Jiajun Wu, Xavier Puig, and C. Karen Liu. Controllable human-object interaction synthesis. In *ECCV*, 2024. 2
- [32] Ming Li, Pan Zhou, Jia-Wei Liu, Jussi Keppo, Min Lin, Shuicheng Yan, and Xiangyu Xu. Instant3d: Instant text-to-3d generation. *Int. J. Comput. Vis.*, 132(10):4456–4472, 2024. 2
- [33] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [34] 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. *arXiv preprint arXiv:2303.05499*, 2023. 5
- [35] Xian Liu, Xiaohang Zhan, Jiayang Tang, Ying Shan, Gang Zeng, Dahua Lin, Xihui Liu, and Ziwei Liu. Humangaussian: Text-driven 3d human generation with gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6646–6657, 2024. 2
- [36] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.*, 34(6):248:1–248:16, 2015. 3, 5
- [37] Tiange Luo, Chris Rockwell, Honglak Lee, and Justin Johnson. Scalable 3d captioning with pretrained models. *arXiv preprint arXiv:2306.07279*, 2023. 2
- [38] Tiange Luo, Justin Johnson, and Honglak Lee. View selection for 3d captioning via diffusion ranking. *arXiv preprint arXiv:2404.07984*, 2024. 2
- [39] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Realfusion: 360° reconstruction of any object from a single image. In *Arxiv*, 2023. 2
- [40] Priyanka Patel and Michael J. Black. CameraHMR: Aligning people with perspective. In *International Conference on 3D Vision (3DV)*, 2025. 5
- [41] Xiaogang Peng, Yiming Xie, Zizhao Wu, Varun Jampani, Deqing Sun, and Huaizu Jiang. Hoi-diff: Text-driven synthesis of 3d human-object interactions using diffusion models. *arXiv preprint arXiv:2312.06553*, 2023. 2, 3
- [42] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 2
- [43] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021. 5
- [44] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. Grounded sam: Assembling open-world models for diverse visual tasks, 2024. 5
- [45] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2
- [46] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: An open large-scale dataset for training next-generation image–text models. In *NeurIPS Datasets and Benchmarks Track*, 2022. The public 2-billion-pair subset is often referred to as “LAION-2D/2B”. 2
- [47] Yichun Shi, Peng Wang, Jianglong Ye, Long Mai, Kejie Li, and Xiao Yang. Mvdream: Multi-view diffusion for 3d generation. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 2
- [48] Sankalp Sinha, Mohammad Sadil Khan, Muhammad Usama, Shino Sam, Didier Stricker, Sk Aziz Ali, and Muhammad Zeshan Afzal. Marvel-40m+: Multi-level visual elaboration for high-fidelity text-to-3d content creation. *arXiv preprint arXiv:2411.17945*, 2024. 3
- [49] Omid Taheri, Nima Ghorbani, Michael J. Black, and Dimitrios Tzionas. GRAB: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision (ECCV)*, 2020. 2
- [50] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. *arXiv preprint arXiv:2303.14184*, 2023. 2
- [51] Tencent Hunyuan3D Team. Hunyuan3d 2.0: Scaling diffusion models for high resolution textured 3d assets generation, 2025. 2, 4
- [52] Jionghao Wang, Yuan Liu, Zhiyang Dou, Zhengming Yu, Yongqing Liang, Cheng Lin, Xin Li, Wenping Wang, Rong Xie, and Li Song. Disentangled clothed avatar generation from text descriptions. *arXiv preprint arXiv:2312.05295*, 2024. 2
- [53] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distilla-

- tion. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. 2
- [54] Qianyang Wu, Ye Shi, Xiaoshui Huang, Jingyi Yu, Lan Xu, and Jingya Wang. Thor: Text to human-object interaction diffusion via relation intervention, 2024. 2, 3
- [55] Zhen Wu, Jiaman Li, Pei Xu, and C. Karen Liu. Human-object interaction from human-level instructions, 2025. 2
- [56] 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. *arXiv preprint arXiv:2412.01506*, 2024. 2, 6, 7
- [57] Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, and Song Han. Sana: Efficient high-resolution image synthesis with linear diffusion transformer, 2024. 3, 5, 8
- [58] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Chore: Contact, human and object reconstruction from a single rgb image. In *European Conference on Computer Vision (ECCV)*. Springer, 2022. 2
- [59] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Visibility aware human-object interaction tracking from single rgb camera. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [60] Xianghui Xie, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Template free reconstruction of human-object interaction with procedural interaction generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 4, 5, 8
- [61] Xianghui Xie, Jan Eric Lenssen, and Gerard Pons-Moll. Intertrack: Tracking human object interaction without object templates. 2024. 2
- [62] Xianghui Xie, Xi Wang, Nikos Athanasiou, Bharat Lal Bhatnagar, Chun-Hao P. Huang, Kaichun Mo, Hao Chen, Xia Jia, Zerui Zhang, Liangxian Cui, Xiao Lin, Bingqiao Qian, Jie Xiao, Wenfei Yang, Hyeongjin Nam, Daniel Sungho Jung, Kihoon Kim, Kyoung Mu Lee, Otmar Hilliges, and Gerard Pons-Moll. RHOBIN Challenge: Reconstruction of human object interaction. *arXiv preprint arXiv:2401.04143*, 2024. 2
- [63] Xianghui Xie, Bowen Wen, Yan Chang, Hesam Rabeti, Jiefeng Li, Ye Yuan, Gerard Pons-Moll, and Stan Birchfield. Cari4d: Category agnostic 4d reconstruction of human-object interaction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2
- [64] Sirui Xu, Zhengyuan Li, Yu-Xiong Wang, and Liang-Yan Gui. Interdiff: Generating 3d human-object interactions with physics-informed diffusion. In *ICCV*, 2023. 2
- [65] Sirui Xu, Ziyin Wang, Yu-Xiong Wang, and Liang-Yan Gui. Interdreamer: Zero-shot text to 3d dynamic human-object interaction. *arXiv preprint arXiv:2403.19652*, 2024. 2
- [66] Sirui Xu, Dongting Li, Yucheng Zhang, Xiyan Xu, Qi Long, Ziyin Wang, Yunzhi Lu, Shuchang Dong, Hezi Jiang, Akshat Gupta, Yu-Xiong Wang, and Liang-Yan Gui. Interact: Advancing large-scale versatile 3d human-object interaction generation. In *CVPR*, 2025. 3
- [67] Yuxuan Xue, Xianghui Xie, Riccardo Marin, and Gerard Pons-Moll. Gen-3diffusion: Realistic image-to-3d generation via 2d & 3d diffusion synergy. *arXiv preprint arXiv:2412.06698*, 2024. 2
- [68] Yuxuan Xue, Xianghui Xie, Riccardo Marin, and Gerard Pons-Moll. Human-3diffusion: Realistic avatar creation via explicit 3d consistent diffusion models. In *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. 2
- [69] Yuxuan Xue, Xianghui Xie, Margaret Kostyrko, and Gerard Pons-Moll. Infinihuman: Infinite 3d human creation with precise control. 2025. 2
- [70] Pradyumna Yalandur-Muralidhar, Yuxuan Xue, Xianghui Xie, Margaret Kostyrko, and Gerard Pons-Moll. Physic: Physically plausible 3d human-scene interaction and contact from a single image. In *ACM SIGGRAPH Asia*, 2025. 2
- [71] Jie Yang, Xuesong Niu, Nan Jiang, Ruimao Zhang, and Huang Siyuan. F-hoi: Toward fine-grained semantic-aligned 3d human-object interactions. *European Conference on Computer Vision*, 2024. 3
- [72] Lixin Yang, Xinyu Zhan, Kailin Li, Wenqiang Xu, Jiefeng Li, and Cewu Lu. CPF: Learning a contact potential field to model the hand-object interaction. In *ICCV*, 2021. 2
- [73] Hongwei Yi, Chun-Hao P. Huang, Dimitrios Tzionas, Muhammed Kocabas, Mohamed Hassan, Siyu Tang, Justus Thies, and Michael J. Black. Human-aware object placement for visual environment reconstruction. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 3959–3970, 2022. 2
- [74] Mert Yuksekgonul, Federico Bianchi, Pratyusha Kalluri, Dan Jurafsky, and James Zou. When and why vision-language models behave like bags-of-words, and what to do about it?, 2023. 7, 8
- [75] Huichao Zhang, Bowen Chen, Hao Yang, Liao Qu, Xu Wang, Li Chen, Chao Long, Feida Zhu, Kang Du, and Min Zheng. Avatarverse: High-quality & stable 3d avatar creation from text and pose. *arXiv preprint arXiv:2308.03610*, 2023. 2
- [76] Juze Zhang, Haimin Luo, Hongdi Yang, Xinru Xu, Qianyang Wu, Ye Shi, Jingyi Yu, Lan Xu, and Jingya Wang. Neuraldome: A neural modeling pipeline on multi-view human-object interactions. In *CVPR*, 2023. 3
- [77] Jinlu Zhang, Yixin Chen, Zan Wang, Jie Yang, Yizhou Wang, and Siyuan Huang. Interactanything: Zero-shot human object interaction synthesis via llm feedback and object affordance parsing, 2025. 2
- [78] Longwen Zhang, Ziyu Wang, Qixuan Zhang, Qiwei Qiu, Anqi Pang, Haoran Jiang, Wei Yang, Lan Xu, and Jingyi Yu. Clay: A controllable large-scale generative model for creating high-quality 3d assets. *ACM Transactions on Graphics (TOG)*, 43(4):1–20, 2024. 2
- [79] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Transactions on Graphics (Proc. SIGGRAPH 2018)*, 37(4):160:1–160:12, 2018. 2

- [80] Thomas Hanwen Zhu, Ruining Li, and Tomas Jakab. DreamHOI: Subject-driven generation of 3d human-object interactions with diffusion priors. *arXiv preprint arXiv:2409.08278*, 2024. [2](#)
- [81] Yiyu Zhuang, Jiayi Lv, Hao Wen, Qing Shuai, Ailing Zeng, Hao Zhu, Shifeng Chen, Yujiu Yang, Xun Cao, and Wei Liu. Idol: Instant photorealistic 3d human creation from a single image. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025. [2](#)