InterDreamer: Zero-Shot Text to 3D Dynamic Human-Object Interaction

Sirui Xu^{1†}, Ziyin Wang^{2†}, Yu-Xiong Wang^{1‡}, and Liang-Yan Gui^{1‡}

¹ University of Illinois at Urbana-Champaign ² Fudan University [†] Equal Contribution [‡] Equal Advising https://sirui-xu.github.io/InterDreamer/

Abstract. Text-conditioned human motion generation has experienced significant advancements with diffusion models trained on extensive motion capture data and corresponding textual annotations. However, extending such success to 3D dynamic human-object interaction (HOI) generation faces notable challenges, primarily due to the lack of large-scale interaction data and comprehensive descriptions that align with these interactions. This paper takes the initiative and showcases the potential of generating human-object interactions without direct training on textinteraction pair data. Our key insight in achieving this is that interaction semantics and dynamics can be decoupled. Being unable to learn interaction semantics through supervised training, we instead leverage pretrained large models, synergizing knowledge from a large language model and a text-to-motion model. While such knowledge offers high-level control over interaction semantics, it cannot grasp the intricacies of lowlevel interaction dynamics. To overcome this issue, we further introduce a world model designed to comprehend simple physics, modeling how human actions influence object motion. By integrating these components, our novel framework, InterDreamer, is able to generate text-aligned 3D HOI sequences in a zero-shot manner. We apply InterDreamer to the BEHAVE and CHAIRS datasets, and our comprehensive experimental analysis demonstrates its capability to generate realistic and coherent interaction sequences that seamlessly align with the text directives.

1 Introduction

Text-guided human motion generation [94] has made unprecedented progress through advancements in diffusion models [31, 85, 86], leading to synthesis outcomes that are more realistic, diverse, and controllable. This progress has further ignited increased interest in exploring expanded tasks related to text-guided human interaction generation, such as social interaction [54] and human-scene interaction [33]. However, many of these explorations are limited in that the dynamics of objects are not involved or cannot be controlled by text. Aiming to bridge such a gap, this paper undertakes the initiative to tackle a more challenging task – generating versatile 3D human-object interactions (HOIs) through language guidance, as illustrated in Fig. 1.



Fig. 1: InterDreamer can generate vivid 3D human-object interaction sequences guided by textual descriptions. Its zero-shot ability is achieved by integrating semantics and dynamics knowledge from large-scale text-motion data (upper left), a large language model (LLM) (bottom left), 3D human-object interaction database (upper middle), and interaction prior (bottom middle). We visualize the generated text-guided interaction sequence (upper right), with the beginning of the sequence unfolded (bottom right). More details are available in https://sirui-xu.github.io/InterDreamer/.

While a direct solution, as suggested by concurrent work [21, 50, 71, 109], would be replicating the success observed in human motion generation and adopting a similar supervised approach for learning text-driven HOIs, it is not scalable. Indeed, even generating social or scene interactions heavily relies on extensive collections of text-interaction pair data [25, 54, 62, 104]. Scaling these methods to address the more complex HOIs outlined in our study could necessitate datasets of comparable magnitude. Achieving this goal appears unattainable by merely annotating existing 3D HOI datasets [7,22,34,36,41,51,129,136], which are relatively limited in size. Although recent studies [21,50,71] have annotated some of these datasets, the volume of text-motion pairs still lags significantly behind that available for existing text-driven motion generation efforts.

An intriguing question naturally emerges: what is the potential of zero-shot learning for text-conditioned HOI generation, which is the main focus of this paper. However, formulating the task in a zero-shot setting presents significant challenges, primarily due to the inability to directly learn the alignment between text and HOI dynamics. Our key observation then is that interaction semantics and dynamics can be decoupled. That is, the high-level semantics of an interaction, aligned with its textual description, can be informed by human motion and the initial object pose. Meanwhile, the low-level dynamics of the interaction – specifically, the subsequent behavior of the object – are governed by the forces exerted by the human, within the constraints of physical laws.

Motivated by these insights, we introduce InterDreamer – a novel framework that synergizes knowledge of interaction semantics and dynamics. As shown in Fig. 1, both of them do not necessarily need to be learned from text-interaction pairs, if they are decoupled, thus leading to the ability for zero-shot generation.

The semantics of interaction, although not available through direct supervised training, can be harnessed from a variety of *prior knowledge* that is independent of text-interaction pair datasets. Specifically, to acquire semantically aligned human motion and initial interaction, we first consult a large language model (LLM), such as GPT-4 [67] and Llama 2 [95], to provide understandings including how humans typically use specific body parts in interactions with particular objects, by exploiting its *in-context learning* capability with *few-shot prompting* [10] and *chain-of-thought prompting* [107]. Both intermediate thoughts and the final thought are then used to (i) generate a semantically aligned human motion with a pre-trained text-to-motion model; and (ii) identify an initial object pose that is harmonious with the generated initial human pose and text description through a retrieval-based algorithm.

While these large models can offer high-level motion semantics modeling, they lack crucial *low-level* dynamics knowledge. Nevertheless, by decoupling interaction dynamics from semantics, a key advantage emerges in our InterDreamer framework: interaction dynamics can be learned from motion capture data *without the necessity of text annotations*. We instantiate this idea by developing a novel *world model*, which predicts the subsequent state of an object affected by the interaction. The key here is to attain *generalizable control signals*. To do so, we exert control over the object through the motion of vertices on the human body. These vertices are solely sampled in regions where contact occurs, *agnostic* to the overall object shape and whole-body motion. Such abstraction empowers the model to learn the fundamental physics from a publicly available 3D HOI dataset [7]. The plausibility of generated interaction is further enhanced by a subsequent optimization procedure on synthesized human and object motion.

To summarize, our contributions are: (i) We initiate a novel task of synthesizing whole-body interactions with dynamic objects guided by textual commands, without access to text-interaction pair data, which is the first to our knowledge. (ii) We introduce a novel framework that decomposes semantics and dynamics, and they can be integrated effortlessly. (iii) Our methodology harnesses knowledge from a large language model (LLM) and a text-to-motion model as external resources, alongside our proposed novel world model. Remarkably, the only component that requires training is the world model, underscoring the ease of use of our framework. Experimental results demonstrate that our zero-shot framework, InterDreamer, is capable of producing semantically aligned and realistic human-object interactions, and generalizes beyond existing HOI datasets.

2 Related Work

Text-Conditioned Human Motion Generation. Significant progress has been witnessed in human motion synthesis tasks, given different kinds of external conditions, including action categories [2, 27, 49, 73], past motions [5, 14, 65, 117, 118, 125], trajectories [38, 39, 80, 97, 113], scene context [11, 29, 33, 91, 99–101, 104, 132, 137, 138], and unconditional generation [76]. Recently, human motion synthesis guided by textual descriptions [1, 6, 15, 20, 25, 26, 42, 45, 59, 60, 74, 75,

4

77, 84, 92, 106, 123, 128, 130, 134, 135, 141] is popular and extended to various applications, including the text-conditioned generation of multiple-person [24, 57, 105] and human-scene interaction [17, 33, 37]. Our goal is to model human and object dynamics concurrently guided by text in a zero-shot setting.

Human-Object Interaction Generation. Synthesizing hand-object interactions [53,124,126,139,140] and single-frame human-object interactions [32,44,72, 102, 112, 127 are popular topics and extended to zero-shot settings [40, 52, 120]. Recently, researchers explore whole-body dynamic interaction generation, in kinematic-based approaches [18, 23, 46–48, 51, 63, 64, 79, 87–89, 98, 109, 110, 116, 119,133] and physics-based approaches [4,8,13,19,30,56,66,68,103,114,115,121]. Current methods in HOI synthesis are often restricted by a narrow scope of actions, the use of non-dynamic objects, and a lack of comprehensive whole-body motion. Our work aims to generate diverse whole-body interactions with various objects, and enables control through language input. Recent datasets [7, 22, 22]34, 36, 41, 51, 90, 129, 136 provide the groundwork for research in this area, and concurrent efforts [21, 50, 71] demonstrate the feasibility of applying supervised learning methods through annotating datasets. However, the amount of data currently available falls short when compared to more extensive text-motion datasets [25, 55, 62]. This discrepancy in data volume limits the capability of supervised methods to capture the complexity of human-object interactions, motivating us to investigate the potential of zero-shot generation.

External Knowledge from LLMs. Large language models (LLMs) are being used for advanced visual tasks, such as editing images based on instructions [9]. In digital humans, they are used to reconstruct 3D human-object interactions [102] and generate human motions [3, 35, 122, 134] as well as humanscene interactions [111]. Our approach is inspired by [102], which uses LLMs to infer contact body parts with a given object for reconstructing 3D humanobject interactions – a task different from ours. Also, our approach utilizes more advanced LLMs, such as GPT-4 [67] or Llama 2 [95], not only to understand contact body parts but also to further narrow the distribution gap between the text descriptions and models in our pipeline. This is achieved by leveraging the in-context learning (ICL) of LLMs with the chain-of-thought prompting [107].

3 Methodology

Problem Formulation. Our goal is to synthesize a sequence of 3D humanobject interactions \boldsymbol{x} that corresponds to a descriptive text p. This sequence is a series of tuples $[(\boldsymbol{h}_1, \boldsymbol{o}_1), (\boldsymbol{h}_2, \boldsymbol{o}_2), \dots, (\boldsymbol{h}_M, \boldsymbol{o}_M)]$, where \boldsymbol{h}_i represents the human pose parameters defined by the SMPL model [58], and \boldsymbol{o}_i defines the object pose in terms of its 3D spatial location and orientation. The sequence length M is variable and is dynamically determined by our text-to-motion model based on the input text p. We do not require text-interaction paired data for training. **Overview.** Our framework, illustrated in Fig. 2, can be conceptualized as a Markov decision process (MDP). We begin by dividing the motion sequence into T segments, each with m frames, where $M = T \times m$. Object motion $\{\boldsymbol{o}_i\}_{i=1}^M$



Fig. 2: An overview of our InterDreamer. (i) Our high-level planning analyzes the description using LLMs and provides guidance to the low-level control. (ii) Our low-level control includes a text-to-motion model that translates text into human actions a_{t+1} , and an interaction retrieval model for extracting the object's initial pose as the first state s_1 . (iii) Our world model executes the actions and outputs the next state s_{t+1} through dynamics modeling. An optimization process is coupled with the dynamics model, projecting the state and action onto valid counterparts s_{t+1}^* and a_{t+1}^* . Solid arrows mean that the process is performed iteratively.

is formulated as a sequence of environmental states $\{s_t\}_{t=1}^T$, and human motion $\{h_i\}_{i=1}^M$ is described as a sequence of actions $\{a_t\}_{t=1}^T$ that interact with the environment. Under such an MDP setup, our pipeline starts with high-level planning L, which, given the textual interaction description p, deciphers the detailed context g = L(p) (Sec. 3.1). Then, a text-to-motion model π translates context g into human actions iteratively, modeled as $a_{t+1} \sim \pi(a_{t+1}|s_t, \{a_i\}_{i=1}^t, g)$ (Sec. 3.2). An interaction retrieval model R proposes an initial object state s_1 , based on the initial action a_1 and context g (Sec. 3.2). After that, a world model P is trained to predict future states s_{t+1} from the current action and state (Sec. 3.3). Our world model further incorporates an optimization process with prior knowledge on human interactions encoded, for both state and action refinement (Sec. 3.4). Notably, the text-to-motion and world models are executed *iteratively* until text-to-motion generates an end frame.

3.1 High-Level Planning

Powered by LLMs' strong reasoning capabilities as well as their common sense, our high-level planning L yields interaction details g = L(p) that cannot be directly extracted in textual descriptions p. The process undertaken by L encompasses three steps: (i) *Determining the object*: The LLM is employed to translate the described objects into corresponding categories from a pre-defined list used in the world model. (ii) Determining initial human-object contact: The LLM infers the body parts involved in the interaction, drawing from a list defined in the SMPL model [58]. (iii) Reducing the distribution gap: The LLM bridges the distribution gap between the freeform textual input and the language used within the text-to-motion model training dataset [25]. This involves standardizing syntax and content according to specific guidelines we design.

In Fig. 2, we demonstrate the prompt used by the LLM, as well as examples, following the few-shot prompting [10]. We define the sequence of intermediate thoughts and the final thought, *i.e.*, the answers to the three questions, as the detailed information g = L(p), which guides the subsequent procedure.

Our high-level planning operates indirectly in the generation of interactions. Nonetheless, it plays a key role in effectively bridging the distribution gap between textual descriptions and subsequent interaction generation. In other words, it narrows the vast range of possible interactions into a more manageable distribution within the capabilities of our framework. As detailed in Sec. 4.1, our high-level planning incorporates GPT-4 [67] or Llama-2 [95] for evaluation.

3.2 Low-Level Control

With the information g derived from the description p, the low-level control aims to create an initial state s_1 and a sequence of human actions $\{a_t\}_{t=1}^T$, such that they correspond with the objectives outlined by g. In this section, we detail the text-to-motion model and interaction retrieval system designed to fulfill the task. **Text-to-Motion.** We construct the text-to-motion model π designed to develop actions to be executed in the world model. At each timestep t, π receives the sequence of previous actions $\{a_i\}_{i=1}^t$ and the text tokens extracted from g = L(p), and produces a preliminary next action a_{t+1} , which will be adjusted to the ultimate action a_{t+1}^* through an optimization process that intertwines it with the object state, introduced in Sec. 3.4. Thus, the overall process coupled with the optimization can be formally defined as $a_{t+1} \sim \pi(a_{t+1}|s_t, \{a_i\}_{i=1}^t, g)$. The initial action $a_1 \sim \pi(a_1|g)$ is influenced by the text from g without prior actions or states, which will be used in the interaction retrieval. As detailed in Sec. 4.1, π is capable of leveraging existing text-to-motion models, including MDM [93], MotionDiffuse [130], ReMoDiffuse [131], and MotionGPT [35].

Interaction Retrieval. The interaction retrieval component R sets up an initial state $s_1 \sim R(a_1, g)$, given the initial action a_1 generated by the text-to-motion model. We present here a user-friendly pipeline that is based on handcraft rules and a learning-based approach in Sec. B.1 of supplementary. First, we build a database by collecting HOI frames from the target dataset, *e.g.*, the BEHAVE [7] or CHAIRS [36] dataset. The indexing key for retrieval is a tuple consisting of the body part in contact and the category of the object involved. The value for retrieval is a per-frame contact map, *i.e.*, a list of K vertex pairs $\{(d_h^i, d_o^i)\}_{i=1}^K$. Here, d_h^i indexes the contact vertex on the human's surface, while d_o^i indexes the corresponding contact vertex on the object's surface. This contact map is linked with their key, thus establishing a searchable record of interactions. During the inference stage, equipped with the body part and object information provided by

the high-level planning (Sec. 3.1), we utilize them as a key to retrieve all relevant contact maps from our database. We sample one map $\{(d_h^i, d_o^i)\}_{i=1}^K$ according to a pre-defined metric, and employ it to ascertain an object state $s_1 \sim R(a_1, g)$, harmonious with the initial human action a_1 , as the initiation of the interaction. More details are provided in Sec. B.1 of supplementary.

3.3 World Model

Our world model combines a dynamics model and the optimization process, dedicated to simulating state transitions affected by applied actions. While drawing inspiration from similar concepts utilized in robotics [83, 108] and autonomous driving systems [43], we use it here for rollout without model-based learning. This model, trained on the 3D HOI dataset, serves a role similar to a physics simulator but is much simpler – it takes the preceding object state s_t along with a pair of consecutive actions a_t and a_{t+1} , and predicts the subsequent object state s_{t+1} . The interplay between the low-level control and the world model ultimately produces a coherent interaction rollout.

In designing the dynamics model, a naïve method would be directly taking the raw actions and past state as its input. This method, however, suffers from a severe generalization problem during inference: the dynamics model is likely to encounter some human actions that do not exist in the training set, since our text-to-motion model is not trained on the HOI dataset. We provide ablation studies as evidence for this claim in Sec. 4. Thus, instead of directly modeling the complex distribution of interactions, we approach the interactions from the contact vertices on the object, as shown in Fig. 2. This locality ensures that the dynamics model remains focused on interactions in the contact region, without being distracted by the motion of body parts that are irrelevant to the object.

Input Representation. Specifically, at each time step t, we abstract the past action as H historical vertex trajectories $\{\{\boldsymbol{v}_i^j\}_{j=1}^N\}_{i=1}^H$, and the future action as F future vertex trajectories $\{\{\boldsymbol{v}_i^j\}_{j=1}^N\}_{i=H+1}^H$, where non-fixed variable N is the number of sampled contact vertices. Note that we train our dynamics model to forecast over a longer duration than the past motion (F > H), while only the foremost future action will be used for autoregressive generation during the inference, to facilitate the long-term generation process, as suggested in [16].

To determine these N vertices, we start with object's signed distance fields $\{\mathbf{sdf}_i\}_{i=1}^H$ over the past H frames, derived from the past state s_t . We then sample vertices that meet the following criteria:

$$|\mathbf{sdf}_i(\boldsymbol{v}_i^j)| \le \delta_1, \quad \forall i = 1, \dots, H, \quad \forall j, \tag{1}$$

$$\|\boldsymbol{v}_i^j - \boldsymbol{v}_i^k\| \ge \delta_2, \quad \forall j \neq k, \tag{2}$$

where δ_1 and δ_2 are two hyperparameters. The objective is to sparsely sample contact vertices, while ensuring they are sufficient to encompass the interaction.

contact vertices, while ensuring they are sufficient to encompass the interaction. We characterize each vertex trajectory $\{v_i^j\}_{i=1}^{H+F}$ with a feature f^j to provide information in addition to motion, which includes (i) vertex coordinates at T-pose, providing information about the position of the human vertex on the body surface; (ii) the vertex-to-object surface distance, indicating vertex's impact on the object; and (iii) the vertex's velocity relative to its nearest object vertex.

Architecture. An overview of the architecture is demonstrated in Fig. 2. An unconditional dynamics block can be initiated as $\mathcal{G}(\boldsymbol{x}_k, \boldsymbol{\Theta})$, mapping the input feature map \boldsymbol{x}_k at the k-th layer to another feature map, with $\boldsymbol{\Theta}$ denoting the dynamics model parameters. To incorporate human vertex controls, we introduce a secondary network $\mathcal{F}(\boldsymbol{y}_k^j, \boldsymbol{\Theta}_v)$ operating on N vertex features $\{\boldsymbol{y}_k^j\}_{j=1}^N$, where $\boldsymbol{\Theta}_v$ is its parameters. With a cross-attention layer Attn, the final dynamics block is formulated as:

$$\boldsymbol{x}_{k+1}, \{\boldsymbol{y}_{k+1}^j\}_{j=1}^N = \operatorname{Attn}(\mathcal{G}(\boldsymbol{x}_k, \boldsymbol{\Theta}), \{\mathcal{F}(\boldsymbol{y}_k^j, \boldsymbol{\Theta}_v)\}_{j=1}^N).$$
(3)

We stack multiple blocks to form the dynamics model. The initial input, \boldsymbol{x}_0 , corresponds to the previous state \boldsymbol{s}_t , while each \boldsymbol{y}_0^j represents the feature of the vertex trajectory, containing both the trajectory $\{\boldsymbol{v}_i^j\}_{i=1}^{H+F}$ and its associated feature vector \boldsymbol{f}^j . The output of this model, \boldsymbol{s}_{t+1} , is preliminary and subject to further optimization as detailed in Sec. 3.4, which will yield the final future state \boldsymbol{s}_{t+1}^* . We utilize the Mean Squared Error loss to train the dynamics model. For more details, please refer to Sec. B.2 of supplementary.

3.4 Optimization

8

Optimization serves as a role to introduce prior knowledge and avoid the accumulation of errors. During inference, we input the initial action a_{t+1} and state s_{t+1} and refine them into the fine-grained action a_{t+1}^* and state s_{t+1}^* . This refinement is achieved through gradient descent on the human and object pose parameters. Our optimization includes several loss terms: a fitting loss to align a_{t+1}^* and s_{t+1}^* with their preliminary counterparts, a velocity loss for temporal smoothness, a contact loss to promote occurring contacts, and a collision loss to reduce penetration in the interaction. For *efficiency*, we perform optimization only if the loss is above a threshold. Specifically, given the reference interaction sequence $\{h_i\}_{i=1}^L$ and $\{o_i\}_{i=1}^L$ of arbitrary length L, derived from previous steps, we apply gradient descent to optimize human pose sequence $\{h_i\}_{i=1}^L$ and object pose sequence $\{o_i^*\}_{i=1}^L$, using the loss function,

$$E_{\rm opt} = \lambda_{\rm fit} E_{\rm fit} + \lambda_{\rm vel} E_{\rm vel} + \lambda_{\rm cont} E_{\rm cont} + \lambda_{\rm pene} E_{\rm pene}, \qquad (4)$$

where λ_{fit} , λ_{vel} , λ_{cont} , and λ_{pene} are hyperparameters. **Fitting Loss.** We minimize the L1 distance between the input and the reference,

$$E_{\text{fit}} = \sum_{i=1}^{L} (\|\boldsymbol{h}_{i}^{*} - \boldsymbol{h}_{i}\|_{1} + \|\boldsymbol{o}_{i}^{*} - \boldsymbol{o}_{i}\|_{1}).$$
(5)

Velocity Loss. We leverage a velocity loss to smooth the interaction sequence,

$$E_{\text{vel}} = \sum_{i=1}^{L-1} (\|\boldsymbol{h}_{i+1}^* - \boldsymbol{h}_i^*\|_1 + \|\boldsymbol{o}_{i+1}^* - \boldsymbol{o}_i^*\|_1).$$
(6)



Fig. 3: Qualitative results on the BEHAVE dataset [7]. The interaction sequences are presented through a time-series visualization where color changes denote progression through frames. Frames are separately visualized when the pelvis remains nearly static. Here, our synergized knowledge comes from GPT-4 [67] and MotionGPT [35].

Contact Loss. We leverage a contact loss to encourage body part to contact the object surface, if they are close to each other in the initial interaction,

$$E_{\text{cont}} = \sum_{i=1}^{L} \sum_{d_h \in \mathcal{T}_i} \min_{d_o} \| \boldsymbol{v}_{\boldsymbol{o}_i^*}[d_o] - \boldsymbol{v}_{\boldsymbol{h}_i^*}[d_h] \|_2.$$
(7)

where $\boldsymbol{v}_{\boldsymbol{h}_i^*}[d_h]$ denotes the vertex on the human body surface, and $\boldsymbol{v}_{\boldsymbol{o}_i^*}[d_o]$ represents the corresponding vertex on the surface of the object. And $\mathcal{T}_i = \{d_h | \min_{d_o} \|\boldsymbol{v}_{\boldsymbol{o}_i}[d_o] - \boldsymbol{v}_{\boldsymbol{h}_i}[d_h] \|_2 \leq \epsilon\}$ includes the index of reference human vertex $\boldsymbol{v}_{\boldsymbol{h}_i}[d_h]$ that is close to the reference object vertex $\boldsymbol{v}_{\boldsymbol{o}_i}[d_o]$, where ϵ is a hyperparameter, d_h and d_o are vertex indices for human mesh and object mesh, respectively. **Penetration Loss.** Given the signed-distance field of the human pose $\mathbf{sdf}_{\boldsymbol{h}_i^*}$, we employ a penetration loss to penalize the body-object interpenetration,

$$E_{\text{pene}} = -\sum_{i=1}^{L} \sum_{d_o} \min(\mathbf{sdf}_{\boldsymbol{h}_i^*}(\boldsymbol{v}_{\boldsymbol{o}_i^*}[d_o]), 0).$$
(8)

4 Experiments

Extensive comparisons evaluate the performance of our InterDreamer across two motion-relevant tasks. Details of the evaluation settings are provided in Sec. 4.1. We present both quantitative (Sec. 4.2) and qualitative results (Sec. 4.3) for our approach. Additionally, we perform ablation studies to verify the efficacy of each component within our framework. These studies also cover the interaction prediction task [116] to evaluate our dynamics model. Additional details and results are presented Sec. C and Sec. D of the supplementary.



A person turns to their right and step on and over a chair and steps down.

Fig. 4: Qualitative results in more challenge scenarios with *free-form input* not from our annotations, showing the ability of our InterDreamer to fit *different object sizes* and handle *complex and long sequences*. Here, our synergized knowledge comes from GPT-4 [67] and MotionGPT [35].

Table 1: Quantitative results on human motion quality on the BEHAVE dataset with our annotation. We show that our high-level planning narrows the *distributional gap* and effectively adapts single human generators into zero-shot human-object interaction generation. To evaluate R-Precision, a batch size of 16 is selected.

Methods	Planning	$\operatorname{R-Precision}^{\uparrow}$		FID↓	MM Dist↓	Multimodality [↑]	Diversity→	
hielioub	(\mathbf{Ours})	Top 1	Top 2	Top 3	110	11111 15150	manning	Direibity
Ground Truth	-	$0.237^{\pm 0.004}$	$0.392^{\pm 0.004}$	$0.496^{\pm 0.005}$	$0.024^{\pm 0.000}$	$4.259^{\pm 0.006}$	-	$6.510^{\pm 0.227}$
MDM [93]	× ✓	$\begin{array}{c} 0.153^{\pm 0.016} \\ \textbf{0.163}^{\pm 0.010} \end{array}$	$\begin{array}{c} 0.279^{\pm 0.026} \\ \textbf{0.307}^{\pm 0.043} \end{array}$	$\begin{array}{c} 0.398^{\pm 0.016} \\ \textbf{0.402}^{\pm 0.019} \end{array}$	$12.279^{\pm 0.217} \\ 10.374^{\pm 0.304}$	$5.351^{\pm 0.057}$ $5.303^{\pm 0.117}$	$7.604^{\pm 0.190} \\ 7.281^{\pm 0.083}$	$7.598^{\pm 0.334} \\ \textbf{7.471}^{\pm 0.427}$
MotionDiffuse [130]	× ✓	$\begin{array}{c} 0.205^{\pm 0.011} \\ \textbf{0.216}^{\pm 0.032} \end{array}$	$\begin{array}{c} 0.351^{\pm 0.002} \\ \textbf{0.369}^{\pm 0.023} \end{array}$	$\begin{array}{c} 0.458^{\pm 0.021} \\ 0.472^{\pm 0.027} \end{array}$	$\begin{array}{c} 10.208^{\pm 0.500} \\ \textbf{9.015}^{\pm 0.403} \end{array}$	$4.837^{\pm 0.064} \\ 4.649^{\pm 0.029}$	$4.520^{\pm 0.163} \\ 4.991^{\pm 0.172}$	$7.323^{\pm 0.412} \\ 7.295^{\pm 0.501}$
ReMoDiffuse [131]	× ✓	$\begin{array}{c} 0.196^{\pm 0.009} \\ 0.223^{\pm 0.006} \end{array}$	$\begin{array}{c} 0.338^{\pm 0.011} \\ \textbf{0.368}^{\pm 0.015} \end{array}$	$\begin{array}{c} 0.448^{\pm 0.012} \\ 0.482^{\pm 0.011} \end{array}$	$\begin{array}{c} 6.385^{\pm0.201} \\ 5.237^{\pm0.174} \end{array}$	$4.855^{\pm 0.029} \\ 4.784^{\pm 0.053}$	$5.889^{\pm 0.524} \\ 6.350^{\pm 0.411}$	$7.160^{\pm 0.306} \\ 7.201^{\pm 0.318}$
MotionGPT [35]	× ✓	$0.233^{\pm 0.003} \\ 0.234^{\pm 0.004}$	$\begin{array}{c} 0.344^{\pm 0.004} \\ \textbf{0.387}^{\pm 0.003} \end{array}$	$\begin{array}{c} 0.457^{\pm 0.005} \\ 0.471^{\pm 0.007} \end{array}$	$5.497^{\pm 0.106} \\ 4.751^{\pm 0.121}$	$5.205^{\pm 0.027} \\ 4.995^{\pm 0.003}$	$1.062^{\pm 0.211} \\ 1.337^{\pm 0.193}$	$8.316^{\pm 0.204} \\ \textbf{7.106}^{\pm 0.487}$

4.1 Experimental Setup

Datasets. We evaluate our model on the BEHAVE dataset [7], which includes recordings of 8 individuals interacting with 20 everyday objects. Our analysis focuses on 18 objects for which interaction sequences are available at 30 Hz. The human pose is modeled using SMPL-H [58, 82], with hand poses set to an average pose *due to the absence of detailed hand pose in the dataset*. Object poses are rotation matrices and translations. We manually segment the long interaction sequences in the test set, and annotate them with descriptions as well as their starting and ending indices, leading to 532 subsequences for evaluation. The CHAIRS [36] dataset encompasses the capture of 46 subjects, represented via SMPL-X [70] bodies, engaging with 81 distinct types of chairs and sofas. Without whole dataset annotations, we use the CHAIRS dataset solely for qualitative evaluation. We further use GPT-4 [67] to rephrase and diversify annotations: 1) *less complexity*: someone holds a backpack and steps left; 2) *our annotation*: a person holds a backpack in front of them with both hands and takes



Fig. 5: Qualitative results on the CHAIRS dataset [36]. Our dynamics model trained on the BEHAVE dataset [7] generalizes well on the CHAIRS dataset unseen in training. Interaction sequences are visualized through a time-series style where color changes denote progression through frames. Frames are separately visualized. Here, high-level planning and low-level control use GPT-4 [67] and MotionGPT [35], respectively.

Table 2: Quantitative results on the zero-shot text-to-interaction task (left) and on the interaction prediction task (right). Our dynamics model with vertex-based control generates interactions with the best quality.

Methods	Zero-shot text-to-interaction		Interaction prediction [116]			
hethous	$\overline{\mathrm{CMD}\downarrow}$	Pene. (10 ⁻² %) \downarrow	Trans. Err. (mm) $\downarrow I$	Rot. Err. $(10^{-3} \text{ rad}) \downarrow \text{Pe}$	ene. $(10^{-2}\%)$	
w/o control	0.424	533	123	256	228	
w/ marker control (InterDiff [116])	0.219	484	123	226	164	
w/ raw control	0.325	957	129	265	218	
w/ vertex control (ours)	0.151	443	119	221	156	

a step to the left; 3) more complexity: with both hands, a person clutches a heavy backpack firmly and brings it close to their body, then steps to the left with their left leg. Metrics. The evaluation metrics are divided into three categories: (i) Human motion quality: The Fréchet Inception Distance (FID) measures the distributional distance between the generated motions and ground truth. The MultiModality (Multimodality) and Diversity metrics assess the variance in generated human motion. R-Precision evaluates the consistency between the text and the generated human motion within the latent feature space. Multi-Modal distance (**MM Dist**) is the distance between the motion feature and the text feature. We follow [25] to generate motion and text features. (ii) Interaction quality: We propose a metric to measure the distance between contact maps of real interactions and those generated (CMD). The per-sequence contact map is defined by the percentage of time that each body part is actively in contact. The detailed formulation is provided in Sec. C of supplementary. We measure the collision (**Pene.** [116]), which calculates the average percentage of object vertices that have non-negative values in the human signed distance fields [69]. (iii) *Object motion accuracy*: The dynamics model's performance in the interaction prediction task [116] is evaluated by the accuracy of predicted object motions, including **Trans. Err.**, the average distance between predicted and ground truth, and Rot. Err., the average distance between the predicted and ground truth. **Baselines.** As we are introducing a new task, there is no established base*line* available at the current stage. Note that it is unfair to compare our work with concurrent supervised learning approaches [21, 50, 71], and their code is not publicly available. To facilitate our comparisons, we develop various baselines to evaluate both our overall pipeline and its individual components. In the



(a) Motion w/ planning v.s. w/o planning

(b) Text feature w/ planning v.s. w/o planning

Fig. 6: (a) Ablation study on the high-level planning. On the *left* are results from MotionGPT [35] based on raw descriptions ("w/o planning"); on the *right* are results with our planning ("w/ planning"). *Free-form* descriptions with the out-of-distribution object name lead to results that bear little resemblance to the description. (b) We visualize CLIP [78] features of descriptions on HumanML3D [25], our initial raw annotations ("w/o planning"), and the annotations processed through our high-level planning framework ("w/ planning"). The CLIP features are processed via t-SNE [61].

Table 3: Ablation study on the high-level planning. Q1 and Q2 ask to identify the object category and the contact body part, respectively. We assess the accuracy by comparing the LLM's responses with labels we annotate. Note that the text input to LLMs may contain ambiguities; for example, the annotation is "hand" when the motion uses "right hand." We include Q1 Acc^{*} and Q2 Acc^{*} excluding ambiguous text.

LLM (# of parameters)	Q1 Acc \uparrow	Q1 Acc* \uparrow	Q2 Acc \uparrow	Q2 Acc* \uparrow
GPT-4 [67]	0.801	0.997	0.703	0.964
Llama-2 (7B) [95]	0.073	0.147	0.436	0.689
Llama-2 (13B) [95]	0.232	0.319	0.662	0.853
Llama-2 (70B) [95]	0.722	0.967	0.798	0.907

context of high-level planning, we utilize GPT-4 [67] and Llama-2 [95], illustrating the effectiveness of our prompts across different language models. For low-level motion generation control, our baselines include MDM [93], Motion-Diffuse [130], ReMoDiffuse [131], and MotionGPT [35], which span a range of text-to-motion approaches trained on HumanML3D [25] and show the generalizability of our framework. Our dynamics model baselines are varied, comprising: InterDreamer with InterDiff [116], which adopts their Interaction Correction module as the dynamics model; InterDreamer without control, which operates object dynamics independently of human motion; and InterDreamer with raw control, utilizing unprocessed human motion to guide the dynamics.

4.2 Quantitative Results

Table 1 presents a comparative analysis of our approach, InterDreamer adopting various text-to-motion models against four counterparts on the BEHAVE [7] dataset. Our approach generalizes across various models and consistently outperforms baselines. Specifically, InterDreamer exhibits superior motion quality, reflected by a significantly lower FID, higher R-Precision, and better diversity, highlighting the benefits of incorporating our planning to reduce the distribution gap in the zero-shot setting.



Fig. 7: Results from the interaction retrieval. We demonstrate that our proposed retrieval approach based on handcraft rules can extract diverse and realistic interactions. Of these, one interaction is sampled based on a pre-defined metric for subsequent steps.



Fig. 8: Ablation study on the dynamics modeling. Given the text description of "A person walks clockwise while holding a small box with left hand," our (b) vertex-based control can synthesize consistent contacts, which (a) the baseline fails to do.

In Table 2, comparing our whole pipeline to baselines without explicit control, InterDreamer achieves better interaction quality in terms of CMD and penetration scores, showing the importance of human influence on object motion. Against methods that utilize direct raw human motion or markers [116] for control, our full method demonstrates enhanced performance by offering more fine-grained guidance and extracting generalizable features.

4.3 Qualitative Results

Fig. 3 displays several results guided by the text that we annotate on the BE-HAVE dataset [7]. Our method exhibits proficiency in interpreting the textual input and synthesizing dynamic, realistic interactions, despite the absence of training with text-interaction paired data. More importantly, our method can process free-form language descriptions of motion that deviate from our annotation, exhibiting zero-shot generalizability in its performance, as illustrated in Fig. 4, where we selectively use more complex sequences of interactive descriptions that are *beyond the scope of the original dataset*. Fig. 5 further exemplifies the zero-shot ability of our method that is able to generalize effectively to the CHAIRS dataset [36], despite our dynamics model not being trained on it. Fig. 7 depicts the retrieval procedure, resulting in a diverse set of interactions that are both high-quality and semantically aligned. More experimental results and the **user study** are presented in Sec. D of supplementary.

4.4 Ablation Study

Adaptability of high-level planning. Is our framework adaptable across different large language models (LLMs)? As illustrated in Table 3, our analysis

14 S. Xu, Z. Wang, Y.-X. Wang, L.-Y. Gui

contains two types of language models: GPT-4 [67], which is accessible through APIs and operates as a black box model; and Llama-2 [95], which is an opensource model. We demonstrate that language models with large parameters exhibit very high accuracy in responding to questions tailored to our prompts, thereby validating the framework's adaptability.

Effectiveness of text-to-motion with high-level planning. In consistency with Table 1, Fig. 6 offers a qualitative comparison of text-to-motion results, contrasting outputs with and without augmentation by LLM-revised text descriptions. The comparison shows that motions generated without LLM-enhanced descriptions often fail to correspond with the intended text, if the text with the object description is not in the distribution of training data on HumanML3D [25]. This underscores the LLM's critical role in bridging the gap in a zero-shot way. On the right-hand side of Fig. 6, we see that planning not only reduces the distributional gap in motion, but also directly reduces it in text. We visualize the CLIP [78] features of descriptions on HumanML3D, our raw annotations, and the annotations processed by high-level planning. The text processed by the planning shows more similarity to the in-distributional text, where the average cosine similarity is 0.932 over 0.913 from the raw annotation.

Effectiveness of world model. In the quantitative evaluation, we show that the performance of our pipeline is enhanced by the tailored design of our world model. Table 2 provides additional evidence of this effectiveness by integrating the proposed world model, as interaction correction within the InterDiff framework [116] in the interaction prediction task. This implementation demonstrates enhanced conditionality in the object dynamics modeling across various tasks, attributed to the vertex-level control. Doing so effectively removes the wholebody complexity, most of which tends to be irrelevant to the interaction. Fig. 8 further indicates that our vertex-based control is able to establish consistent interactions over time, while guidance from motion features is not robust.

5 Conclusion

We introduce the novel task of text-guided 3D human-object interaction generation, and we aim to achieve this without reliance on text-interaction pair data. To this end, we present InterDreamer that decouples interaction dynamics from semantics, where high-level planning and low-level control are introduced to generate semantically aligned human motion and initial object pose, while a world model is responsible for the object dynamics guided by the interaction. Our approach demonstrates promising effectiveness in this novel task, suggesting its considerable potential for various real-world applications.

Limitations. The current utilization of dynamics modeling could be enhanced. A prospective improvement involves incorporating model-based learning techniques, which empower the agent to more effectively interact with the environment and learn a broader range of skills. The generated results may not be physically plausible and hand poses are rough because they are missing from the dataset, but could be improved by integrating a physics simulator.

References

- Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. In: 3DV (2019)
- Athanasiou, N., Petrovich, M., Black, M.J., Varol, G.: Teach: Temporal action composition for 3d humans. In: 3DV (2022)
- 3. Athanasiou, N., Petrovich, M., Black, M.J., Varol, G.: SINC: Spatial composition of 3d human motions for simultaneous action generation. In: ICCV (2023)
- 4. Bae, J., Won, J., Lim, D., Min, C.H., Kim, Y.M.: Pmp: Learning to physically interact with environments using part-wise motion priors. In: SIGGRAPH (2023)
- 5. Barquero, G., Escalera, S., Palmero, C.: BeLFusion: Latent diffusion for behaviordriven human motion prediction. In: ICCV (2023)
- Barquero, G., Escalera, S., Palmero, C.: Seamless human motion composition with blended positional encodings. In: CVPR (2024)
- Bhatnagar, B.L., Xie, X., Petrov, I., Sminchisescu, C., Theobalt, C., Pons-Moll, G.: BEHAVE: Dataset and method for tracking human object interactions. In: CVPR (2022)
- Braun, J., Christen, S., Kocabas, M., Aksan, E., Hilliges, O.: Physically plausible full-body hand-object interaction synthesis. arXiv preprint arXiv:2309.07907 (2023)
- Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023)
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are fewshot learners. In: NeurIPS (2020)
- 11. Cao, Z., Gao, H., Mangalam, K., Cai, Q.Z., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: ECCV (2020)
- 12. Casas, D., Comino-Trinidad, M.: SMPLitex: A generative model and dataset for 3d human texture estimation from single image. In: BMVC (2023)
- 13. Chao, Y.W., Yang, J., Chen, W., Deng, J.: Learning to sit: Synthesizing humanchair interactions via hierarchical control. In: AAAI (2021)
- 14. Chen, L.H., Zhang, J., Li, Y., Pang, Y., Xia, X., Liu, T.: HumanMAC: Masked motion completion for human motion prediction. In: ICCV (2023)
- 15. Chen, X., Jiang, B., Liu, W., Huang, Z., Fu, B., Chen, T., Yu, G.: Executing your commands via motion diffusion in latent space. In: CVPR (2023)
- Chi, C., Feng, S., Du, Y., Xu, Z., Cousineau, E., Burchfiel, B., Song, S.: Diffusion policy: Visuomotor policy learning via action diffusion. In: RSS (2023)
- Cong, P., Dou, Z.W., Ren, Y., Yin, W., Cheng, K., Sun, Y., Long, X., Zhu, X., Ma, Y.: LaserHuman: Language-guided scene-aware human motion generation in free environment. arXiv preprint arXiv:2403.13307 (2024)
- Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F.: Context-aware human motion prediction. In: CVPR (2020)
- 19. Cui, J., Liu, T., Liu, N., Yang, Y., Zhu, Y., Huang, S.: AnySkill: Learning openvocabulary physical skill for interactive agents. In: CVPR (2024)
- Dabral, R., Mughal, M.H., Golyanik, V., Theobalt, C.: MoFusion: A framework for denoising-diffusion-based motion synthesis. In: CVPR. pp. 9760–9770 (2023)
- Diller, C., Dai, A.: CG-HOI: Contact-guided 3d human-object interaction generation. In: CVPR (2024)
- Fan, Z., Taheri, O., Tzionas, D., Kocabas, M., Kaufmann, M., Black, M.J., Hilliges, O.: ARCTIC: A dataset for dexterous bimanual hand-object manipulation. In: CVPR (2023)

- Ghosh, A., Dabral, R., Golyanik, V., Theobalt, C., Slusallek, P.: IMoS: Intentdriven full-body motion synthesis for human-object interactions. arXiv preprint arXiv:2212.07555 (2022)
- Ghosh, A., Dabral, R., Golyanik, V., Theobalt, C., Slusallek, P.: ReMoS: Reactive 3d motion synthesis for two-person interactions. arXiv preprint arXiv:2311.17057 (2023)
- 25. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: CVPR (2022)
- 26. Guo, C., Zuo, X., Wang, S., Cheng, L.: Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In: ECCV (2022)
- 27. Guo, C., Zuo, X., Wang, S., Zou, S., Sun, Q., Deng, A., Gong, M., Cheng, L.: Action2motion: Conditioned generation of 3d human motions. In: ACMMM (2020)
- Han, S., Joo, H.: CHORUS: Learning canonicalized 3d human-object spatial relations from unbounded synthesized images. In: ICCV (2023)
- Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3d scenes by learning human-scene interaction. In: CVPR (2021)
- Hassan, M., Guo, Y., Wang, T., Black, M., Fidler, S., Peng, X.B.: Synthesizing physical character-scene interactions. In: SIGGRAPH (2023)
- Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020)
- Hou, Z., Yu, B., Tao, D.: Compositional 3d human-object neural animation. arXiv preprint arXiv:2304.14070 (2023)
- 33. Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based generation, optimization, and planning in 3d scenes. In: CVPR (2023)
- 34. Huang, Y., Taheri, O., Black, M.J., Tzionas, D.: InterCap: Joint markerless 3D tracking of humans and objects in interaction. In: GCPR (2022)
- Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: MotionGPT: Human motion as a foreign language. In: NeurIPS (2023)
- Jiang, N., Liu, T., Cao, Z., Cui, J., Chen, Y., Wang, H., Zhu, Y., Huang, S.: CHAIRS: Towards full-body articulated human-object interaction. In: ICCV (2023)
- Jiang, N., Zhang, Z., Li, H., Ma, X., Wang, Z., Chen, Y., Liu, T., Zhu, Y., Huang, S.: Scaling up dynamic human-scene interaction modeling. In: CVPR (2024)
- Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: GMD: Controllable human motion synthesis via guided diffusion models. In: ICCV (2023)
- Kaufmann, M., Aksan, E., Song, J., Pece, F., Ziegler, R., Hilliges, O.: Convolutional autoencoders for human motion infilling. In: 3DV (2020)
- Kim, H., Han, S., Kwon, P., Joo, H.: Zero-shot learning for the primitives of 3d affordance in general objects. arXiv preprint arXiv:2401.12978 (2024)
- Kim, J., Kim, J., Na, J., Joo, H.: ParaHome: Parameterizing everyday home activities towards 3d generative modeling of human-object interactions. arXiv preprint arXiv:2401.10232 (2024)
- Kim, J., Kim, J., Choi, S.: Flame: Free-form language-based motion synthesis & editing. In: AAAI (2023)
- 43. Kim, S.W., Zhou, Y., Philion, J., Torralba, A., Fidler, S.: Learning to simulate dynamic environments with gamegan. In: CVPR (2020)
- 44. Kim, T., Saito, S., Joo, H.: NCHO: Unsupervised learning for neural 3d composition of humans and objects. In: ICCV (2023)
- 45. Kong, H., Gong, K., Lian, D., Mi, M.B., Wang, X.: Priority-centric human motion generation in discrete latent space. In: ICCV (2023)

- Krebs, F., Meixner, A., Patzer, I., Asfour, T.: The kit bimanual manipulation dataset. In: Humanoids (2021)
- Kulkarni, N., Rempe, D., Genova, K., Kundu, A., Johnson, J., Fouhey, D., Guibas, L.: NIFTY: Neural object interaction fields for guided human motion synthesis. arXiv preprint arXiv:2307.07511 (2023)
- 48. Lee, J., Joo, H.: Locomotion-Action-Manipulation: Synthesizing human-scene interactions in complex 3d environments. In: ICCV (2023)
- 49. Lee, T., Moon, G., Lee, K.M.: Multiact: Long-term 3d human motion generation from multiple action labels. In: AAAI (2023)
- Li, J., Clegg, A., Mottaghi, R., Wu, J., Puig, X., Liu, C.K.: Controllable humanobject interaction synthesis. arXiv preprint arXiv:2312.03913 (2023)
- Li, J., Wu, J., Liu, C.K.: Object motion guided human motion synthesis. ACM Transactions on Graphics (TOG) 42(6), 1–11 (2023)
- 52. Li, L., Dai, A.: GenZI: Zero-shot 3d human-scene interaction generation. In: CVPR (2024)
- Li, Q., Wang, J., Loy, C.C., Dai, B.: Task-oriented human-object interactions generation with implicit neural representations. arXiv preprint arXiv:2303.13129 (2023)
- Liang, H., Zhang, W., Li, W., Yu, J., Xu, L.: InterGen: Diffusion-based multi-human motion generation under complex interactions. arXiv preprint arXiv:2304.05684 (2023)
- 55. Lin, J., Zeng, A., Lu, S., Cai, Y., Zhang, R., Wang, H., Zhang, L.: Motion-X: A large-scale 3d expressive whole-body human motion dataset. In: NeurIPS (2023)
- Liu, L., Hodgins, J.: Learning basketball dribbling skills using trajectory optimization and deep reinforcement learning. ACM Transactions on Graphics (TOG) 37(4), 1–14 (2018)
- 57. Liu, Y., Chen, C., Yi, L.: Interactive humanoid: Online full-body motion reaction synthesis with social affordance canonicalization and forecasting. arXiv preprint arXiv:2312.08983 (2023)
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM transactions on graphics (2015)
- Lu, S., Chen, L.H., Zeng, A., Lin, J., Zhang, R., Zhang, L., Shum, H.Y.: HumanTOMATO: Text-aligned whole-body motion generation. arXiv preprint arXiv:2310.12978 (2023)
- Ma, S., Cao, Q., Zhang, J., Tao, D.: Contact-aware human motion generation from textual descriptions. arXiv preprint arXiv:2403.15709 (2024)
- Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research 9(11) (2008)
- Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of motion capture as surface shapes. In: ICCV (2019)
- Mandery, C., Terlemez, O., Do, M., Vahrenkamp, N., Asfour, T.: The kit wholebody human motion database. In: ICAR (2015)
- 64. Mandery, C., Terlemez, O., Do, M., Vahrenkamp, N., Asfour, T.: Unifying representations and large-scale whole-body motion databases for studying human motion. IEEE Transactions on Robotics **32**(4), 796–809 (2016)
- 65. Mao, W., Liu, M., Salzmann, M.: Generating smooth pose sequences for diverse human motion prediction. In: CVPR (2021)
- Merel, J., Tunyasuvunakool, S., Ahuja, A., Tassa, Y., Hasenclever, L., Pham, V., Erez, T., Wayne, G., Heess, N.: Catch & carry: reusable neural controllers for vision-guided whole-body tasks. ACM Transactions on Graphics (TOG) 39(4), 39–1 (2020)

18 S. Xu, Z. Wang, Y.-X. Wang, L.-Y. Gui

- 67. OpenAI: ChatGPT. https://chat.openai.com/ (2023)
- Pan, L., Wang, J., Huang, B., Zhang, J., Wang, H., Tang, X., Wang, Y.: Synthesizing physically plausible human motions in 3d scenes. arXiv preprint arXiv:2308.09036 (2023)
- Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR (2019)
- Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3D hands, face, and body from a single image. In: CVPR (2019)
- Peng, X., Xie, Y., Wu, Z., Jampani, V., Sun, D., Jiang, H.: HOI-Diff: Text-driven synthesis of 3d human-object interactions using diffusion models. arXiv preprint arXiv:2312.06553 (2023)
- 72. Petrov, I.A., Marin, R., Chibane, J., Pons-Moll, G.: Object pop-up: Can we infer 3d objects and their poses from human interactions alone? In: CVPR (2023)
- 73. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3d human motion synthesis with transformer vae. In: ICCV (2021)
- 74. Petrovich, M., Black, M.J., Varol, G.: TEMOS: Generating diverse human motions from textual descriptions. In: ECCV (2022)
- 75. Petrovich, M., Black, M.J., Varol, G.: TMR: Text-to-motion retrieval using contrastive 3d human motion synthesis. In: ICCV (2023)
- Raab, S., Leibovitch, I., Li, P., Aberman, K., Sorkine-Hornung, O., Cohen-Or, D.: MoDi: Unconditional motion synthesis from diverse data. In: CVPR (2023)
- Raab, S., Leibovitch, I., Tevet, G., Arar, M., Bermano, A.H., Cohen-Or, D.: Single motion diffusion. arXiv preprint arXiv:2302.05905 (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
- 79. Razali, H., Demiris, Y.: Action-conditioned generation of bimanual object manipulation sequences. In: AAAI (2023)
- Rempe, D., Luo, Z., Bin Peng, X., Yuan, Y., Kitani, K., Kreis, K., Fidler, S., Litany, O.: Trace and pace: Controllable pedestrian animation via guided trajectory diffusion. In: CVPR (2023)
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
- 82. Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. ACM Transactions on Graphics **36**(6) (2017)
- Seo, Y., Hafner, D., Liu, H., Liu, F., James, S., Lee, K., Abbeel, P.: Masked world models for visual control. In: CoRL (2023)
- Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. arXiv preprint arXiv:2303.01418 (2023)
- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: ICML (2015)
- Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
- Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for characterscene interactions. ACM Trans. Graph. 38(6), 209–1 (2019)
- Starke, S., Zhao, Y., Komura, T., Zaman, K.: Local motion phases for learning multi-contact character movements. ACM Transactions on Graphics (TOG) 39(4), 54–1 (2020)

- Taheri, O., Choutas, V., Black, M.J., Tzionas, D.: GOAL: Generating 4d wholebody motion for hand-object grasping. In: CVPR (2022)
- Taheri, O., Ghorbani, N., Black, M.J., Tzionas, D.: GRAB: A dataset of wholebody human grasping of objects. In: ECCV (2020)
- Tendulkar, P., Surís, D., Vondrick, C.: FLEX: Full-body grasping without fullbody grasps. In: CVPR (2023)
- 92. Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. In: ECCV (2022)
- Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-Or, D., Bermano, A.H.: Human motion diffusion model. arXiv preprint arXiv:2209.14916 (2022)
- 94. Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-or, D., Bermano, A.H.: Human motion diffusion model. In: ICLR (2023)
- 95. Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al.: Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 (2023)
- 96. Turk, A.M.: Amazon mechanical turk. Retrieved August 17, 2012 (2012)
- 97. Wan, W., Dou, Z., Komura, T., Wang, W., Jayaraman, D., Liu, L.: Tlcontrol: Trajectory and language control for human motion synthesis. arXiv preprint arXiv:2311.17135 (2023)
- 98. Wan, W., Yang, L., Liu, L., Zhang, Z., Jia, R., Choi, Y.K., Pan, J., Theobalt, C., Komura, T., Wang, W.: Learn to predict how humans manipulate large-sized objects from interactive motions. IEEE Robotics and Automation Letters (2022)
- 99. Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: CVPR (2021)
- 100. Wang, J., Rong, Y., Liu, J., Yan, S., Lin, D., Dai, B.: Towards diverse and natural scene-aware 3d human motion synthesis. In: CVPR (2022)
- Wang, J., Yan, S., Dai, B., Lin, D.: Scene-aware generative network for human motion synthesis. In: CVPR (2021)
- Wang, X., Li, G., Kuo, Y.L., Kocabas, M., Aksan, E., Hilliges, O.: Reconstructing action-conditioned human-object interactions using commonsense knowledge priors. In: 3DV (2022)
- Wang, Y., Lin, J., Zeng, A., Luo, Z., Zhang, J., Zhang, L.: PhysHOI: Physics-based imitation of dynamic human-object interaction. arXiv preprint arXiv:2312.04393 (2023)
- 104. Wang, Z., Chen, Y., Liu, T., Zhu, Y., Liang, W., Huang, S.: HUMANISE: Language-conditioned human motion generation in 3d scenes. In: NeurIPS (2022)
- 105. Wang, Z., Wang, J., Lin, D., Dai, B.: InterControl: Generate human motion interactions by controlling every joint. arXiv preprint arXiv:2311.15864 (2023)
- 106. Wei, D., Sun, X., Sun, H., Li, B., Hu, S., Li, W., Lu, J.: Understanding text-driven motion synthesis with keyframe collaboration via diffusion models. arXiv preprint arXiv:2305.13773 (2023)
- 107. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. In: NeurIPS (2022)
- Wu, P., Escontrela, A., Hafner, D., Abbeel, P., Goldberg, K.: Daydreamer: World models for physical robot learning. In: CoRL (2023)
- Wu, Q., Shi, Y., Huang, X., Yu, J., Xu, L., Wang, J.: THOR: Text to human-object interaction diffusion via relation intervention. arXiv preprint arXiv:2403.11208 (2024)
- 110. Wu, Y., Wang, J., Zhang, Y., Zhang, S., Hilliges, O., Yu, F., Tang, S.: SAGA: Stochastic whole-body grasping with contact. In: ECCV (2022)

- 111. Xiao, Z., Wang, T., Wang, J., Cao, J., Zhang, W., Dai, B., Lin, D., Pang, J.: Unified human-scene interaction via prompted chain-of-contacts. arXiv preprint arXiv:2309.07918 (2023)
- 112. Xie, X., Bhatnagar, B.L., Pons-Moll, G.: Chore: Contact, human and object reconstruction from a single rgb image. In: ECCV (2022)
- Xie, Y., Jampani, V., Zhong, L., Sun, D., Jiang, H.: OmniControl: Control any joint at any time for human motion generation. arXiv preprint arXiv:2310.08580 (2023)
- 114. Xie, Z., Starke, S., Ling, H.Y., van de Panne, M.: Learning soccer juggling skills with layer-wise mixture-of-experts. In: SIGGRAPH (2022)
- 115. Xie, Z., Tseng, J., Starke, S., van de Panne, M., Liu, C.K.: Hierarchical planning and control for box loco-manipulation. arXiv preprint arXiv:2306.09532 (2023)
- 116. Xu, S., Li, Z., Wang, Y.X., Gui, L.Y.: InterDiff: Generating 3d human-object interactions with physics-informed diffusion. In: ICCV (2023)
- 117. Xu, S., Wang, Y.X., Gui, L.Y.: Diverse human motion prediction guided by multilevel spatial-temporal anchors. In: ECCV (2022)
- Xu, S., Wang, Y.X., Gui, L.: Stochastic multi-person 3d motion forecasting. In: ICLR (2023)
- Xu, X., Joo, H., Mori, G., Savva, M.: D3D-HOI: Dynamic 3d human-object interactions from videos. arXiv preprint arXiv:2108.08420 (2021)
- 120. Yang, Y., Zhai, W., Luo, H., Cao, Y., Zha, Z.J.: LEMON: Learning 3d humanobject interaction relation from 2d images. In: CVPR (2024)
- 121. Yang, Z., Yin, K., Liu, L.: Learning to use chopsticks in diverse gripping styles. ACM Transactions on Graphics (TOG) 41(4), 1–17 (2022)
- 122. Yao, H., Song, Z., Zhou, Y., Ao, T., Chen, B., Liu, L.: MoConVQ: Unified physics-based motion control via scalable discrete representations. arXiv preprint arXiv:2310.10198 (2023)
- Yazdian, P.J., Liu, E., Cheng, L., Lim, A.: MotionScript: Natural language descriptions for expressive 3d human motions. arXiv preprint arXiv:2312.12634 (2023)
- 124. Ye, Y., Li, X., Gupta, A., De Mello, S., Birchfield, S., Song, J., Tulsiani, S., Liu, S.: Affordance diffusion: Synthesizing hand-object interactions. In: CVPR (2023)
- 125. Yuan, Y., Kitani, K.: DLow: Diversifying latent flows for diverse human motion prediction. In: ECCV (2020)
- 126. Zhang, H., Christen, S., Fan, Z., Zheng, L., Hwangbo, J., Song, J., Hilliges, O.: ArtiGrasp: Physically plausible synthesis of bi-manual dexterous grasping and articulation. arXiv preprint arXiv:2309.03891 (2023)
- 127. Zhang, J.Y., Pepose, S., Joo, H., Ramanan, D., Malik, J., Kanazawa, A.: Perceiving 3d human-object spatial arrangements from a single image in the wild. In: ECCV (2020)
- 128. Zhang, J., Zhang, Y., Cun, X., Zhang, Y., Zhao, H., Lu, H., Shen, X., Shan, Y.: Generating human motion from textual descriptions with discrete representations. In: CVPR (2023)
- 129. Zhang, J., Luo, H., Yang, H., Xu, X., Wu, Q., Shi, Y., Yu, J., Xu, L., Wang, J.: NeuralDome: A neural modeling pipeline on multi-view human-object interactions. In: CVPR (2023)
- 130. Zhang, M., Cai, Z., Pan, L., Hong, F., Guo, X., Yang, L., Liu, Z.: MotionDiffuse: Text-driven human motion generation with diffusion model. arXiv preprint arXiv:2208.15001 (2022)
- Zhang, M., Guo, X., Pan, L., Cai, Z., Hong, F., Li, H., Yang, L., Liu, Z.: ReMoDiffuse: Retrieval-augmented motion diffusion model. In: ICCV (2023)

- 132. Zhang, W., Dabral, R., Leimkühler, T., Golyanik, V., Habermann, M., Theobalt, C.: ROAM: Robust and object-aware motion generation using neural pose descriptors. arXiv preprint arXiv:2308.12969 (2023)
- Zhang, X., Bhatnagar, B.L., Starke, S., Guzov, V., Pons-Moll, G.: COUCH: Towards controllable human-chair interactions. In: ECCV (2022)
- 134. Zhang, Y., Huang, D., Liu, B., Tang, S., Lu, Y., Chen, L., Bai, L., Chu, Q., Yu, N., Ouyang, W.: Motiongpt: Finetuned llms are general-purpose motion generators. arXiv preprint arXiv:2306.10900 (2023)
- 135. Zhang, Z., Liu, R., Aberman, K., Hanocka, R.: TEDi: Temporally-entangled diffusion for long-term motion synthesis. arXiv preprint arXiv:2307.15042 (2023)
- 136. Zhao, C., Zhang, J., Du, J., Shan, Z., Wang, J., Yu, J., Wang, J., Xu, L.: I'M HOI: Inertia-aware monocular capture of 3d human-object interactions. In: CVPR (2024)
- 137. Zhao, K., Wang, S., Zhang, Y., Beeler, T., Tang, S.: Compositional human-scene interaction synthesis with semantic control. In: ECCV (2022)
- Zhao, K., Zhang, Y., Wang, S., Beeler, T., Tang, S.: Synthesizing diverse human motions in 3d indoor scenes. In: ICCV (2023)
- Zheng, J., Zheng, Q., Fang, L., Liu, Y., Yi, L.: CAMS: Canonicalized manipulation spaces for category-level functional hand-object manipulation synthesis. In: CVPR (2023)
- 140. Zhou, K., Bhatnagar, B.L., Lenssen, J.E., Pons-Moll, G.: Toch: Spatio-temporal object-to-hand correspondence for motion refinement. In: ECCV (2022)
- 141. Zhou, W., Dou, Z., Cao, Z., Liao, Z., Wang, J., Wang, W., Liu, Y., Komura, T., Wang, W., Liu, L.: EMDM: Efficient motion diffusion model for fast, high-quality motion generation. arXiv preprint arXiv:2312.02256 (2023)

Supplementary Material

In this supplementary material, we include additional method details and experimental results: (i) We provide a demo video, explained in Sec. A. (ii) We present additional details of interaction retrieval, world model, and optimization in Sec. B. (iii) We provide implementation details and additional information on the experimental setup in Sec. C. (iv) We provide additional qualitative experiments in Sec. D.

A Visualization Video

Beyond the qualitative results presented in the main paper, we include a video on the project website that offers more detailed visualizations of the task, further illustrating the efficacy of our approach. These demos highlight (i) We conduct a qualitative comparison of our approach with existing text-to-HOI work [21,71] within the framework of supervised learning. Note that as our setting is the zero-shot generation, it is unfair to compare our work with these approaches; we include the comparison here for additional reference. We evaluate our method by directly testing our trained model on the annotated data available from their websites, specifically retrieving their generated videos for direct comparison. Remarkably, even without training on these datasets, our method generates results that demonstrate high-quality interactions. (ii) The smoothness of the interaction sequences generated by our method. (iii) Remarkably, our method is proficient in creating sequences that maintain *consistent* contact where contact remains largely unchanged during interactions, such as carrying a box and sitting in a chair. (iv) It is even capable of synthesizing complex interactions involving dynamically-changing contact, such as the handover and throwing of objects. (v) We contrast our framework, which incorporates vertex control, against generation processes that rely on raw motion control. The results demonstrate strong generalizability of our method.

B Additional Details of Methodology

B.1 Low-Level Control

Handcraft Interaction Retrieval. In Sec. 3.2 of the main paper, we detail the construction of the interaction database and emphasize the use of body parts and object categories as keys to fetch semantically-aligned contact maps. Same as the main paper, we define a contact map as a list of K index pairs of vertices $\{(d_h^i, d_o^i)\}_{i=1}^K$. This section delves into the methodology for outlining an optimization process to generate the object initial pose s_1 given contact maps and the initial human pose a_1 , and choose one pose based on a predefined metric.

Let $\boldsymbol{v}_{\boldsymbol{h}_1}[d_o]$ denote the vertex on the surface of the object, and $\boldsymbol{v}_{\boldsymbol{o}_1}[d_h]$ represent the corresponding vertex on the human body surface, where d_o and d_h

are the indices of vertices. Specifically, to optimize s_1 , the overall optimization objective is given by,

$$E_{\rm opt} = \lambda_{\rm fit} E_{\rm fit} + \lambda_{\rm cont} E_{\rm cont} + \lambda_{\rm pene} E_{\rm pene}, \qquad (9)$$

where λ_{fit} , λ_{cont} , and λ_{pene} are hyperparameters.

Fitting Loss. To project a contact map to a object pose, we minimize the L2 distance between the human vertices and the object vertices indicated by the contact map,

$$E_{\rm fit} = \sum_{i=1}^{K} \|\boldsymbol{v}_{\boldsymbol{o}_1}[d_{\boldsymbol{o}}^i] - \boldsymbol{v}_{\boldsymbol{h}_1}[d_{\boldsymbol{h}}^i]\|_2.$$
(10)

Contact Loss. We leverage a contact loss to encourage body part to contact the object surface,

$$E_{\text{cont}} = \sum_{d_h} \min_{d_o} \| \boldsymbol{v}_{o_1}[d_o] - \boldsymbol{v}_{h_1}[d_h] \|_2.$$
(11)

where d_h here indicates the vertices index in the certain body part. **Penetration Loss.** Given the signed-distance field of the human pose \mathbf{sdf}_{h_1} , we employ a penetration loss to penalize the body-object interpenetration,

$$E_{\text{pene}} = -\sum_{d_o} \min(\mathbf{sdf}_{h_1}(\boldsymbol{v}_{o_1}[d_o]), 0).$$
(12)

The metric for determining the final pose selection is given by the expression $\mathbb{1}(E_{\text{pene}} = 0)/E_{\text{cont}}$. We sample one pose from the set of poses generated by all contact maps, based on the metric.

Learning-based Interaction Retrieval. Our interaction retrieval can also be achieved by integrating knowledge from several learning-based algorithms. Although being more complicated, the retrieval can be done without handcraft rules. Our pipeline can be divided into followings. (i) Given the text prompt tand the initial human pose a_1 , we synthesize corresponding images via Stable Diffusion [81]. (ii) We follow [28] filter out images with low quality in interaction (iii) An off-the-shelf model LEMON [120] is employed to obtain object affordance and human contact, given the generated image paired with human pose a_1 and object template. The output $\{(l_h^i)_{i=1}^M, (l_o^i)_{i=1}^N\}$ indicates the contact vert indexes of human and object respectively, and the output T_1 indicates the estimated object translation, which is used for initialization in the optimization. (iv) To acquire the object pose, we utilize the optimization to minimize the Chamfer distance between the human vertices and the object vertices, indicated by the contact vertices obtained in the last step.

$$E_{\rm fit} = \sum_{j} \min_{k} \| \boldsymbol{v}_{\boldsymbol{o}_1}[l_o^k] - \boldsymbol{v}_{\boldsymbol{h}_1}[l_h^j] \|_2.$$
(13)

24 S. Xu, Z. Wang, Y.-X. Wang, L.-Y. Gui

B.2 World Model

In Sec. 3.3 of the main paper, we introduce the input representation for the dynamics model. The architecture of the world is also introduced in detail. In this section, we introduce detail with particular emphasis on scenarios in which the initial state s_1 is presented.

In the particular instance where the time step t = 1, the state vector s_1 encapsulates a single frame. Consequently, we employ two distinct models for dynamics prediction. For predictions originating from the initial state, the history motion encompasses a single time step H = 1. In contrast, for predictions for subsequent states, the historical interval covering m time steps, where m denoting the frame count per segment.

C Additional Details of Experimental Setup



Fig. A: We use Amazon Mechanical Turk [96] to build an annotation platform. We provide instructions to guide the annotator to split a long sequence into several short sub-sequences with their start and end frames, and then annotate each sub-sequence. We inform annotators that our collected data are used for text-motion generation when they accept this work.

Datasets. We have included a screenshot of our annotation platform in Fig. A. Upon acceptance of our work, we plan to update the Acknowledgements section with the identities of the annotators. Our annotations are further diversified by the GPT-4 [67]. The prompt used for this purpose is: I'm going to give you an essay, and I want you to give me three sentences in English of varying degrees of complexity, like the following example: "A person lifts something with both hands, dumps it to the right, and then puts it back down." A person reaches to the left for something, then reaches to the right as if to dump it, and then

puts it back to the left. A person picks up something with both hands, tilts it, and then puts it down. To make the text more complex, you can add more detailed adjectives and adverbs and complicate sentence structure and verbs. The input text is "A man leaned his upper right thigh against a small table and stretched his left hand and left foot outward." Please give me three texts that vary in complexity but keep the meaning of the sentence the same.

Metrics. In Sec. 4.1, we introduce the metrics employed in this paper. This section details the formula for the metric **CMD** proposed. The formulations for other metrics are available in the existing literature [25,116]. CMD quantifies the discrepancy between the contact maps of ground truth interactions and those synthesized one. In this context, a contact map is characterized by the proportion of time each body part $\{p_i\}_{i=1}^{P}$ maintains active contact. Here, p_i denotes the percentage of time during which the body part *i* is less than a threshold distance from the object. And the metric is defined as,

$$CMD = \frac{1}{P} \sum_{i=1}^{P} \|\boldsymbol{p}_i - \boldsymbol{p}_i^{\text{GT}}\|_1, \qquad (14)$$

where p_i^{GT} is from the ground truth contact map, P is the number of the body parts defined in the SMPL [58], and we define the distance threshold as 0.03 m. **Implementation Details.** The segment in the MDP contains m = 4 frames. The dynamics model, which includes 2 dynamics blocks as described in the main paper, is trained on the BEHAVE training set [7], with a batch size of 32, a latent dimension of 64, and for 500 epochs. For rollout after the initial step t > 1, our dynamics model is trained to predict over a longer timeframe ($F = 3 \times m = 12$), exceeding the past motion duration (H = m = 4). For the initial step t = 1, we train a separate dynamics model to forecast a duration of F = 15 given the past motion over H = 1 frame, consistent with Sec. B.2. The optimization process is conducted over 300 epochs, utilizing a learning rate of 0.01.

D Additional Qualtitative Results

User Study. We conduct a double-blind user study. Given two annotations, we use our dynamics model with InterDiff [116] and with our proposed vertex control to generate 8 samples respectively. We design pairwise evaluations. Considering one pair from ours and another one from InterDiff, human judges are asked to determine which interaction has the better interaction quality. From the results of 11 human evaluations in Table A, our approach has a success rate of 61.4% against the baseline.

Interaction Retrieval. In addition to the results in the main paper and the demo video, we here visualize the intermediate retrieval results. Fig. **B** depicts both the handcraft retrieval and learning-based retrieval procedure, resulting in a diverse set of interactions that are both high-quality and semantically aligned.

25

26 S. Xu, Z. Wang, Y.-X. Wang, L.-Y. Gui



Fig. B: Qualitative results from the interaction retrieval. We demonstrate that our proposed handcraft interaction retrieval (unified color person) and learning-based interaction retrieval (textured person, where textures are from [12]) can extract diverse and realistic interactions.

Table A: User study on the zero-shot text-to-interaction task. Under pairwise human voting results, our generated results significantly outperform the baseline considering the interaction fidelity.

Method	w/ InterDiff [116] vs. w/	vertex control (ours)
w/ InterDiff [116] vs.	N/A	38.6%
w/ vertex control (ours)	61.4%	N/A

E Potential Negative Societal Impact

Some potential negative societal impacts include: (i) Our approach can be used to synthesize realistic human motion interacting with objects, which could potentially lead to the creation of misinformation. (ii) Our approach evaluates on real behavioral information, which may raise privacy concerns. However, our model utilizes a processed representation (SMPL [58]) of the human motion that retains minimal identifying details, in contrast to raw data or images. This aspect can be positively regarded as a feature that enhances privacy.