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

Sirui Xu[†] **Ziyin Wang**[†] **Yu-Xiong Wang**[‡] **Liang-Yan Gui**[‡]
University of Illinois Urbana-Champaign
[†] Equal Contribution [‡] Equal Advising
{siruixu2, ziyin, yxw, lgui}@illinois.edu
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 text-interaction 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 pre-trained 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 low-level interaction dynamics. To overcome this issue, we 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 without relying on paired text-interaction data. We apply InterDreamer to the BEHAVE, OMOMO, 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 has made unprecedented progress through advancements in diffusion models [41, 105, 106, 131], leading to synthesis outcomes that are realistic, diverse, and controllable. This progress has ignited an increased interest in exploring expanded tasks related to text-guided human interaction generation, such as social interaction [69] and human-scene interaction [44]. However, many of these explorations are limited in that the dynamics of objects is not involved or text-guided. Aiming to bridge such a gap, this paper tackles a more challenging task – *generating versatile 3D human-object interactions (HOIs) through language guidance*, as illustrated in Figure 1.

Although a direct solution, as suggested by the concurrent work [28, 65, 91, 107, 136, 140], 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. As can be observed, generating social or scene interactions is heavily dependent on extensive collections of text-interaction pair data [34, 69, 83, 130], and scaling these methods to address more complex HOIs outlined in our study could require datasets of comparable magnitude. Achieving this goal appears unattainable by merely annotating existing 3D HOI datasets [7, 30, 45, 47, 53, 66, 169, 170, 177, 180], which are relatively limited in

38th Conference on Neural Information Processing Systems (NeurIPS 2024).

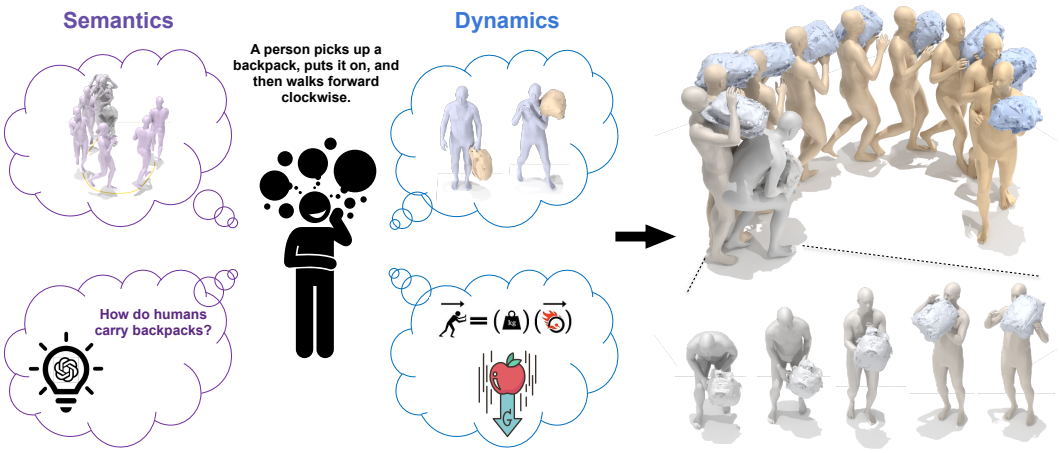

Figure 1: InterDreamer generates vivid 3D human-object interaction sequences guided by text descriptions, by synergizing semantics and dynamics knowledge from large-scale text-motion data (upper left), a large language model (bottom left), human-object interaction data (upper middle), and prior knowledge (bottom middle) from simple physics. We visualize the generated text-guided interaction sequence (upper right), with the beginning of the sequence unfolded (bottom right).

size. Although recent studies [28, 65, 91, 155] have annotated some of these datasets, the volume of text-interaction pairs still lags behind that available for existing text-driven motion generation efforts.

An intriguing question naturally arises: given the limited annotations of the text, *what is the potential of learning for text-conditioned HOI generation without text supervision*, which is the main focus of this paper. However, formulating the task in such a setting presents significant challenges, primarily due to the inability to directly learn the alignment between text and HOI dynamics. Our key observation 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 – is 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 (Figure 1), both of which do not necessarily require learning from text-interaction pairs, if they are decoupled.

The semantics of interaction, although not available through direct supervised training, can be harnessed from *prior knowledge* without text-interaction pair datasets. Specifically, to acquire semantically aligned interaction, we first consult a large language model (LLM), such as GPT-4 [88] and Llama 2 [120], to provide understanding 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* [134]. The intermediate thoughts and the final thought are then used to (**i**) generate 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 human pose and text description, following a philosophy similar to *retrieval-augmented generation* [62].

While these large models can offer high-level motion semantic 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 *world model*, which predicts the subsequent state of an object affected by the interaction. The key here is to reach *generalizable representations* in different motion and objects. 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 simple dynamics from a publicly available 3D HOI dataset BEHAVE [7], and generalize naturally to other datasets [47, 66]. The plausibility of the generated interaction is further enhanced by a subsequent optimization procedure on the synthesized human and object motion.

To summarize, our contributions are: (**i**) We address the task of synthesizing whole-body interactions with dynamic objects guided by textual commands, achieving this notably without the need for paired text-interaction data, a novel paradigm to the best of our knowledge. (**ii**) We introduce a framework that decomposes semantics and dynamics, and they can be easily integrated. Specifically, it harnesses knowledge from a large language model (LLM) and a text-to-motion model as external resources, alongside our proposed world model. Remarkably, the only component that requires additional training is the world model, which highlights the *ease of use* of our framework. (**iii**) Experimental results demonstrate that our 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, 36, 61, 93], past motion [5, 17, 86, 110, 149, 150, 163], trajectories [31, 50, 51, 100, 122, 145], scene context [12, 29, 39, 44, 113, 114, 125–127, 130, 153, 175, 182, 183], and without condition [96]. Recently, human motion synthesis guided by textual descriptions [1, 6, 18, 24, 26, 34, 35, 49, 54, 57, 64, 71, 77, 81, 94, 95, 97, 104, 116, 133, 160, 162, 168, 172, 174, 178, 179, 181, 185, 187] is popular and extended to various applications, including the text-conditioned generation of multiple-person [33, 43, 63, 75, 132] and human-scene interaction [14, 21, 44, 48]. Our goal is to model human and object dynamics concurrently guided by text.

Human-Object Interaction Generation. Synthesizing hand-object interactions [11, 15, 20, 68, 73, 79, 80, 119, 137, 161, 165, 167, 184, 186] and single-frame human-object interactions [25, 42, 56, 92, 128, 143, 152, 154, 166] are popular topics and extended to zero-shot settings [52, 67, 156, 157]. Recently, researchers explore whole-body dynamic interaction generation, in kinematic-based approaches [22, 27, 32, 38, 58–60, 66, 84, 85, 99, 107–109, 111, 123, 136, 139, 140, 148, 151, 176] and physics-based approaches [4, 8, 16, 23, 40, 72, 74, 78, 87, 89, 115, 117, 124, 129, 142, 146, 147, 158]. 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, 30, 45, 47, 53, 66, 112, 138, 144, 155, 164, 169, 170, 180] provide the groundwork for research in this area, and concurrent efforts [28, 65, 91] demonstrate the feasibility of applying supervised learning methods via annotating datasets. However, the amount of data currently available fall short when compared to more extensive text-motion datasets [34, 70, 83]. 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 [128] and generate human motion [3, 46, 159, 178] as well as human-scene interactions [141]. Our approach is inspired by [128], which uses LLMs to infer contact body parts with a given object for reconstructing 3D human-object interactions – a task different from ours. Our approach utilizes GPT-4 [88] or Llama 2 [120], to not only understand contact body parts but also narrow the distribution gap between different tasks, and provide knowledge for interaction retrieval. This is accomplished by utilizing the in-context learning capabilities of LLMs [22] and their support for retrieval-augmented generation [62].

3 Methodology

Problem Formulation. Our goal is to synthesize a sequence of 3D human-object interactions $\boldsymbol{x}$ that satisfies a descriptive text p. This sequence is a series of tuples $[(\boldsymbol{h}_1, \boldsymbol{o}_1), (\boldsymbol{h}_2, \boldsymbol{o}_2), \ldots, (\boldsymbol{h}_M, \boldsymbol{o}_M)]$, where $\boldsymbol{h}_i$ represents the human pose parameters defined in the SMPL model [76], while the shape of the human is unified the same as [34]. $\boldsymbol{o}_i$ defines the pose of the rigid object in terms of its 3D spatial position 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 supervision for training.

Overview. Our framework, illustrated in Figure 2, can be conceptualized as a Markov decision process (MDP). We begin by dividing the motion sequence into T segments, each with m frames,

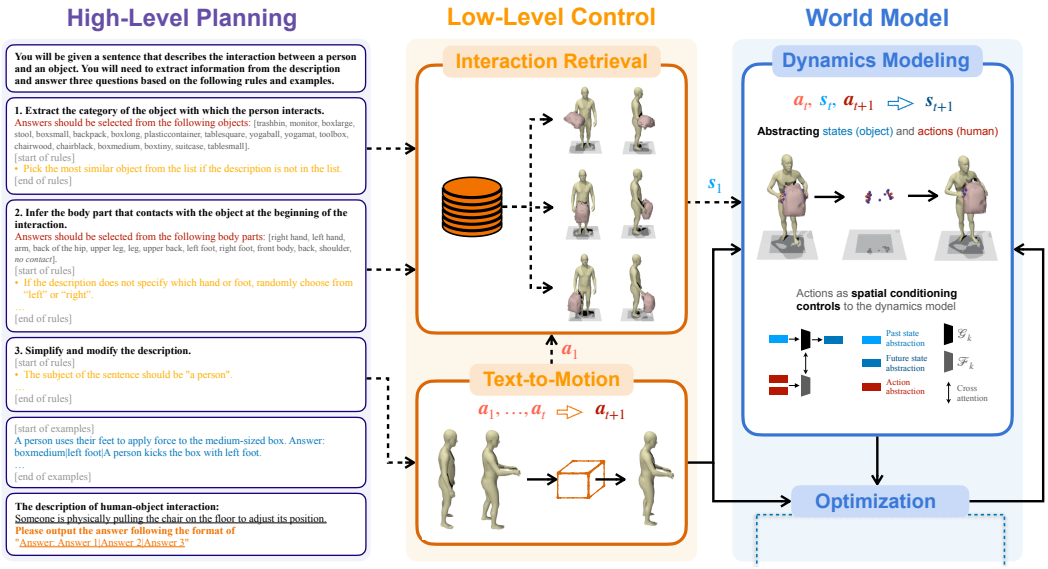

Figure 2: **An overview of our InterDreamer. (i)** Our high-level planning analyzes the description using LLMs and provides guidance to the low-level control. **(ii)** Our low-level control includes a text-to-motion model that translates text into human actions $\boldsymbol{a}_{t+1}$, and an interaction retrieval model that extracts the object's first state $\boldsymbol{s}_1$. **(iii)** Our world model executes actions to output the next state $\boldsymbol{s}_{t+1}$. We achieve this by abstracting the problem as predicting the motion of contact vertices – represented by red spheres for humans and blue spheres for objects on the top right – using human vertices as controls for the prediction of object vertices. An optimization process is coupled with the dynamics model, projecting the state and action onto valid counterparts. Solid arrows mean that the process is performed iteratively.

where $M = T \times m$. Object motion $\{\boldsymbol{o}_i\}_{i=1}^{M}$ can be seen as a sequence of environmental states $\{\boldsymbol{s}_t\}_{t=1}^{T}$, and human motion $\{\boldsymbol{h}_i\}_{i=1}^{M}$ is described as a sequence of actions $\{\boldsymbol{a}_t\}_{t=1}^{T}$ that interact with the environment. Under such an MDP setup, our framework starts with high-level planning L, which deciphers textual interaction description p by $g = L(p)$ (Sec. 3.1). Then, a text-to-motion model π translates context g into human actions, modeled as $\boldsymbol{a}_{t+1} \sim \pi(\boldsymbol{a}_{t+1}|\boldsymbol{s}_t, \{\boldsymbol{a}_i\}_{i=1}^{t}, g)$ (Sec. 3.2). The interaction retrieval R proposes an initial object state $\boldsymbol{s}_1 \sim R(\boldsymbol{s}_1|\boldsymbol{a}_1, g)$, based on the initial action $\boldsymbol{a}_1$ and context g (Sec. 3.2). After that, a world model P is trained to predict future states $\boldsymbol{s}_{t+1} \sim P(\boldsymbol{s}_{t+1}|\boldsymbol{a}_t, \boldsymbol{s}_t, \boldsymbol{a}_{t+1})$ from the current action and state (Sec. 3.3). Our world model incorporates an optimization process, 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

Leveraging LLMs' strong reasoning capabilities and inherent common sense, our high-level planning L yields interaction details $g = L(p)$ that cannot be naïvely extracted in textual descriptions p. The process undertaken by L encompasses three steps: **(i)** *Determining the object*: the LLM is employed to translate described objects into corresponding categories from a predefined list. **(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 [76]. And most importantly, **(iii)** *reducing the distribution gap*: the LLM bridges the distribution gap between the free-form textual input and the language used within the training data of the text-to-motion model [34]. This involves standardizing syntax and content according to designed guidelines. In Figure 2, we demonstrate the prompt we used with the few-shot prompting [10]. We define intermediate thoughts and the final thought, *i.e.*, answers to three questions, as detailed information $g = L(p)$, which guides the subsequent procedure, structuring the entire framework with a philosophy similar to retrieval-augmented generation [62]. Our high-level planning operates indirectly in the generation of interactions. Nonetheless, it narrows the vast range of possible interactions in the real world into a more manageable distribution within the capabilities of our framework. We incorporate GPT-4 [88] and Llama-2 [120] for evaluation.

3.2 Low-Level Control

With the information g derived from the description p, the low-level control aims to create a sequence of human actions $\{\boldsymbol{a}_t\}_{t=1}^T$ by a text-to-motion model, and an initial state $\boldsymbol{s}_1$ by interaction retrieval, such that they correspond to the objectives outlined by g.

Text-to-Motion. We utilize a text-to-motion model π to develop actions to be executed in the world model. At each timestep t, π receives the sequence of previous actions $\{\boldsymbol{a}_i\}_{i=1}^t$ and the text tokens encoded from the rewritten description in $g = L(p)$, and produces a next action $\boldsymbol{a}_{t+1}$, which later in Sec. 3.4 will be adjusted through an optimization process that intertwines actions with the object state $\boldsymbol{s}_t$. Thus, the overall process can be formally defined as $\boldsymbol{a}_{t+1} \sim \pi(\boldsymbol{a}_{t+1}|\boldsymbol{s}_t, \{\boldsymbol{a}_i\}_{i=1}^t, g)$, while the initial action $\boldsymbol{a}_1 \sim \pi(\boldsymbol{a}_1|g)$ is influenced merely by context g without prior actions or states, which will be used in interaction retrieval. π builds upon existing text-to-motion models, where we evaluate MDM [118], MotionDiffuse [172], ReMoDiffuse [173], and MotionGPT [46].

Interaction Retrieval. The interaction retrieval component R establishes the initial state $\boldsymbol{s}_1 \sim R(\boldsymbol{s}_1|\boldsymbol{a}_1, g)$, based on the initial action $\boldsymbol{a}_1$ generated by the text-to-motion model. We propose a user-friendly pipeline for this purpose built on handcrafted rules. First, we create databases by extracting HOI frames from the training sets of each target datasets — BEHAVE [7], OMOMO [66], and CHAIRS [47]. The indexing key for retrieval is a tuple consisting of the body part in contact and the category of the involved object. Each retrieval value is a per-frame contact map, represented by a list of K vertex pairs $\{(d_h^i, d_o^i)\}_{i=1}^K$. Here, d_h^i refers to the contact vertex on the human surface, while d_o^i refers to the corresponding contact vertex on the object surface. This contact map is linked to its corresponding key, creating a searchable record of interactions. During the inference stage, using the body part and object information provided by the high-level planning (Sec. 3.1), we retrieve all relevant contact maps from the database. We then sample one map $\{(d_h^i, d_o^i)\}_{i=1}^K$ and use it to establish the object state $\boldsymbol{s}_1 \sim R(\boldsymbol{s}_1|\boldsymbol{a}_1, g)$, thus initializing the interaction. Further details including how we ensure consistency between the sampled state and human action are provided in Sec. B.1 of the Appendix. We also discuss an alternative learning-based approach in Sec. B.1.

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 [103, 135] and autonomous driving systems [55], we use it here to generate HOI trajectories. This model, trained on the training set of a 3D HOI dataset such as BEHAVE [7], serves a similar role as a simulator but is much simpler – it takes the preceding object state $\boldsymbol{s}_t$ along with a pair of consecutive actions $\boldsymbol{a}_t$ and $\boldsymbol{a}_{t+1}$, and predicts the subsequent object state $\boldsymbol{s}_{t+1}$. 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 raw actions, states, and object geometry as input. However, this suffers from a severe generalization problem during inference: the dynamics model is likely to encounter human actions and object geometry that do not exist in the training set, since our text-to-motion model is not trained with object interaction data. To overcome this limitation, we instead focus on encoding interactions through the contact vertices on the object, which capture both the action and object geometry, as shown in Figure 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 and geometry details that are irrelevant to the interaction.

Input Representation. Specifically, at each timestep t, we abstract the past actions as H historical vertex trajectories $\{\{\boldsymbol{v}_i^j\}_{j=1}^N\}_{i=1}^H$, and the future actions as $F = m$ future vertex trajectories $\{\{\boldsymbol{v}_i^j\}_{j=1}^N\}_{i=H+1}^{H+F}$, where non-fixed variable N is the number of sampled contact vertices, and m is the length of segments as mentioned in the overview of Sec 3. Note that we train our dynamics model to forecast over a longer duration than the past motion ($F > H$), only the foremost future action will be used for autoregressive generation during the inference, as suggested in [19]. 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 $\boldsymbol{s}_t$. We then sample vertices that meet the following criteria: $|\mathbf{sdf}_i(\boldsymbol{v}_i^j)| \leq \delta_1, \mathbf{sdf}_i(\boldsymbol{v}_i^k)| \leq \delta_1, \forall i = 1, \ldots, H, \forall j$, and $\|\boldsymbol{v}_i^j - \boldsymbol{v}_i^k\| \geq \delta_2, \forall j \neq k$, where δ_1 and δ_2 are two hyperparameters. The objective is to sparsely sample contact vertices while ensuring that they are sufficient to encompass the interaction. We characterize each vertex trajectory $\{\boldsymbol{v}_i^j\}_{i=1}^{H+F}$

Table 1: **Quantitative results** on evaluating the dynamics model. Our dynamics model with vertex-based action generates interactions of the best quality.

Methods	Text-to-Interaction		Interaction Prediction [148]		
	CMD ↓	Pene. (10^{-2}%) ↓	Trans. Err. (mm) ↓	Rot. Err. (10^{-3} rad) ↓	Pene. (10^{-2}%) ↓
w/o action	0.424	533	123	256	228
contact markers as action (InterDiff [148])	0.219	484	123	226	164
human motion as action	0.325	957	129	265	218
contact vertices as action (**ours**)	**0.151**	**443**	**119**	**221**	**156**

Table 2: **Quantitative results** on human motion quality given our annotation on the BEHAVE [7] dataset. We show that our high-level planning effectively adapts single human generators into human-object interaction generation. To evaluate R-Precision, a batch size of 16 is selected.

Methods	Planning (**Ours**)	R-Precision↑			FID↓	MM Dist↓	Multimodality↑	Diversity→
		Top 1	Top 2	Top 3				
Ground Truth	-	$0.237^{\pm0.004}$	$0.392^{\pm0.004}$	$0.496^{\pm0.005}$	$0.024^{\pm0.000}$	$4.259^{\pm0.006}$	-	$6.510^{\pm0.227}$
MDM [118]	×	$0.153^{\pm0.016}$	$0.279^{\pm0.026}$	$0.398^{\pm0.016}$	$12.279^{\pm0.217}$	$5.351^{\pm0.057}$	$\mathbf{7.604}^{\pm0.190}$	$7.598^{\pm0.334}$
	✓	$\mathbf{0.163}^{\pm0.010}$	$\mathbf{0.307}^{\pm0.043}$	$\mathbf{0.402}^{\pm0.019}$	$\mathbf{10.374}^{\pm0.304}$	$\mathbf{5.303}^{\pm0.117}$	$7.281^{\pm0.083}$	$\mathbf{7.471}^{\pm0.427}$
MotionDiffuse [172]	×	$0.205^{\pm0.011}$	$0.351^{\pm0.002}$	$0.458^{\pm0.021}$	$10.208^{\pm0.500}$	$4.837^{\pm0.064}$	$4.520^{\pm0.163}$	$7.323^{\pm0.412}$
	✓	$\mathbf{0.216}^{\pm0.032}$	$\mathbf{0.369}^{\pm0.023}$	$\mathbf{0.472}^{\pm0.027}$	$\mathbf{9.015}^{\pm0.403}$	$\mathbf{4.649}^{\pm0.029}$	$\mathbf{4.991}^{\pm0.172}$	$\mathbf{7.295}^{\pm0.501}$
ReMoDiffuse [173]	×	$0.196^{\pm0.009}$	$0.338^{\pm0.011}$	$0.448^{\pm0.012}$	$6.385^{\pm0.201}$	$4.855^{\pm0.029}$	$5.889^{\pm0.524}$	$\mathbf{7.160}^{\pm0.306}$
	✓	$\mathbf{0.223}^{\pm0.006}$	$\mathbf{0.368}^{\pm0.015}$	$\mathbf{0.482}^{\pm0.011}$	$\mathbf{5.237}^{\pm0.174}$	$\mathbf{4.784}^{\pm0.053}$	$\mathbf{6.350}^{\pm0.411}$	$7.201^{\pm0.318}$
MotionGPT [46]	×	$0.233^{\pm0.003}$	$0.344^{\pm0.004}$	$0.457^{\pm0.005}$	$5.497^{\pm0.106}$	$5.205^{\pm0.027}$	$1.062^{\pm0.211}$	$8.316^{\pm0.204}$
	✓	$\mathbf{0.234}^{\pm0.004}$	$\mathbf{0.387}^{\pm0.003}$	$\mathbf{0.471}^{\pm0.007}$	$\mathbf{4.751}^{\pm0.121}$	$\mathbf{4.995}^{\pm0.003}$	$\mathbf{1.337}^{\pm0.193}$	$\mathbf{7.106}^{\pm0.487}$

with a feature $\boldsymbol{f}^j$ to provide (i) human vertex coordinates at T-pose, providing information about the position of the human vertex on the body surface; (ii) the vertex-to-object surface vector, indicating vertex's impact on the object as well as inherently including information related to the object's shape; and (iii) the vertex's velocity relative to its nearest object vertex. Thus, the model needs to learn how the features of human action $\boldsymbol{f}^j$ affect the evolution of the state of the object.

Architecture. As demonstrated in Figure 2, the network comprises two components: $\mathcal{G}$ that operates without contact vertex conditions, applicable in scenarios where no contact occurs, and $\mathcal{F}$, which incorporates contact vertex conditions into the object trajectory when contact is present. The k-th layer of $\mathcal{G}$ can be initiated as $\mathcal{G}_k(\boldsymbol{x}_k, \boldsymbol{\Theta})$, mapping the input feature map $\boldsymbol{x}_k$ at the k-th layer to another feature map, with Θ denoting the MLP's parameters. To incorporate human vertex controls, we introduce a second network $\mathcal{F}_k(\boldsymbol{y}_k^j, \boldsymbol{\Theta}_v)$ 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, a dynamics block is formulated as: $\boldsymbol{x}_{k+1}, \{\boldsymbol{y}_{k+1}^j\}_{j=1}^N = \mathrm{Attn}(\mathcal{G}_k(\boldsymbol{x}_k, \boldsymbol{\Theta}), \{\mathcal{F}_k(\boldsymbol{y}_k^j, \boldsymbol{\Theta}_v)\}_{j=1}^N)$. We stack multiple dynamics blocks to form the 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 is preliminary and subject to further optimization as introduced in Sec. 3.4, which will yield the final future state. We utilize the Mean Squared Error loss to train the dynamics model. For more explanations, please refer to Sec. B.2 of the Appendix.

3.4 Optimization

Optimization plays a role in introducing prior knowledge and avoiding the accumulation of errors. During inference, we input the action $\boldsymbol{a}_{t+1}$ and state $\boldsymbol{s}_{t+1}$ and refine them. 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 optimized results with their preliminary one, a velocity loss for temporal smoothness, a contact loss to promote occurring contact, and a collision loss to reduce penetration. We provide detailed formulations in Sec. B.3 of the Appendix.

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 (Sec. 4.3) results for our approach. Additionally, we perform ablation studies to verify the efficacy of each component within our framework. These studies also cover the

Table 3: **Quantitative results** on human motion quality on the OMOMO [66] dataset with their provided annotation. We show that our high-level planning narrows the distribution gap and adapts single human generators into human-object interaction generation. To evaluate R-Precision, a batch size of 32 is selected.

Methods	Planning (**Ours**)	R-Precision↑			FID↓	MM Dist↓	Multimodality↑	Diversity→
		Top 1	Top 2	Top 3				
Ground Truth	-	$0.044^{\pm0.004}$	$0.095^{\pm0.008}$	$0.151^{\pm0.009}$	$0.000^{\pm0.000}$	$6.858^{\pm0.006}$	-	$5.660^{\pm0.110}$
MDM [118]	×	$0.056^{\pm0.005}$	$0.096^{\pm0.007}$	$0.135^{\pm0.006}$	$16.638^{\pm0.631}$	$7.110^{\pm0.063}$	$2.446^{\pm0.456}$	$\mathbf{5.862}^{\pm0.520}$
	✓	$\mathbf{0.062}^{\pm0.006}$	$\mathbf{0.109}^{\pm0.004}$	$\mathbf{0.155}^{\pm0.008}$	$\mathbf{15.735}^{\pm0.285}$	$\mathbf{6.889}^{\pm0.082}$	$\mathbf{2.663}^{\pm0.520}$	$6.461^{\pm0.841}$
MotionDiffuse [172]	×	$0.048^{\pm0.006}$	$0.094^{\pm0.008}$	$0.143^{\pm0.013}$	$15.442^{\pm0.231}$	$\mathbf{5.799}^{\pm0.054}$	$1.658^{\pm0.209}$	$5.981^{\pm0.516}$
	✓	$\mathbf{0.075}^{\pm0.005}$	$\mathbf{0.141}^{\pm0.015}$	$\mathbf{0.189}^{\pm0.009}$	$\mathbf{10.815}^{\pm0.093}$	$5.916^{\pm0.094}$	$\mathbf{1.677}^{\pm0.264}$	$\mathbf{5.718}^{\pm0.522}$
ReMoDiffuse [173]	×	$0.062^{\pm0.003}$	$0.111^{\pm0.005}$	$0.160^{\pm0.012}$	$15.479^{\pm0.209}$	$5.690^{\pm0.049}$	$1.179^{\pm0.145}$	$6.032^{\pm0.540}$
	✓	$\mathbf{0.067}^{\pm0.004}$	$\mathbf{0.127}^{\pm0.006}$	$\mathbf{0.174}^{\pm0.006}$	$\mathbf{14.560}^{\pm0.080}$	$\mathbf{5.678}^{\pm0.033}$	$\mathbf{1.193}^{\pm0.202}$	$\mathbf{5.368}^{\pm0.417}$
MotionGPT [46]	×	$0.061^{\pm0.005}$	$0.114^{\pm0.006}$	$0.152^{\pm0.006}$	$18.472^{\pm0.528}$	$6.358^{\pm0.076}$	$\mathbf{4.553}^{\pm0.244}$	$\mathbf{6.726}^{\pm0.156}$
	✓	$\mathbf{0.064}^{\pm0.007}$	$\mathbf{0.121}^{\pm0.007}$	$\mathbf{0.164}^{\pm0.009}$	$\mathbf{17.512}^{\pm0.498}$	$\mathbf{6.287}^{\pm0.041}$	$4.470^{\pm0.191}$	$7.048^{\pm0.169}$

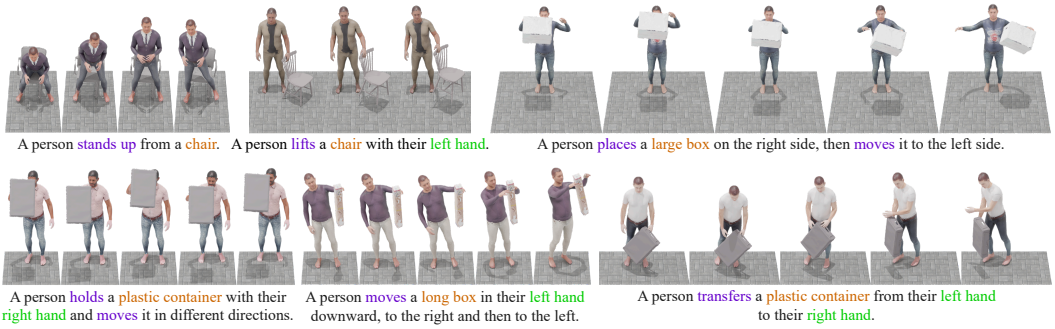

Figure 3: **Qualitative results** on free-form text input. The interaction sequences, with textures from [13], are presented through a time-series visualization.

interaction prediction task [148] to evaluate our dynamics model. Additional details and results are presented in Sec. C and Sec. D of the Appendix. Please refer to our website for video results.

4.1 Experimental Setup

Datasets. We use BEHAVE [7], CHAIRS [47], and OMOMO [66] datasets for quantitative evaluation. The BEHAVE dataset includes recordings of 8 individuals interacting with 20 everyday objects, and our analysis focuses on 18 objects for which interaction sequences are available at 30 Hz. The human pose is modeled using SMPL-H [102], with hand poses set to an average pose *due to the absence of detailed hand pose in the dataset*. 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

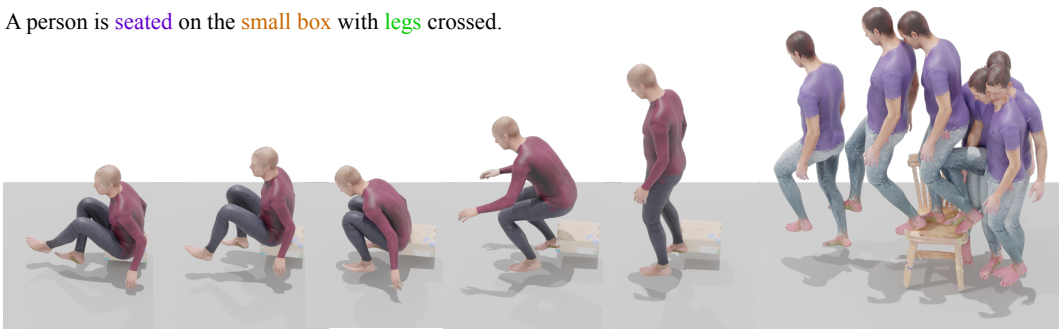

Figure 4: **Qualitative results** in more challenge scenarios with *free-form input not* from our annotations, showing the ability of our InterDreamer to fit *object sizes* and handle *complex and long sequences*. Here, our synergized models are GPT-4 [88] and MotionGPT [46].

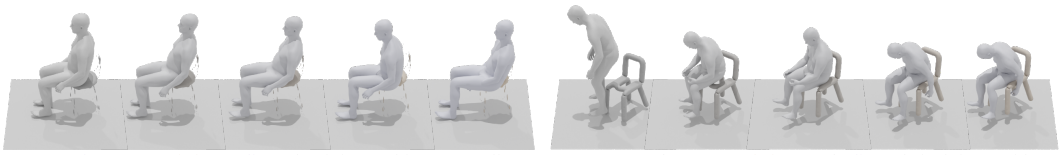

Figure 5: **Qualitative results** on the CHAIRS [47] dataset. Our dynamics model trained on the BEHAVE [7] dataset generalizes well on the CHAIRS objects unseen in training. Frames are separately visualized. Here, our synergized models are GPT-4 [88] and MotionGPT [46].

Figure 6: Results from the interaction retrieval. We demonstrate that our proposed retrieval approach based on handcraft rules can extract diverse and realistic interactions.

subsequences for evaluation. For qualitative evaluation, we go beyond using annotations specifically created and employ free – form text to demonstrate our model's capability on out-of-distribution inputs. To assess our model's performance with novel objects, we expand our retrieval database to include objects from the OMOMO [66] and CHAIRS [47] datasets, while we do not fine-tune the dynamics model on them–a direct qualitative evaluation without additional adaptation.

Metrics. The evaluation metrics are divided into three categories: (**i**) *Human motion quality*: The Fréchet Inception Distance (**FID**) measures the distance between the generated motion 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 space. MultiModal distance (**MM Dist**) is the distance between the motion feature and the text feature. We follow [34] to generate motion and text features. (**ii**) *Interaction quality*: We propose **CMD** to measure the distance between contact maps of real interactions and those generated. 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 the Appendix. We also measure the collision (**Pene.** [148]), which calculates the average percentage of object vertices that have non-negative values in the human signed distance fields [90]. (**iii**) *Object motion accuracy*: The dynamics model's performance in the interaction prediction task [148] is evaluated by the accuracy of predicted object motion, including **Trans. Err.**, the average distance between predicted and ground truth, and **Rot. Err.**, the average distance between the predicted and ground truth.

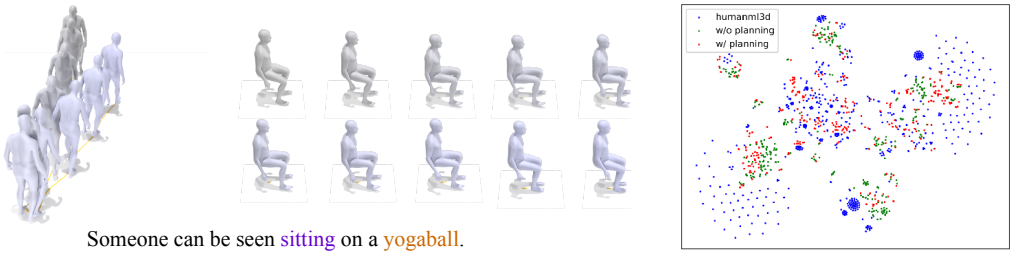

(a) Low-level control w/o planning v.s. w/ planning

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

Figure 7: (**a**) **Ablation study** on the high-level planning. On the *left* are results from MotionGPT [46] using free-form descriptions, and on the *right* are results with our planning additionally. Without planning, the motion generative model struggles to interpret free-form HOI descriptions and generate semantically-aligned motion. (**b**) We visualize CLIP [98] features of text on HumanML3D [34] via t-SNE [82], raw HOI descriptions ("w/o planning"), and HOI descriptions processed through our high-level planning ("w/ planning"). See Table 5 for quantitative measurements.

(a) human motion as action (b) contact vertices as action

Figure 8: **Ablation study** on the dynamics model. 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.

Baselines. Most recent research on text-to-HOI synthesis follows a supervised learning approach [28, 91], making direct quantitative comparisons unfair. Therefore, we primarily focus on qualitative comparisons with these methods. To enable quantitative evaluation, we develop a range of baselines to assess both the overall performance of our framework and the effectiveness of its individual components. In the context of high-level planning, we utilize GPT-4 [88] and Llama-2 [120], illustrating the effectiveness of our prompts across different language models. For low-level motion generation control, our baselines include MDM [118], MotionDiffuse [172], ReMoDiffuse [173], and MotionGPT [46], which span a range of text-to-motion approaches trained on HumanML3D [34] and show the generalizability of our framework. To evaluate the dynamics model, we include different design choices: (**i**) unconditional dynamics model which operates object dynamics independently of human motion; (**ii**) using human marker features as actions to the dynamics model, similar to [148]; (**iii**) using unprocessed human motion and object pointcloud motion as input to the dynamics model; (**iv**) our proposed vertex-based actions where only the contacting vertices are used for control.

4.2 Quantitative Results

In Table 1, comparing our framework to baselines with unconditional dynamics model, 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 for action features, our method demonstrates enhanced performance by offering more fine-grained guidance and extracting generalizable features for dynamics modeling. Tables 2 and 3 present a comparative analysis of our approach of combining high-level planning with low-level control, where we adopt various text-to-motion models against their counterparts without high-level planning on the BEHAVE and OMOMO datasets. Our approach consistently outperforms baselines. Specifically, InterDreamer exhibits superior motion quality, reflected by a lower FID, higher R-Precision, and better diversity, highlighting the benefits of incorporating our planning to reduce the distribution gap for the motion generator to generalize in the HOI synthesis task.

4.3 Qualitative Results

Figure 3 displays several results guided by the free-form text. 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, as illustrated in Figure 4, we selectively use more complex sequences of interactive descriptions that are *beyond the scope of the existing HOI dataset*. Figure 5 further exemplifies our method that is able to generalize effectively to the CHAIRS dataset, despite our dynamics model not being trained on it. Figure 6 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 the Appendix.

Table 4: **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 ↑	Q1 Acc* ↑	Q2 Acc ↑	Q2 Acc* ↑
GPT-4 [88]	0.801	0.997	0.703	0.964
Llama-2 (7B) [120]	0.073	0.147	0.436	0.689
Llama-2 (13B) [120]	0.232	0.319	0.662	0.853
Llama-2 (70B) [120]	0.722	0.967	0.798	0.907

Table 5: **Quantitative comparison** of text similarity. The text processed by high-level planning is more similar to text in HumanML3D [34] on average, while addressing the distributional gap significantly for challenging out-of-distributional descriptions, compared to text without planning.

Sim. to [34]↑	Average	Out-of-Dist.
w/o planning	0.913	0.838
w/ planning	**0.932**	**0.927**

4.4 Ablation Study

Adaptability of High-Level Planning. Is our framework adaptable across different large language models (LLMs)? As illustrated in Table 4, our analysis contains two types of language models: GPT-4 [88], which is accessible through APIs and operates as a black box model; and Llama-2 [120], an open-source model. We demonstrate that language models with large parameters exhibit very high accuracy in responding to questions tailored to our prompts, validating the framework's adaptability.

Effectiveness of High-Level Planning with Low-Level Control. In consistency with Table 2, Figure 7 offers a qualitative comparison of text-to-motion results, contrasting results with and without LLM-revised text descriptions. The comparison shows that motion generated without LLM-enhanced descriptions often fails to correspond to the intended text, if the text is too challenging, *e.g.,* not in the distribution of HumanML3D [34], which is used to train text-to-motion models. This underscores the LLM's critical role in bridging the distribution gap. In Figure 7(b), we visualize the CLIP [98] features of descriptions from HumanML3D, our raw annotations, and those processed by high-level planning. Quantitative evidence is provided in Table 5. Text processed through high-level planning demonstrates greater similarity to the HumanML3D dataset. Additionally, we test on more challenging out-of-distribution text, selecting examples with an average cosine similarity to HumanML3D text of less than 0.85. High-level planning successfully rephrases these texts, significantly increasing their similarity. For example, in Figure 7(a), the text "Someone can be seen sitting on a yoga ball" has a cosine similarity of 0.874 to the closest in-distribution text, while the rephrased text after planning, "A person is seated on an object," achieves a similarity of 0.958.

Effectiveness of World Model. In the quantitative evaluation, we show that the performance of our framework is enhanced by the tailored design of our world model. Table 1 provides additional evidence of the effectiveness by integrating the proposed world model, as interaction correction within the InterDiff framework [148] in the interaction prediction task. This implementation demonstrates enhanced conditionality in the object dynamics modeling across two tasks, attributed to the vertex-level condition as actions. Doing so effectively removes the whole-body complexity, most of which tends to be irrelevant to the interaction. Figure 8 further indicates that our vertex-based condition can establish consistent interactions over time, while the condition by raw motion is not robust.

5 Conclusion

We tackle the task of text-guided 3D human-object interaction generation, aiming to accomplish this without relying on paired text-interaction data. To this end, we present InterDreamer that decouples interaction dynamics from semantics, formulating the task as retrieval-augmented generation and Markov decision process, 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 effectiveness in this novel task, suggesting its 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 and learn a broader range of skills. The results may not be physically plausible and lead to artifacts in some cases, for example, foot skating. Hand poses are rough because they are missing from the dataset, but could be improved by integrating a physics simulator.

Acknowledgments

We thank Qiusi Zhan for supporting the implementation and evaluation of high-level planning, Jiangwei Yu and Xiyan Xu for their efforts in dataset annotation, and Xiang Li for the discussions. This work was supported in part by NSF Grant 2106825, NIFA Award 2020-67021-32799, the IBM-Illinois Discovery Accelerator Institute, the Toyota Research Institute, and the Jump ARCHES endowment through the Health Care Engineering Systems Center at Illinois and the OSF Foundation. This work used computational resources on NCSA Delta and PTI Jetstream2 through allocations CIS220014, CIS230012, CIS230013, and CIS240311 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, and on TACC Frontera through the National Artificial Intelligence Research Resource (NAIRR) Pilot.

References

[1] Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. In: 3DV (2019) 3

[2] Athanasiou, N., Petrovich, M., Black, M.J., Varol, G.: Teach: Temporal action composition for 3d humans. In: 3DV (2022) 3

[3] Athanasiou, N., Petrovich, M., Black, M.J., Varol, G.: SINC: Spatial composition of 3d human motions for simultaneous action generation. In: ICCV (2023) 3

[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) 3

[5] Barquero, G., Escalera, S., Palmero, C.: BeLFusion: Latent diffusion for behavior-driven human motion prediction. In: ICCV (2023) 3

[6] Barquero, G., Escalera, S., Palmero, C.: Seamless human motion composition with blended positional encodings. In: CVPR (2024) 3

[7] 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) 1, 2, 3, 5, 6, 7, 8, 22, 25

[8] Braun, J., Christen, S., Kocabas, M., Aksan, E., Hilliges, O.: Physically plausible full-body hand-object interaction synthesis. arXiv preprint arXiv:2309.07907 (2023) 3

[9] Brooks, T., Holynski, A., Efros, A.A.: Instructpix2pix: Learning to follow image editing instructions. In: CVPR (2023) 3

[10] 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 few-shot learners. In: NeurIPS (2020) 2, 4

[11] Cao, J., Liu, J., Kitani, K., Zhou, Y.: Multi-modal diffusion for hand-object grasp generation. arXiv preprint arXiv:2409.04560 (2024) 3

[12] Cao, Z., Gao, H., Mangalam, K., Cai, Q.Z., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: ECCV (2020) 3

[13] Casas, D., Comino-Trinidad, M.: SMPLitex: A generative model and dataset for 3d human texture estimation from single image. In: BMVC (2023) 7

[14] Cen, Z., Pi, H., Peng, S., Shen, Z., Yang, M., Zhu, S., Bao, H., Zhou, X.: Generating human motion in 3d scenes from text descriptions. In: CVPR (2024) 3

[15] Cha, J., Kim, J., Yoon, J.S., Baek, S.: Text2HOI: Text-guided 3d motion generation for hand-object interaction. In: CVPR (2024) 3

[16] Chao, Y.W., Yang, J., Chen, W., Deng, J.: Learning to sit: Synthesizing human-chair interactions via hierarchical control. In: AAAI (2021) 3

[17] Chen, L.H., Zhang, J., Li, Y., Pang, Y., Xia, X., Liu, T.: HumanMAC: Masked motion completion for human motion prediction. In: ICCV (2023) 3

[18] 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) 3

[19] 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) 5

[20] Christen, S., Hampali, S., Sener, F., Remelli, E., Hodan, T., Sauser, E., Ma, S., Tekin, B.: Diffh2o: Diffusion-based synthesis of hand-object interactions from textual descriptions. arXiv preprint arXiv:2403.17827 (2024) 3

[21] 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) 3

[22] Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F.: Context-aware human motion prediction. In: CVPR (2020) 3

[23] Cui, J., Liu, T., Liu, N., Yang, Y., Zhu, Y., Huang, S.: AnySkill: Learning open-vocabulary physical skill for interactive agents. In: CVPR (2024) 3

[24] Dabral, R., Mughal, M.H., Golyanik, V., Theobalt, C.: MoFusion: A framework for denoising-diffusion-based motion synthesis. In: CVPR. pp. 9760–9770 (2023) 3

[25] Dai, S., Li, W., Sun, H., Huang, H., Ma, C., Huang, H., Xu, K., Hu, R.: InterFusion: Text-driven generation of 3d human-object interaction. In: ECCV (2024) 3

[26] Dai, W., Chen, L.H., Wang, J., Liu, J., Dai, B., Tang, Y.: Motionlcm: Real-time controllable motion generation via latent consistency model. arXiv preprint arXiv:2404.19759 (2024) 3

[27] Daiya, D., Conover, D., Bera, A.: COLLAGE: Collaborative human-agent interaction generation using hierarchical latent diffusion and language models. arXiv preprint arXiv:2409.20502 (2024) 3

[28] Diller, C., Dai, A.: CG-HOI: Contact-guided 3d human-object interaction generation. In: CVPR (2024) 1, 2, 3, 9, 21

[29] Diller, C., Funkhouser, T., Dai, A.: FutureHuman3D: Forecasting complex long-term 3d human behavior from video observations. In: CVPR (2024) 3

[30] 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) 1, 3

[31] Feng, H., Ma, W., Gao, Q., Zheng, X., Xue, N., Xu, H.: Stratified avatar generation from sparse observations. In: CVPR (2024) 3

[32] Ghosh, A., Dabral, R., Golyanik, V., Theobalt, C., Slusallek, P.: IMoS: Intent-driven full-body motion synthesis for human-object interactions. arXiv preprint arXiv:2212.07555 (2022) 3

[33] 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) 3

[34] 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) 1, 3, 4, 8, 9, 10, 25

[35] 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) 3

[36] 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) 3

[37] Han, S., Joo, H.: CHORUS: Learning canonicalized 3d human-object spatial relations from unbounded synthesized images. In: ICCV (2023) 22

[38] Hassan, M., Ceylan, D., Villegas, R., Saito, J., Yang, J., Zhou, Y., Black, M.J.: Stochastic scene-aware motion prediction. In: ICCV (2021) 3

[39] Hassan, M., Ghosh, P., Tesch, J., Tzionas, D., Black, M.J.: Populating 3d scenes by learning human-scene interaction. In: CVPR (2021) 3

[40] Hassan, M., Guo, Y., Wang, T., Black, M., Fidler, S., Peng, X.B.: Synthesizing physical character-scene interactions. In: SIGGRAPH (2023) 3

[41] Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: NeurIPS (2020) 1

[42] Hou, Z., Yu, B., Tao, D.: Compositional 3d human-object neural animation. arXiv preprint arXiv:2304.14070 (2023) 3

[43] Huang, B., Li, C., Xu, C., Pan, L., Wang, Y., Lee, G.H.: Closely interactive human reconstruction with proxemics and physics-guided adaption. In: CVPR (2024) 3

[44] 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) 1, 3

[45] Huang, Y., Taheri, O., Black, M.J., Tzionas, D.: InterCap: Joint markerless 3D tracking of humans and objects in interaction. In: GCPR (2022) 1, 3

[46] Jiang, B., Chen, X., Liu, W., Yu, J., Yu, G., Chen, T.: MotionGPT: Human motion as a foreign language. In: NeurIPS (2023) 3, 5, 6, 7, 8, 9, 26

[47] 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) 1, 2, 3, 5, 7, 8, 22

[48] 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) 3

[49] Jin, P., Li, H., Cheng, Z., Li, K., Yu, R., Liu, C., Ji, X., Yuan, L., Chen, J.: Local action-guided motion diffusion model for text-to-motion generation. arXiv preprint arXiv:2407.10528 (2024) 3

[50] Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: GMD: Controllable human motion synthesis via guided diffusion models. In: ICCV (2023) 3

[51] Kaufmann, M., Aksan, E., Song, J., Pece, F., Ziegler, R., Hilliges, O.: Convolutional autoencoders for human motion infilling. In: 3DV (2020) 3

[52] 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) 3

[53] 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) 1, 3

[54] Kim, J., Kim, J., Choi, S.: Flame: Free-form language-based motion synthesis & editing. In: AAAI (2023) 3

[55] Kim, S.W., Zhou, Y., Philion, J., Torralba, A., Fidler, S.: Learning to simulate dynamic environments with gamegan. In: CVPR (2020) 5

[56] Kim, T., Saito, S., Joo, H.: NCHO: Unsupervised learning for neural 3d composition of humans and objects. In: ICCV (2023) 3

[57] Kong, H., Gong, K., Lian, D., Mi, M.B., Wang, X.: Priority-centric human motion generation in discrete latent space. In: ICCV (2023) 3

[58] Krebs, F., Meixner, A., Patzer, I., Asfour, T.: The kit bimanual manipulation dataset. In: Humanoids (2021) 3

[59] 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)

[60] Lee, J., Joo, H.: Locomotion-Action-Manipulation: Synthesizing human-scene interactions in complex 3d environments. In: ICCV (2023) 3

[61] Lee, T., Moon, G., Lee, K.M.: Multiact: Long-term 3d human motion generation from multiple action labels. In: AAAI (2023) 3

[62] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., et al.: Retrieval-augmented generation for knowledge-intensive nlp tasks. In: NeurIPS (2020) 2, 3, 4

[63] Li, B., Ho, E.S., Shum, H.P., Wang, H.: Two-person interaction augmentation with skeleton priors. In: CVPR (2024) 3

[64] Li, C., Chibane, J., He, Y., Pearl, N., Geiger, A., Pons-Moll, G.: Unimotion: Unifying 3d human motion synthesis and understanding. arXiv preprint arXiv:2409.15904 (2024) 3

[65] Li, J., Clegg, A., Mottaghi, R., Wu, J., Puig, X., Liu, C.K.: Controllable human-object interaction synthesis. arXiv preprint arXiv:2312.03913 (2023) 1, 2, 3

[66] Li, J., Wu, J., Liu, C.K.: Object motion guided human motion synthesis. ACM Transactions on Graphics (TOG) **42**(6), 1–11 (2023) 1, 2, 3, 5, 7, 8, 22, 25, 26

[67] Li, L., Dai, A.: GenZI: Zero-shot 3d human-scene interaction generation. In: CVPR (2024) 3

[68] 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) 3

[69] 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) 1

[70] 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) 3

[71] Liu, H., Zhan, X., Huang, S., Mu, T.J., Shan, Y.: Programmable motion generation for open-set motion control tasks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1399–1408 (2024) 3

[72] 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) 3

[73] Liu, S., Zhou, Y., Yang, J., Gupta, S., Wang, S.: Contactgen: Generative contact modeling for grasp generation. In: ICCV (2023) 3

[74] Liu, Y., Chen, C., Ding, C., Yi, L.: PhysReaction: Physically plausible real-time humanoid reaction synthesis via forward dynamics guided 4d imitation. arXiv preprint arXiv:2404.01081 (2024) 3

[75] 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) 3

[76] Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: A skinned multi-person linear model. ACM transactions on graphics (2015) 3, 4, 25, 26

[77] 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) 3

[78] Luo, Z., Wang, J., Liu, K., Zhang, H., Tessler, C., Wang, J., Yuan, Y., Cao, J., Lin, Z., Wang, F., et al.: SMPLOlympics: Sports environments for physically simulated humanoids. arXiv preprint arXiv:2407.00187 (2024) 3

[79] Ma, J., Chen, X., Bao, W., Xu, J., Wang, H.: MADiff: Motion-aware mamba diffusion models for hand trajectory prediction on egocentric videos. arXiv preprint arXiv:2409.02638 (2024) 3

[80] Ma, J., Xu, J., Chen, X., Wang, H.: Diff-IP2D: Diffusion-based hand-object interaction prediction on egocentric videos. arXiv preprint arXiv:2405.04370 (2024) 3

[81] Ma, S., Cao, Q., Zhang, J., Tao, D.: Contact-aware human motion generation from textual descriptions. arXiv preprint arXiv:2403.15709 (2024) 3

[82] Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. Journal of machine learning research **9**(11) (2008) 8

[83] Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: AMASS: Archive of motion capture as surface shapes. In: ICCV (2019) 1, 3

[84] Mandery, C., Terlemez, O., Do, M., Vahrenkamp, N., Asfour, T.: The kit whole-body human motion database. In: ICAR (2015) 3

[85] 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) 3

[86] Mao, W., Liu, M., Salzmann, M.: Generating smooth pose sequences for diverse human motion prediction. In: CVPR (2021) 3

[87] 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) 3

[88] OpenAI: ChatGPT. https://chat.openai.com/ (2023) 2, 3, 4, 7, 8, 9, 10, 23, 26

[89] 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) 3

[90] Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning continuous signed distance functions for shape representation. In: CVPR (2019) 8

[91] 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) 1, 2, 3, 9, 21

[92] 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) 3

[93] Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3d human motion synthesis with transformer vae. In: ICCV (2021) 3

[94] Petrovich, M., Black, M.J., Varol, G.: TEMOS: Generating diverse human motions from textual descriptions. In: ECCV (2022) 3

[95] Petrovich, M., Black, M.J., Varol, G.: TMR: Text-to-motion retrieval using contrastive 3d human motion synthesis. In: ICCV (2023) 3

[96] Raab, S., Leibovitch, I., Li, P., Aberman, K., Sorkine-Hornung, O., Cohen-Or, D.: MoDi: Unconditional motion synthesis from diverse data. In: CVPR (2023) 3

[97] Raab, S., Leibovitch, I., Tevet, G., Arar, M., Bermano, A.H., Cohen-Or, D.: Single motion diffusion. arXiv preprint arXiv:2302.05905 (2023) 3

[98] 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) 8, 10

[99] Razali, H., Demiris, Y.: Action-conditioned generation of bimanual object manipulation sequences. In: AAAI (2023) 3

[100] 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) 3

[101] Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022) 22

[102] Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. ACM Transactions on Graphics **36**(6) (2017) 7

[103] Seo, Y., Hafner, D., Liu, H., Liu, F., James, S., Lee, K., Abbeel, P.: Masked world models for visual control. In: CoRL (2023) 5

[104] Shafir, Y., Tevet, G., Kapon, R., Bermano, A.H.: Human motion diffusion as a generative prior. arXiv preprint arXiv:2303.01418 (2023) 3

[105] Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: ICML (2015) 1

[106] Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020) 1

[107] Song, W., Zhang, X., Li, S., Gao, Y., Hao, A., Hou, X., Chen, C., Li, N., Qin, H.: HOIAnimator: Generating text-prompt human-object animations using novel perceptive diffusion models. In: CVPR (2024) 1, 3

[108] Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. ACM Trans. Graph. **38**(6), 209–1 (2019)

[109] 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) 3

[110] Sun, J., Chowdhary, G.: Towards globally consistent stochastic human motion prediction via motion diffusion. arXiv preprint arXiv:2305.12554 (2023) 3

[111] Taheri, O., Choutas, V., Black, M.J., Tzionas, D.: GOAL: Generating 4d whole-body motion for hand-object grasping. In: CVPR (2022) 3

[112] Taheri, O., Ghorbani, N., Black, M.J., Tzionas, D.: GRAB: A dataset of whole-body human grasping of objects. In: ECCV (2020) 3

[113] Tang, J., Wang, J., Ji, K., Xu, L., Yu, J., Shi, Y.: A unified diffusion framework for scene-aware human motion estimation from sparse signals. In: CVPR (2024) 3

[114] Tendulkar, P., Surís, D., Vondrick, C.: FLEX: Full-body grasping without full-body grasps. In: CVPR (2023) 3

[115] Tessler, C., Guo, Y., Nabati, O., Chechik, G., Peng, X.B.: MaskedMimic: Unified physics-based character control through masked motion inpainting. arXiv preprint arXiv:2409.14393 (2024) 3

[116] Tevet, G., Gordon, B., Hertz, A., Bermano, A.H., Cohen-Or, D.: Motionclip: Exposing human motion generation to clip space. In: ECCV (2022) 3

[117] Tevet, G., Raab, S., Cohan, S., Reda, D., Luo, Z., Peng, X.B., Bermano, A.H., van de Panne, M.: CLoSD: Closing the loop between simulation and diffusion for multi-task character control. arXiv preprint arXiv:2410.03441 (2024) 3

[118] Tevet, G., Raab, S., Gordon, B., Shafir, Y., Cohen-Or, D., Bermano, A.H.: Human motion diffusion model. arXiv preprint arXiv:2209.14916 (2022) 5, 6, 7, 9

[119] Tian, J., Yang, L., Ji, R., Ma, Y., Xu, L., Yu, J., Shi, Y., Wang, J.: Gaze-guided hand-object interaction synthesis: Benchmark and method. arXiv preprint arXiv:2403.16169 (2024) 3

[120] 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) 2, 3, 4, 9, 10

[121] Turk, A.M.: Amazon mechanical turk. Retrieved August **17**, 2012 (2012) 23

[122] 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) 3

[123] 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) 3

[124] Wang, J., Hodgins, J., Won, J.: Strategy and skill learning for physics-based table tennis animation. In: SIGGRAPH (2024) 3

[125] Wang, J., Xu, H., Xu, J., Liu, S., Wang, X.: Synthesizing long-term 3d human motion and interaction in 3d scenes. In: CVPR (2021) 3

[126] 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)

[127] Wang, J., Yan, S., Dai, B., Lin, D.: Scene-aware generative network for human motion synthesis. In: CVPR (2021) 3

[128] 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) 3

[129] 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) 3

[130] Wang, Z., Chen, Y., Liu, T., Zhu, Y., Liang, W., Huang, S.: HUMANISE: Language-conditioned human motion generation in 3d scenes. In: NeurIPS (2022) 1, 3

[131] Wang, Z., Li, D., Jiang, R.: Diffusion models in 3d vision: A survey. arXiv preprint arXiv:2410.04738 (2024) 1

[132] Wang, Z., Wang, J., Lin, D., Dai, B.: InterControl: Generate human motion interactions by controlling every joint. arXiv preprint arXiv:2311.15864 (2023) 3

[133] 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) 3

[134] 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) 2

[135] Wu, P., Escontrela, A., Hafner, D., Abbeel, P., Goldberg, K.: Daydreamer: World models for physical robot learning. In: CoRL (2023) 5

[136] 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) 1, 3

[137] Wu, Q., Dou, Z., Xu, S., Shimada, S., Wang, C., Yu, Z., Liu, Y., Lin, C., Cao, Z., Komura, T., et al.: DICE: End-to-end deformation capture of hand-face interactions from a single image. arXiv preprint arXiv:2406.17988 (2024) 3

[138] Wu, S., Liu, Y., Li, L., Bi, M., Zeng, W., Yang, X.: HIMO: A new benchmark for full-body human interacting with multiple objects. In: ECCV (2024) 3

[139] Wu, Y., Wang, J., Zhang, Y., Zhang, S., Hilliges, O., Yu, F., Tang, S.: SAGA: Stochastic whole-body grasping with contact. In: ECCV (2022) 3

[140] Wu, Z., Li, J., Liu, C.K.: Human-object interaction from human-level instructions. arXiv preprint arXiv:2406.17840 (2024) 1, 3

[141] 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) 3

[142] 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. In: ICLR (2024) 3

[143] Xie, X., Bhatnagar, B.L., Pons-Moll, G.: Chore: Contact, human and object reconstruction from a single rgb image. In: ECCV (2022) 3

[144] Xie, X., Lenssen, J.E., Pons-Moll, G.: InterTrack: Tracking human object interaction without object templates. arXiv preprint arXiv:2408.13953 (2024) 3

[145] 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) 3

[146] Xie, Z., Starke, S., Ling, H.Y., van de Panne, M.: Learning soccer juggling skills with layer-wise mixture-of-experts. In: SIGGRAPH (2022) 3

[147] 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) 3

[148] Xu, S., Li, Z., Wang, Y.X., Gui, L.Y.: InterDiff: Generating 3d human-object interactions with physics-informed diffusion. In: ICCV (2023) 3, 6, 7, 8, 9, 10, 25

[149] Xu, S., Wang, Y.X., Gui, L.Y.: Diverse human motion prediction guided by multi-level spatial-temporal anchors. In: ECCV (2022) 3

[150] Xu, S., Wang, Y.X., Gui, L.: Stochastic multi-person 3d motion forecasting. In: ICLR (2023) 3

[151] Xu, X., Joo, H., Mori, G., Savva, M.: D3D-HOI: Dynamic 3d human-object interactions from videos. arXiv preprint arXiv:2108.08420 (2021) 3

[152] Xu, Z., Chen, Q., Peng, Y., Liu, Y.: Semantic-aware human object interaction image generation. In: ICML (2024) 3

[153] Xue, K., Seo, H.: Shape conditioned human motion generation with diffusion model. arXiv preprint arXiv:2405.06778 (2024) 3

[154] Yang, C., Kang, C., Kong, K., Oh, H., Kang, S.J.: Person in Place: Generating associative skeleton-guidance maps for human-object interaction image editing. In: CVPR (2024) 3

[155] Yang, J., Niu, X., Jiang, N., Zhang, R., Huang, S.: F-HOI: Toward fine-grained semantic-aligned 3d human-object interactions. arXiv preprint arXiv:2407.12435 (2024) 2, 3

[156] Yang, Y., Zhai, W., Luo, H., Cao, Y., Zha, Z.J.: LEMON: Learning 3d human-object interaction relation from 2d images. In: CVPR (2024) 3, 22

[157] Yang, Y., Zhai, W., Wang, C., Yu, C., Cao, Y., Zha, Z.J.: EgoChoir: Capturing 3d human-object interaction regions from egocentric views. arXiv preprint arXiv:2405.13659 (2024) 3

[158] Yang, Z., Yin, K., Liu, L.: Learning to use chopsticks in diverse gripping styles. ACM Transactions on Graphics (TOG) **41**(4), 1–17 (2022) 3

[159] 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) 3

[160] Yazdian, P.J., Liu, E., Cheng, L., Lim, A.: MotionScript: Natural language descriptions for expressive 3d human motions. arXiv preprint arXiv:2312.12634 (2023) 3

[161] 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) 3

[162] Yuan, W., Shen, W., He, Y., Dong, Y., Gu, X., Dong, Z., Bo, L., Huang, Q.: Mogents: Motion generation based on spatial-temporal joint modeling. In: ECCV (2024) 3

[163] Yuan, Y., Kitani, K.: DLow: Diversifying latent flows for diverse human motion prediction. In: ECCV (2020) 3

[164] Zhang, C., Liu, Y., Xing, R., Tang, B., Yi, L.: Core4d: A 4d human-object-human interaction dataset for collaborative object rearrangement. arXiv preprint arXiv:2406.19353 (2024) 3

[165] 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) 3

[166] 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) 3

[167] Zhang, J., Zhang, Y., An, L., Li, M., Zhang, H., Hu, Z., Liu, Y.: ManiDext: Hand-object manipulation synthesis via continuous correspondence embeddings and residual-guided diffusion. arXiv preprint arXiv:2409.09300 (2024) 3

[168] 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) 3

[169] 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) 1, 3

[170] Zhang, J., Zhang, J., Song, Z., Shi, Z., Zhao, C., Shi, Y., Yu, J., Xu, L., Wang, J.: Hoi-mˆ 3: Capture multiple humans and objects interaction within contextual environment. In: CVPR (2024) 1, 3

[171] Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: ICCV (2023) 22

[172] 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) 3, 5, 6, 7, 9

[173] 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) 5, 6, 7, 9

[174] Zhang, M., Li, H., Cai, Z., Ren, J., Yang, L., Liu, Z.: Finemogen: Fine-grained spatio-temporal motion generation and editing. In: NeurIPS (2024) 3

[175] 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) 3

[176] Zhang, X., Bhatnagar, B.L., Starke, S., Guzov, V., Pons-Moll, G.: COUCH: Towards controllable human-chair interactions. In: ECCV (2022) 3

[177] Zhang, X., Bhatnagar, B.L., Starke, S., Petrov, I., Guzov, V., Dhamo, H., Pérez-Pellitero, E., Pons-Moll, G.: FORCE: Dataset and method for intuitive physics guided human-object interaction. arXiv preprint arXiv:2403.11237 (2024) 1

[178] 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) 3

[179] Zhang, Z., Liu, R., Aberman, K., Hanocka, R.: TEDi: Temporally-entangled diffusion for long-term motion synthesis. arXiv preprint arXiv:2307.15042 (2023) 3

[180] 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) 1, 3

[181] Zhao, K., Li, G., Tang, S.: DART: A diffusion-based autoregressive motion model for real-time text-driven motion control. arXiv preprint arXiv:2410.05260 (2024) 3

[182] Zhao, K., Wang, S., Zhang, Y., Beeler, T., Tang, S.: Compositional human-scene interaction synthesis with semantic control. In: ECCV (2022) 3

[183] Zhao, K., Zhang, Y., Wang, S., Beeler, T., Tang, S.: Synthesizing diverse human motions in 3d indoor scenes. In: ICCV (2023) 3

[184] 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) 3

[185] Zhong, L., Xie, Y., Jampani, V., Sun, D., Jiang, H.: SMooDi: Stylized motion diffusion model. arXiv preprint arXiv:2407.12783 (2024) 3

[186] Zhou, K., Bhatnagar, B.L., Lenssen, J.E., Pons-Moll, G.: Toch: Spatio-temporal object-to-hand correspondence for motion refinement. In: ECCV (2022) 3

[187] 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) 3