


Text2HOI: Text-guided 3D Motion Generation for Hand-Object Interaction

Junuk Cha¹Jihyeon Kim^{1,2†}Jae Shin Yoon^{3*}Seungryul Baek^{1*}¹UNIST²KETI³Adobe Research

"Open a box with the right hand."



"Hand over an apple with both hands."




Figure 1. Given a text and a canonical object mesh as prompts, we generate 3D motion for hand-object interaction without requiring object trajectory and initial hand pose. We represent the right hand with a light skin color and the left hand with a dark skin color. The articulation of a box in the first row is controlled by estimating an angle for the pre-defined axis of the box.

Abstract

This paper introduces the first text-guided work for generating the sequence of hand-object interaction in 3D. The main challenge arises from the lack of labeled data where existing ground-truth datasets are nowhere near generalizable in interaction type and object category, which inhibits the modeling of diverse 3D hand-object interaction with the correct physical implication (e.g., contacts and semantics) from text prompts. To address this challenge, we propose to decompose the interaction generation task into two subtasks: hand-object contact generation; and hand-object motion generation. For contact generation, a VAE-based network takes as input a text and an object mesh, and generates the probability of contacts between the surfaces of hands and the object during the interaction. The network learns a variety of local geometry structure of diverse objects that is independent of the objects' category, and thus, it is applicable to general objects. For motion generation, a Transformer-based diffusion model utilizes this 3D contact map as a strong prior for generating physically

plausible hand-object motion as a function of text prompts by learning from the augmented labeled dataset; where we annotate text labels from many existing 3D hand and object motion data. Finally, we further introduce a hand refiner module that minimizes the distance between the object surface and hand joints to improve the temporal stability of the object-hand contacts and to suppress the penetration artifacts. In the experiments, we demonstrate that our method can generate more realistic and diverse interactions compared to other baseline methods. We also show that our method is applicable to unseen objects. We will release our model and newly labeled data as a strong foundation for future research. Codes and data are available in: <https://github.com/JunukCha/Text2HOI>.

1. Introduction

Imagine handing over an apple on a table to your friends: you might first grab it and convey this to them. During a social interaction, the hand pose and motion are often defined as a function of object's pose, shape, and category. While existing works [3, 8, 9, 15, 21, 27, 30, 31] have been successful in modeling diverse and realistic 3D human body motions from a text prompt (where there exists no text-guided hand motion genera-

This research was conducted when Jihyeon Kim was a graduate student (Master candidate) at UNIST†. Co-last authors*.

tion works), the context of object interaction has been often missing, which significantly limits the expressiveness in the semantics of the generated motion sequence. In this paper, we propose a first work that can generate realistic and assorted hand-object motions in 3D from a text prompt as shown in Fig. 1. Our work can be used for various applications such as generating surgical simulations, interactive control of a character for gaming, and future path planning between a robot hand and objects for robotics.

Learning to generate a sequence of 3D hand-object interaction from a text prompt is extremely challenging due to the scarcity of the dataset: the diversity of existing datasets for a sequence of 3D meshes and associated text labels is far behind the one of real-world distribution which is determined by a number of parameters such as hand type (*e.g.*, left or right), object’s category and structure, scale, contact regions, and so on. A generative model learned from such limited data will fail in the diverse modeling of physically and semantically plausible 3D hand-object interaction.

To overcome this challenge, we propose to decompose the interaction generation task into two subtasks, “object contact map generation” and “hand-object motion generation”, where the models dedicated to each task learn a general geometry representation from the augmented dataset, which leads to the significant improvement in the generalizability and physical plausibility of the combined pipeline.

For contact map generation, we newly develop a contact map prediction network that encodes a local geometry surface of a 3D object mesh along with a target motion text; and generates a 3D contact map—3D probability map at the object’s surface that describes the potential regions contacted by hand meshes during the interaction—along with the general geometric features. Since the local geometry representation is category-agnostic, the network is applicable to general objects. By adding condition of the scale information, our contact map generation module is, in nature, able to decode scale-variant probability, *e.g.*, if the object’s scale is smaller, the region of the predicted contact probability becomes wider, reflecting the natural tendency to grasp smaller objects over a wider area.

For motion generation, a Transformer-based diffusion model utilizes the contact map and geometric features as strong guidance to generate the sequence of 3D hand and object movements from a text prompt. Unlike a conventional diffusion process [11], the model is designed to directly estimate the final sample at each step, which allows us to apply explicit geometric loss (*e.g.*, relative distance or orientation) to improve the geometric correctness. Our diffusion model learns the augmented data where we perform extensive manual annotation of the text labels from external motion datasets [5, 16, 25].

Using these two modules, we introduce the first text-guided hand-object interaction generation framework that generates the 3D interaction in a compositional way. Given a text prompt, canonical 3D object mesh, and object’s scale, our VAE-based

contact predictor generates a 3D contact map, and geometry features. Our Transformer-based diffusion model encodes the contact, text, and geometry information with frame-wise and agent-wise (*i.e.*, object, and left and right hand) positional embedding to decode realistic 3D hand-object interaction. Finally, our new Transformer-based refiner module pushes the physical correctness of the 3D interaction in a single feed-forward manner by refining the contacts and suppressing the penetration artifacts.

In the experiments, we validate our model on three datasets (H2O [16], GRAB [25], and ARCTIC [5]) where our method outperforms other baseline methods in terms of accuracy, diversity, and physical realism by large margins. We also demonstrate that our compositional framework enables the application of our method to new objects that are not seen during training.

Our contributions can be summarized as follows:

- To the best of our knowledge, we propose the first approach that can generate a sequence of 3D hand-object interaction in various styles and lengths from a text prompt.
- We propose a novel compositional framework that enables the modeling of high-quality hand-object interaction from limited data.
- We introduce a new fast and efficient hand refinement module that improves physical realism (*e.g.*, penetration-free interaction) without any test-time optimization.
- We annotate text labels from existing hands and object motion datasets, which will be made public.

2. Related Work

Text to human motion generation. Thanks to the user-friendly nature of textual inputs, there has been substantial progress in the field of text-guided human motion generation [3, 8, 9, 12, 15, 17, 19, 21, 27, 30, 31, 33]. Guo *et al.* [8] proposed the *text2length* and *text2motion* modules to generate human motion in varying time length, while remaining realistic and faithful to the text. Tevet *et al.* [27] introduced a Motion Diffusion Model (MDM) for generating natural and expressive human motion, utilizing the geometric losses and Transformer-based approach that predicts the sample instead of noise in each diffusion step. Recently, Liang *et al.* [19] presented a method that can generate interactive motion between two people. But it cannot handle three or more multi-agents.

Hand and object motion generation. Existing approaches [1, 2, 4, 7, 10, 13, 14, 18, 20, 34] focus on grasping the stationary object. They are limited in their ability to manipulate the object and are therefore inadequate to generate a natural hand-object motion. Ghosh *et al.* [6] proposed a method for generating full-body motion in interaction with 3D objects, which is guided by action labels, while it requires an optimization stage for full performance. To generate hand and manipulated object motion, Zhang *et al.* [29] proposed a network that relies on the current hand pose, past and future trajectories of both hands and object,

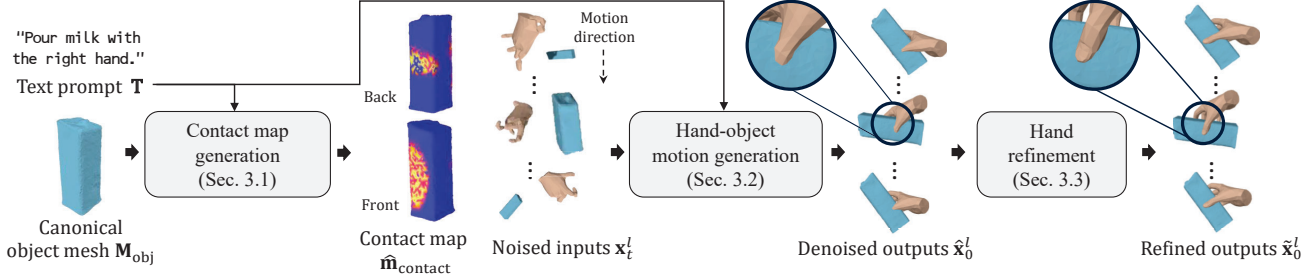


Figure 2. **Schematic diagram of the overall framework.** Given a text prompt and a canonical object mesh prompt, our aim is to generate the 3D motion for hand-object interaction. We first generate a contact map from the canonical object mesh conditioned by the text prompt and object’s scale. The hand-object motion generation module removes the noise from the inputs for the denoised outputs to align with the predicted contact map and the text prompt. The denoised outputs exhibit artifacts, including the penetration. To address these artifacts, the hand refinement module adjusts the generated (denoised) hand pose parameters to restrain the penetration and to improve contact interactions.

and diverse spatial representations. Zheng *et al.* [32] generate the hand-object motion covering both rigid and articulated objects, given an initial hand pose, object geometry, and sparse sequences of object poses. While plausible, these methods [29, 32] require the 3D object sequence as inputs, which is often not available from a user. In addition, they cannot utilize text modality.

3. Method

Our goal is to generate hand-object interacting motions given a text prompt \mathbf{T} and a canonical object mesh \mathbf{M}_{obj} . To address them, we design our framework with three stages, as shown in Fig. 2. First, we use the canonical object mesh \mathbf{M}_{obj} combined with the text feature $f^{\text{CLIP}}(\mathbf{T})$ via the CLIP text encoder f^{CLIP} [23] to estimate the contact map $\hat{\mathbf{m}}_{\text{contact}}$ that provides a strong prior for relative 3D locations of hands and an object. Then, we use the Transformer-based diffusion model to denoise the noised input data $\{\mathbf{x}_t^l\}_{l=1}^L$ at the t -th diffusion time-step, where L is the overall sequence length. By conditioning the text features $f^{\text{CLIP}}(\mathbf{T})$, contact map $\hat{\mathbf{m}}_{\text{contact}}$, object features \mathbf{F}_{obj} and scale s_{obj} on the diffusion model, we estimate the denoised sample $\hat{\mathbf{x}}_0$ from the noised one \mathbf{x}_t . Lastly, hand refiner improves the initial generated hand-object motions considering penetration and contact between hands and an object.

3.1. Contact map prediction

To generate natural motions for hand-object interaction, it is crucial to understand contact points between hands and an object. For this, we design the contact prediction network f^{contact} that encodes contact points on the surfaces of the object mesh \mathbf{M}_{obj} along with a text prompt \mathbf{T} and object’s scale s_{obj} .

We first compute s_{obj} which represents the maximum distance from center of object mesh to its vertices. We then sample N -point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$ from the vertices of canonical object mesh using the farthest point sampling (FPS) algorithm [22]. Subsequently, we normalize \mathbf{P} to \mathbf{P}_{norm} by dividing it with s_{obj} . The contact prediction network f^{contact} receives the normalized

point cloud \mathbf{P}_{norm} , text features $f^{\text{CLIP}}(\mathbf{T})$, object’s scale s_{obj} , and Gaussian random noise vector $\mathbf{z}_{\text{contact}} \in \mathbb{R}^{64}$, and produces the contact map $\hat{\mathbf{m}}_{\text{contact}} \in \mathbb{R}^{N \times 1}$. In the middle of f^{contact} , we obtain the object features $\mathbf{F}_{\text{obj}} \in \mathbb{R}^{1,024}$. To train f^{contact} , we use the combination of binary cross-entropy loss, dice loss and kullback-leibler (KL) divergence loss following [18].

3.2. Text-to-3D hand-object motion generation

Our text-to-3D hand-object interaction generator (Text2HOI) f^{THOI} , whose architecture is the Transformer encoder [28], is trained via the diffusion-based approach [11].

3.2.1 Preliminaries.

The 3D hand-object motion is represented as $\mathbf{x}_0 = \{\mathbf{x}_{0,\text{lhand}}^l, \mathbf{x}_{0,\text{rhand}}^l, \mathbf{x}_{0,\text{obj}}^l\}_{l=1}^{L_{\text{max}}}$, where l denotes the frame index. This motion comprises $3 \cdot L_{\text{max}}$ elements, which accounts for the maximum motion length L_{max} of three agents (*i.e.*, left and right hands and an object): For left and right hands, $\mathbf{x}_{0,\text{lhand}}^l \in \mathbb{R}^{99}$ and $\mathbf{x}_{0,\text{rhand}}^l \in \mathbb{R}^{99}$ are composed of 99-dimensional vectors by flattening and concatenating the 3D hand translation parameters $\mathbf{t}_h^l \in \mathbb{R}^3$ and MANO hand pose parameters $\theta^l \in \mathbb{R}^{16 \times 6}$ in 6D representation [35]. For an object, $\mathbf{x}_{0,\text{obj}}^l \in \mathbb{R}^{10}$ is 10-dimensional vector that concatenates the 3D object translation $\mathbf{t}_o^l \in \mathbb{R}^3$, object rotation $\mathbf{r}^l \in \mathbb{R}^6$ [35], and object articulation angle $\alpha^l \in \mathbb{R}^1$.

The 3D hand-object interaction \mathbf{x}_0 is used to generate the mesh of hands, and to deform the mesh of objects: The left and right hand meshes are generated from $\mathbf{x}_{0,\text{lhand}}$ and $\mathbf{x}_{0,\text{rhand}}$ by feeding them to the MANO layer [24] to output the hand vertices $\mathbf{V}_{\text{lhand}}, \mathbf{V}_{\text{rhand}} \in \mathbb{R}^{L \times V \times 3}$, and hand joints $\mathbf{J}_{\text{lhand}}, \mathbf{J}_{\text{rhand}} \in \mathbb{R}^{L \times J \times 3}$ in 3D global space, where $V = 778$ and $J = 21$. A deformed object’s point cloud $\mathbf{P}_{\text{def}} \in \mathbb{R}^{L \times N \times 3}$ is generated in 3D global space by transforming the object’s point clouds \mathbf{P} with the translation, rotation and articulation angles in $\mathbf{x}_{0,\text{obj}}$. The notation $\hat{\cdot}$ and $\tilde{\cdot}$ indicate that these values are derived from the estimated $\hat{\mathbf{x}}_0$ and refined $\tilde{\mathbf{x}}_0$, respectively.

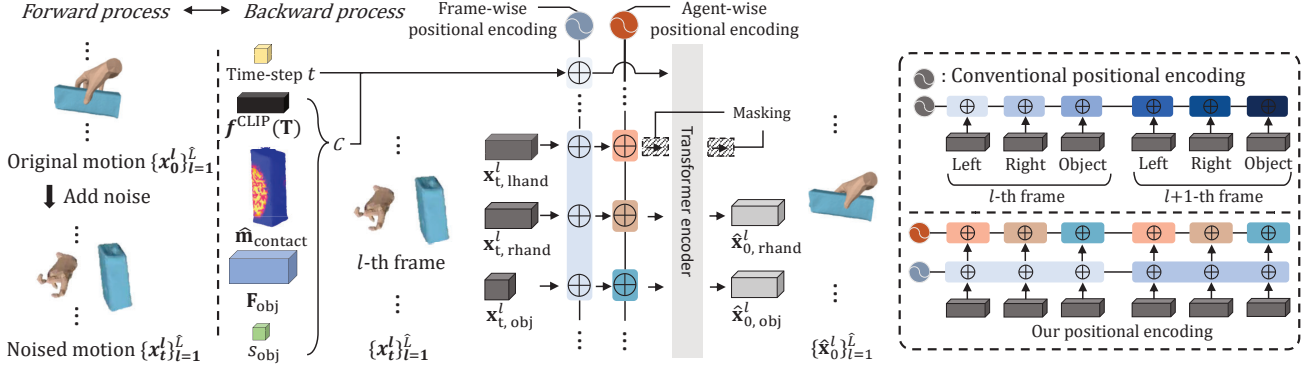


Figure 3. The details of the text-to-3D hand-object motion generation in our framework. In the forward process, we generate the noised motion $\{\mathbf{x}_t^l\}_{l=1}^{\hat{L}}$ by adding the noise to the original (ground-truth) motion $\{\mathbf{x}_0^l\}_{l=1}^{\hat{L}}$. In the backward process, the Transformer encoder denoises the noised motion $\{\mathbf{x}_t^l\}_{l=1}^{\hat{L}}$, using various conditions c including text features $f^{\text{CLIP}}(\mathbf{T})$, contact map $\hat{\mathbf{m}}_{\text{contact}}$, object features \mathbf{F}_{obj} , and object’s scale s_{obj} . The right panel illustrates a comparison between conventional positional encoding, which can only differentiate each patch, and our proposed encoding, which provides detailed differentiation of both frames and agents. A unique positional encoding value is assigned for each box, distinguished by different colors.

3.2.2 Forward process.

Our forward process is formulated as:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t \quad (1)$$

following [11], where t is the diffusion time-step, \mathbf{x}_0 is the original 3D hand-object motion, \mathbf{x}_t is the noised 3D hand-object motion at the t -th diffusion time-step, and $\bar{\alpha}_t \in (0,1)$ is a set of constant hyper-parameters. The noise ϵ_t is randomly sampled from the Gaussian distribution at each diffusion-time step t .

3.2.3 Backward process.

In the backward process, the text-to-3D hand-object interaction generator (Text2HOI) f^{THOI} denoises the noised motion \mathbf{x}_t to reconstruct the original (ground-truth) motion \mathbf{x}_0 : $\hat{\mathbf{x}}_0 = f^{\text{THOI}}(\mathbf{x}_t, t, c)$, where c denotes the conditions, as described in [27]. Since we exploit the Transformer encoder as the architecture, the noised signal \mathbf{x}_t needs to be first converted to the proper input embedding \mathbf{X}_t . Similarly, the output of Transformer architecture $\hat{\mathbf{X}}_t$ also needs to be converted to the denoised signal $\hat{\mathbf{x}}_0$. Furthermore, the text features $f^{\text{CLIP}}(\mathbf{T})$, object features \mathbf{F}_{obj} , estimated contact map $\hat{\mathbf{m}}_{\text{contact}}$ and object’s scale s_{obj} are merged together to constitute the conditional signals \mathbf{X}_{cond} , which will be detailed in the remainder of the section:

Transformer input generation. The forwarded signal $\mathbf{x}_t^l = \{\mathbf{x}_{t,\text{lhand}}^l, \mathbf{x}_{t,\text{rhand}}^l, \mathbf{x}_{t,\text{obj}}^l\}$ is passed through corresponding fully connected layers (*i.e.*, $f^{\text{in,lhand}}$, $f^{\text{in,rhand}}$, and $f^{\text{in,obj}}$), respectively to obtain the input to the Transformer encoder, $\mathbf{X}_t^l = \{\mathbf{X}_{t,\text{lhand}}^l \in \mathbb{R}^{512}, \mathbf{X}_{t,\text{rhand}}^l \in \mathbb{R}^{512}, \mathbf{X}_{t,\text{obj}}^l \in \mathbb{R}^{512}\}$, respectively. Then, we apply two types of positional encoding: frame-wise and agent-wise. Frame-wise positional encoding adds an sinusoidal value to \mathbf{X}_t^l which varies according to the motion length index l ; while irrespective to the type of agents.

Agent-wise positional encoding adds distinct encoding values for each agent (left hand, right hand, and object), which are consistent across different frames, to $\mathbf{X}_{t,\text{lhand}}$, $\mathbf{X}_{t,\text{rhand}}$, and $\mathbf{X}_{t,\text{obj}}$. These are designed to help the Transformer encoder to better understand the input data. The detail pipeline of these positional encodings is shown in the right bottom panel of Fig. 3.

The Transformer encoder has a maximum input capacity of 451. The first input is reserved for the conditioning, while the remaining inputs accommodate the maximum motion length L_{max} of 150 frames, involving three distinct agents: left hand, right hand and object. We mask out all inputs except for the first $1+3\hat{L}$ inputs where \hat{L} is the estimated length of sequence and subsequently, we mask inputs which are not belonging to the estimated hand type \mathbf{H}^* (see Sec. 4.1 for details about how \hat{L} and \mathbf{H}^* are estimated).

Conditional input generation. To generate denoised hand-object motions conditioned on the text prompt \mathbf{T} and canonical object mesh \mathbf{M}_{obj} , we need to generate the conditional input for the t -th diffusion time-step. Conditional input $\mathbf{X}_{t,\text{cond}}$ is generated by:

$$\mathbf{X}_{t,\text{cond}} = \mathbf{X}_{\text{cond}} + t_{\text{emb}} \quad (2)$$

where the diffusion time-step embedding $t_{\text{emb}} = f^{\text{ts}}(t)$ is obtained by applying the diffusion time-step t to the time-step embedding fully-connected layer f^{ts} and the condition embedding \mathbf{X}_{cond} is generated as follows:

$$\mathbf{X}_{\text{cond}} = \mathbf{X}_{\text{text}}^{\text{cond}} + \mathbf{X}_{\text{obj}}^{\text{cond}} \quad (3)$$

where the text condition $\mathbf{X}_{\text{text}}^{\text{cond}} = f^{\text{text}}(f^{\text{CLIP}}(\mathbf{T}))$ is generated by applying the text feature $f^{\text{CLIP}}(\mathbf{T})$ to the fc layer f^{text} . The object condition $\mathbf{X}_{\text{obj}}^{\text{cond}} = f^{\text{obj}}(\{\mathbf{F}_{\text{obj}}, \hat{\mathbf{m}}_{\text{contact}}, s_{\text{obj}}\})$ is obtained by concatenating object feature \mathbf{F}_{obj} , contact map $\hat{\mathbf{m}}_{\text{contact}}$ and object’s scale s_{obj} , and feeding them to the fc layer f^{obj} .

Transformer output conversion. Masked inputs $\mathbf{X}_t = \{\mathbf{X}_{t,\text{cond}}, \mathbf{X}_t^1, \mathbf{X}_t^2, \dots, \mathbf{X}_t^{\hat{L}}\}$ are fed to the Transformer encoder to estimate the outputs $\hat{\mathbf{X}}_0 = \{\hat{\mathbf{X}}_0^l\}_{l=1}^{\hat{L}}$, where $\hat{\mathbf{X}}_0^l = \{\hat{\mathbf{x}}_{0,\text{lhand}}^l, \hat{\mathbf{x}}_{0,\text{rhand}}^l, \hat{\mathbf{x}}_{0,\text{obj}}^l\}$. Each outputs— $\hat{\mathbf{X}}_{0,\text{lhand}}^l$, $\hat{\mathbf{X}}_{0,\text{rhand}}^l$, and $\hat{\mathbf{X}}_{0,\text{obj}}^l$ —are passed through its own dedicated fully connected layer, denoted as $f^{\text{out,lhand}}$, $f^{\text{out,rhand}}$, and $f^{\text{out,obj}}$, respectively, to obtain the denoised signal $\hat{\mathbf{x}}_0 = \{\hat{\mathbf{x}}_{0,\text{lhand}}^l, \hat{\mathbf{x}}_{0,\text{rhand}}^l, \hat{\mathbf{x}}_{0,\text{obj}}^l\}$.

Training. Note that the losses related to left and right hands are activated by indicator functions $\mathbb{1}_{\text{left}}$ and $\mathbb{1}_{\text{right}}$, respectively, which are derived from the hand type \mathbf{H}^* . The f^{THOI} is trained with loss functions as follows:

$$L_{\text{THOI}}(f^{\text{THOI}}) = L_{\text{diff}}(f^{\text{THOI}}) + L_{\text{dm}}(f^{\text{THOI}}) + L_{\text{ro}}(f^{\text{THOI}}) \quad (4)$$

where

$$L_{\text{diff}}(f^{\text{THOI}}) = E_{\mathbf{x}_t \sim q(\mathbf{x}_0|c), t \sim [1, T]} \|\mathbf{x}_0 - f^{\text{THOI}}(\mathbf{x}_t, t, c)\|_2^2 \quad (5)$$

is the loss which is used to reconstruct \mathbf{x}_0 from \mathbf{x}_t similar to [27]. We have two more losses (*i.e.*, L_{dm} , L_{ro}) to make the f^{THOI} to generate more accurate hand-object motions. The distance map loss L_{dm} , proposed in [19], is employed in our hand-object motion generation to align the estimated distance map with ground-truth distance map as follows:

$$L_{\text{dm}}(f^{\text{THOI}}) = \sum_{i=1}^{\hat{L} \times J \times N} \left\{ \mathbb{1}_{\text{left}} \cdot \left((\hat{\mathbf{d}}_{\text{left}}^i - \mathbf{d}_{\text{left}}^i) \cdot I(\mathbf{d}_{\text{left}}^i < \tau) \right)^2 + \mathbb{1}_{\text{right}} \cdot \left((\hat{\mathbf{d}}_{\text{right}}^i - \mathbf{d}_{\text{right}}^i) \cdot I(\mathbf{d}_{\text{right}}^i < \tau) \right)^2 \right\} \quad (6)$$

where $\hat{\mathbf{d}}_{\text{left}}^i$ and $\hat{\mathbf{d}}_{\text{right}}^i$ denote the i -th element of $\hat{\mathbf{d}}_{\text{left}}$ and $\hat{\mathbf{d}}_{\text{right}} \in \mathbb{R}^{\hat{L} \times J \times N}$, respectively. These represent the estimated distance maps between the J hand joints (left $\hat{\mathbf{J}}_{\text{lhand}}$ and right $\hat{\mathbf{J}}_{\text{rhand}}$) and the N object points $\hat{\mathbf{P}}_{\text{def}}$ across a sequence of \hat{L} frames, derived from their 3D global positions. $\mathbf{d}_{\text{left}}^i$ and $\mathbf{d}_{\text{right}}^i$ denote the i -th element of \mathbf{d}_{left} and $\mathbf{d}_{\text{right}} \in \mathbb{R}^{\hat{L} \times J \times N}$ which are the ground-truth distance maps obtained for left and right hands, respectively. The indicator function $I(\cdot)$ outputs 1 when the statement is true and 0, otherwise. It activates the loss only when the hand-object distance is below the distance threshold τ , where it is empirically set as $2cm$.

In the relative orientation loss L_{ro} , we consider the 3D relative rotation as follows, as hands and objects exhibit severe rotation changes:

$$L_{\text{ro}}(f^{\text{THOI}}) = \mathbb{1}_{\text{left}} \cdot \|R(\hat{\mathbf{x}}_{0,\text{lhand}}, \hat{\mathbf{x}}_{0,\text{obj}}) - R(\mathbf{x}_{0,\text{lhand}}, \mathbf{x}_{0,\text{obj}})\|_2^2 + \mathbb{1}_{\text{right}} \cdot \|R(\hat{\mathbf{x}}_{0,\text{rhand}}, \hat{\mathbf{x}}_{0,\text{obj}}) - R(\mathbf{x}_{0,\text{rhand}}, \mathbf{x}_{0,\text{obj}})\|_2^2, \quad (7)$$

where $R(\cdot, \cdot)$ indicates the 3D relative orientation between hand and object.

Sampling. At each time-step t , the model f^{THOI} predicts a clean motion, denoted as $\hat{\mathbf{x}}_0 = f^{\text{THOI}}(\mathbf{x}_t, t, c)$, and then re-noise $\hat{\mathbf{x}}_0$ to \mathbf{x}_{t-1} [27]. This procedure is conducted repeatedly, starting from $t=T$ to $t=1$.

3.3. Hand refinement network

We propose a hand refinement network f^{ref} that considers the contact and penetration between hands and an object generated from Text2HOI f^{THOI} in Sec. 3.2. The architecture of f^{ref} is similar to that of f^{THOI} : 1) it employs a Transformer encoder architecture, and 2) it utilizes frame-wise and agent-wise position encoding. The main differences between f^{THOI} and f^{ref} are that f^{ref} does not involve the diffusion mechanism; it does not receive any conditions as input; and it refines only hand motions.

Inputs and outputs. The hand refinement network receives several inputs: Text2HOI’s hand output $\hat{\mathbf{x}}_{0,\text{hand}}$, hand joints $\hat{\mathbf{J}}_{\text{hand}}$, predicted contact map $\hat{\mathbf{m}}_{\text{contact}}$, deformed object’s point cloud $\hat{\mathbf{P}}_{\text{def}}$, and distance-based attention map \mathbf{m}_{att} . The attention map $\mathbf{m}_{\text{att}} = \exp(-50 \times \mathbf{D})$ is defined as [26], where $\mathbf{D} \in \mathbb{R}^{J \times 3}$ represents the 3D displacement between J hand joints $\hat{\mathbf{J}}_{\text{hand}}$ and the nearest object points in $\hat{\mathbf{P}}_{\text{def}}$. These components, denoted as $\hat{\mathbf{x}}_{0,\text{hand}}$, $\hat{\mathbf{J}}_{\text{hand}}$, $\hat{\mathbf{m}}_{\text{contact}}$, $\hat{\mathbf{P}}_{\text{def}}$, and \mathbf{m}_{att} , are flattened and concatenated to form the hand refiner input. As indicated in Sec. 3.2, these inputs are masked using \mathbf{H}^* . Then, f^{ref} outputs the refined hand motions $\tilde{\mathbf{x}}_{\text{hand}}$. They are masked using \mathbf{H}^* for loss calculation and result visualization.

Training. The hand refinement network is trained using the loss function L_{refine} as follows:

$$L_{\text{refine}}(f^{\text{ref}}) = L_{\text{simple}}(f^{\text{ref}}) + L_{\text{penet}}(f^{\text{ref}}) + \lambda_1 L_{\text{contact}}(f^{\text{ref}}), \quad (8)$$

where λ_1 is set as 5. The simple L2 loss is expressed as follows:

$$L_{\text{simple}}(f^{\text{ref}}) = \|\tilde{\mathbf{x}}_{\text{hand}} - \mathbf{x}_{\text{hand}}\|_2^2, \quad (9)$$

where \mathbf{x}_{hand} denotes the ground-truth hand motions. The penetration loss L_{penet} [13] is applied only on hand vertices that penetrate the object surfaces as follows:

$$L_{\text{penet}}(f^{\text{ref}}) = \mathbb{1}_{\text{left}} \cdot \|d(\tilde{v}_{\text{lhand}}, \hat{p}_{\text{obj}}^{\text{left}})\|^2 + \mathbb{1}_{\text{right}} \cdot \|d(\tilde{v}_{\text{rhand}}, \hat{p}_{\text{obj}}^{\text{right}})\|^2, \quad (10)$$

where $d(\cdot, \cdot)$ denotes the Euclidean distance between two points, $\tilde{v}_{\text{lhand}} \in \tilde{\mathbf{V}}_{\text{lhand}}$ and $\tilde{v}_{\text{rhand}} \in \tilde{\mathbf{V}}_{\text{rhand}}$ are hand vertices that penetrate the object surface, and $\hat{p}_{\text{obj}}^{\text{left}} \in \hat{P}_{\text{def}}$ and $\hat{p}_{\text{obj}}^{\text{right}} \in \hat{P}_{\text{def}}$ denote the object points closest to \tilde{v}_{lhand} and \tilde{v}_{rhand} , respectively. The contact loss L_{contact} [13] is applied to the joints that are sufficiently close to the object surface, as follows:

$$L_{\text{contact}}(f^{\text{ref}}) = \mathbb{1}_{\text{left}} \cdot \|d(\tilde{j}_{\text{lhand}}, \hat{c}_{\text{obj}}^{\text{left}})\|^2 + \mathbb{1}_{\text{right}} \cdot \|d(\tilde{j}_{\text{rhand}}, \hat{c}_{\text{obj}}^{\text{right}})\|^2, \quad (11)$$

where $\tilde{j}_{\text{lhand}} \in \tilde{\mathbf{J}}_{\text{lhand}}$ and $\tilde{j}_{\text{rhand}} \in \tilde{\mathbf{J}}_{\text{rhand}}$ represent the hand joints that are within a distance threshold τ from the object surface, respectively. $\hat{c}_{\text{obj}}^{\text{left}} \in \hat{P}_{\text{def}}$ and $\hat{c}_{\text{obj}}^{\text{right}} \in \hat{P}_{\text{def}}$ represent object points closest to \tilde{j}_{lhand} and \tilde{j}_{rhand} , respectively.

Table 1. Comparison on H2O, GRAB, and ARCTIC datasets. † denotes our produced results. → denotes that the higher value of the metric, the closer to the GT distribution. Best results are emphasized in bold.

Method	Accuracy (top-3) ↑	FID ↓	H2O		
			Diversity →	Multimodality ↑	Physical realism ↑
GT	0.9920 ± 0.0003	-	0.6057 ± 0.0050	0.2067 ± 0.0024	0.4790 ± 0.0002
T2M† [8]	0.6463 ± 0.0014	0.3439 ± 0.0006	0.3475 ± 0.0040	0.0634 ± 0.0022	0.3890 ± 0.016
MDM† [27]	0.5832 ± 0.0011	0.3015 ± 0.0011	0.5127 ± 0.0054	0.1738 ± 0.0049	0.5572 ± 0.0013
IMOS† [6]	0.5518 ± 0.0026	0.2945 ± 0.0011	0.4076 ± 0.0056	0.1798 ± 0.0115	0.3532 ± 0.0026
Ours	0.8295 ± 0.0015	0.1744 ± 0.0013	0.5365 ± 0.0073	0.2469 ± 0.0081	0.7574 ± 0.0022
GRAB					
GT	0.9994 ± 0.0001	-	0.8557 ± 0.0054	0.4390 ± 0.0045	0.8084 ± 0.0002
T2M† [8]	0.1897 ± 0.0007	0.7886 ± 0.0005	0.5712 ± 0.0078	0.0964 ± 0.0027	0.5844 ± 0.0002
MDM† [27]	0.5127 ± 0.0009	0.6023 ± 0.0011	0.8012 ± 0.0054	0.5194 ± 0.0145	0.7382 ± 0.0004
IMOS† [6]	0.4097 ± 0.0005	0.6147 ± 0.0003	0.6861 ± 0.0060	0.2845 ± 0.0036	0.6418 ± 0.0014
Ours	0.9218 ± 0.0010	0.3017 ± 0.0004	0.8351 ± 0.0061	0.5216 ± 0.0131	0.8839 ± 0.0005
ARCTIC					
GT	0.9997 ± 0.0001	-	0.5916 ± 0.0037	0.3279 ± 0.0038	0.9573 ± 0.0000
T2M† [8]	0.5234 ± 0.0015	0.3599 ± 0.0005	0.3301 ± 0.0023	0.0849 ± 0.0017	0.0143 ± 0.0001
MDM† [27]	0.5572 ± 0.0012	0.3025 ± 0.0006	0.4984 ± 0.0039	0.2632 ± 0.0065	0.7043 ± 0.0009
IMOS† [6]	0.8190 ± 0.0039	0.1826 ± 0.0005	0.5702 ± 0.0039	0.2741 ± 0.0049	0.7569 ± 0.0023
Ours	0.9205 ± 0.0012	0.1329 ± 0.0006	0.5758 ± 0.0042	0.3170 ± 0.0068	0.8760 ± 0.0009

4. Experiments

4.1. Implementation details

We use $T = 1,000$ noising steps and a cosine noise schedule. We use sinusoidal positional encoding for frame-wise and agent-wise positional encodings. We set the maximum length of motion sequences, denoted as L_{\max} , to 150 frames. Further details about network architecture can be found in the supplemental material.

Hand-type selection. We use the CLIP text encoder [23] f^{CLIP} to calculate cosine similarity between the input text prompt \mathbf{T} and predefined prompt templates $\Gamma(\mathbf{H}) = \text{“A photo of } \mathbf{H}\text{”}$, where $\mathbf{H} \in \{\text{left hand, right hand, both hands}\}$. The hand type \mathbf{H}^* with the highest cosine similarity to \mathbf{T} is selected for masking in the Transformer’s inputs, outputs, and losses (see Secs. 3.2, 3.3 and supplemental material.).

Motion length prediction. To obtain proper motion length $\hat{L} \leq L_{\max}$, we design a motion-length prediction network f^{Length} . It receives the text feature vector $f^{\text{CLIP}}(\mathbf{T})$ and Gaussian random noise $\mathbf{n} \in \mathbb{R}^{64}$, to predict the appropriate motion length \hat{L} for the text prompt \mathbf{T} . To train f^{Length} , we use the loss function $L_{\text{length}} = \|\hat{L} - L\|^2$, where L is ground truth.

4.2. Dataset

We use H2O [16], GRAB [25], and ARCTIC [5] in our experiment, which collects hand-object mesh sequences. We automatically generate text prompts by exploiting action labels for H2O and GRAB datasets; while we manually label text

prompts for ARCTIC. The characteristics of three datasets and details of our annotation process are further illustrated in the supplemental material.

4.3. Evaluation metrics and baselines

Evaluation metrics. We use the metrics of accuracy, frechet inception distance (FID), diversity, and multi-modality, as used in IMOS [6]. The accuracy serves as an indicator of how well the model generates motions and is evaluated by the pre-trained action classifier. We train a standard RNN-based action classifier to extract motion features and classify the action from the motions, as in IMOS [6]. The FID quantifies feature-space distances between real and generated motions, capturing the dissimilarity. The diversity reflects the range of distinct motions, and multi-modality measures the average variance of motions for an individual text prompt. To assess the physical realism of generated hand-object motions, we employ a physical model following the approach in ManipNet [29], assigning a realism score of 0 (unreal) or 1 (real) for measuring the realism of each frame. Experiments are conducted 20 times to establish the robustness, and we reported results within a 95% confidence interval.

Baselines. We compare our approach with three existing text-to-human motion generation methods: T2M [8], MDM [27], and IMOS [6]. T2M [8] employs a temporal VAE-based architecture and MDM [27] utilizes a diffusion model. IMOS [6] is designed to first generate human body and arm motions conditioned on both action labels and past body motions. It then optimizes object rotation and translation based on their history to generate body and arm motion. Since they were

Table 2. Ablation study on the positional encoding, losses, and conditions for ‘Ours w/o f^{ref} ’ and ablation study on losses for ‘Ours’.

Method	f^{ref}	GRAB				
		Accuracy (top-3) \uparrow	FID \downarrow	Diversity \rightarrow	Multimodality \uparrow	Physical realism \uparrow
GT	-	0.9994 \pm 0.0001	-	0.8557 \pm 0.0054	0.4390 \pm 0.0045	0.8084 \pm 0.0002
w/o frame-wise & agent-wise PE	\times	0.8294 \pm 0.0016	0.3461 \pm 0.0018	0.7814 \pm 0.063	0.4776 \pm 0.0194	0.8024 \pm 0.0007
w/o agent-wise PE	\times	0.8314 \pm 0.0012	0.3412 \pm 0.0006	0.8011 \pm 0.067	0.4755 \pm 0.0122	0.8221 \pm 0.0009
w/o L_{dm} & L_{ro}	\times	0.8289 \pm 0.0038	0.3416 \pm 0.0020	0.7887 \pm 0.0640	0.4654 \pm 0.0150	0.7490 \pm 0.0006
w/o L_{ro}	\times	0.8272 \pm 0.0020	0.3407 \pm 0.0015	0.7997 \pm 0.0079	0.4627 \pm 0.0104	0.8247 \pm 0.0011
w/o L_{dm}	\times	0.8202 \pm 0.0017	0.3444 \pm 0.0007	0.8156 \pm 0.0070	0.4819 \pm 0.0125	0.6410 \pm 0.0010
w/o $\hat{\mathbf{m}}_{\text{contact}}$ & s_{obj}	\times	0.8197 \pm 0.0009	0.3428 \pm 0.0012	0.7994 \pm 0.0055	0.4305 \pm 0.0121	0.7815 \pm 0.0006
w/o s_{obj}	\times	0.8274 \pm 0.0013	0.3413 \pm 0.0010	0.7963 \pm 0.0054	0.4405 \pm 0.0139	0.8018 \pm 0.0005
w/o $\hat{\mathbf{m}}_{\text{contact}}$	\times	0.8277 \pm 0.0027	0.3411 \pm 0.0006	0.8012 \pm 0.0067	0.4455 \pm 0.0115	0.7892 \pm 0.0009
Ours w/o f^{ref}	\times	0.8411 \pm 0.0009	0.3321 \pm 0.0006	0.8143 \pm 0.0050	0.4989 \pm 0.0154	0.8312 \pm 0.0005
w/o L_{penet} & L_{contact}	\checkmark	0.8838 \pm 0.0014	0.3234 \pm 0.0007	0.8277 \pm 0.0068	0.5111 \pm 0.014	0.6249 \pm 0.0008
w/o L_{contact}	\checkmark	0.8827 \pm 0.0008	0.3114 \pm 0.0013	0.8301 \pm 0.0048	0.4808 \pm 0.0151	0.1467 \pm 0.0005
w/o L_{penet}	\checkmark	0.8941 \pm 0.0009	0.3024 \pm 0.0005	0.8267 \pm 0.0061	0.5182 \pm 0.0099	0.8782 \pm 0.0006
Ours	\checkmark	0.9218 \pm 0.0010	0.3017 \pm 0.0004	0.8351 \pm 0.0061	0.5216 \pm 0.0131	0.8839 \pm 0.0005

“Use a hammer with the right hand.”

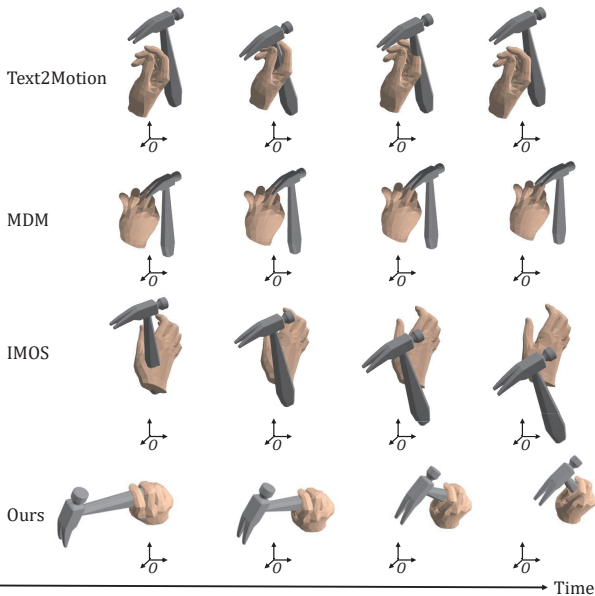


Figure 4. We compare our generated hand-object motions with other baselines’ results. Each row show the results of Text2Motion [8], MDM [27], IMOS [6], and ours.

originally designed for generating individual human motions from text prompts, to ensure a fair comparison, we re-train the methods using hand-object motion data, allowing them to generate hand-object motions from text prompts.

4.4. Experimental results

Comparison to other methods. We compare our method with other state-of-the-art methods (*i.e.*, T2M [8], MDM [27], and IMOS [6]), as shown in Tab. 1. For all datasets, our method outperforms other baselines in multiple measures. Particularly, our method demonstrates exceptional performance in generating physically realistic hand-object motions, as evidenced by achieving the highest score in Physical realism compared to other

approaches. The similarity of our distribution to the ground truth (GT) distribution in terms of Diversity, along with our highest scores in Multimodality and Accuracy, demonstrates our model’s capability to generate motions that are both diverse and accurate, and are well-aligned with text prompts. We compare our qualitative results with other baselines in Fig. 4. It shows that our method, compared to others, outperforms in generating motions where hand and object interact realistically, and these motions align closely with text prompts “Use a hammer with the right hand.”. The right hand well grabs the hammer and mimics the motion of driving something into a wall.

Qualitative results are shown in Fig. 5. Our method generates realistic hand-object motions that are closely aligned with the input text prompts, effectively handling even unseen objects. Please refer to the supplemental material for more visualizations including video results, and text-guided and scale-variant contact maps.

4.5. Ablation study

We conduct several ablation studies on GRAB dataset, to validate the effectiveness of our modules. The results are demonstrated in Tab. 2.

Position encoding. We introduce two types of positional encodings: frame-wise and agent-wise, which assists the Transformer to interpret inputs in a more distinct way. Seeing the results ‘w/o frame-wise & agent-wise PE’ and ‘w/o agent-wise PE’, we can conclude that by leveraging the specialized positional encoding, f^{THOI} is capable of generating more realistic hand-object motions.

Losses. We remove the distance map loss L_{dm} and relative orientation loss L_{ro} in our approach, and see how the performance changes: Seeing ‘w/o L_{dm} & L_{ro} ’, ‘w/o L_{dm} ’ and ‘w/o L_{ro} ’, we can conclude that these losses induce better results by facilitating model’s understanding of 3D relationship between hands and an object.

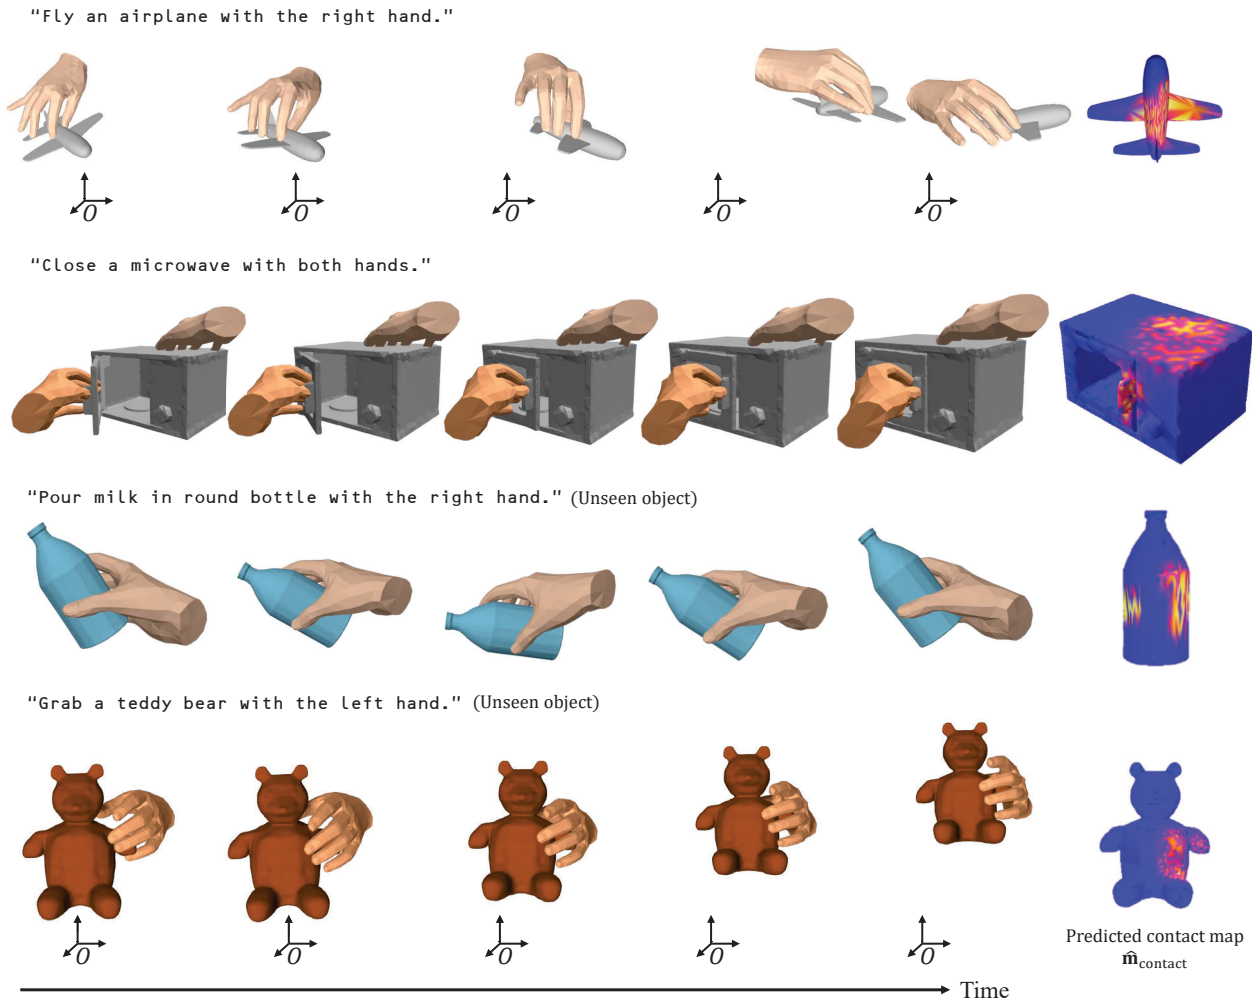


Figure 5. We demonstrate the generated hand-object motions and the predicted contact map results. The first and second rows show the results with objects seen during training. The third and fourth rows show the results with objects unseen during training.

Condition inputs. We remove the contact map \hat{m}_{contact} and scale of the object s_{obj} conditions from the original pipeline, and see how the performance changes. Seeing ‘w/o $\hat{m}_{\text{contact}} \& s_{\text{obj}}$ ’, ‘w/o s_{obj} ’ and ‘w/o \hat{m}_{contact} ’, we can observe that gradually including additional conditions aids in generating more appropriate hand poses to the object.

Refiner. Compared to the ‘Ours w/o f^{refiner} ’ that does not involve the refiner, ‘Ours’ provides far better performance, especially in the physical realism of hand and object motions. Also, we demonstrate the effect of losses L_{penet} and L_{contact} by removing them in ‘w/o $L_{\text{penet}} \& L_{\text{contact}}$ ’, ‘w/o L_{contact} ’ and ‘w/o L_{penet} ’, as shown in the same table. Involving more losses consistently improve the performance.

5. Conclusion

In this paper, we propose a novel method for generating the sequence of 3D hand-object interaction from a text prompt and a canonical object mesh. This is achieved through the

three-staged framework that 1) estimates the text-guided and scale-variant contact maps; 2) generates hand-object motions based on a Transformer-based diffusion mechanism; and 3) refines the interaction by considering the penetration and contacts between hands and an object. In experiments, we validate our effectiveness of hand-object interaction generation by comparing it to three baselines where our method outperforms previous methods with strong physical plausibility and accuracy.

Acknowledgements. This work was supported by IITP grants (No. 2020-0-01336 Artificial intelligence graduate school program (UNIST) 10%; No. 2021-0-02068 AI innovation hub 10%; No. 2022-0-00264 Comprehensive video understanding and generation with knowledge-based deep logic neural network 20%) and the NRF grant (No. RS-2023-00252630 20%), all funded by the Korean government (MSIT). This work was also supported by Korea Institute of Marine Science & Technology Promotion (KIMST) funded by Ministry of Oceans and Fisheries (RS-2022-KS221674) 20% and received support from AI Center, CJ Corporation (20%).

References

- [1] Samarth Brahmabhatt, Cusuh Ham, Charles C Kemp, and James Hays. Contactdb: Analyzing and predicting grasp contact via thermal imaging. In *CVPR*, 2019. 2
- [2] Samarth Brahmabhatt, Ankur Handa, James Hays, and Dieter Fox. Contactgrasp: Functional multi-finger grasp synthesis from contact. In *IROS*, 2019. 2
- [3] Xin Chen, Biao Jiang, Wen Liu, Zilong Huang, Bin Fu, Tao Chen, and Gang Yu. Executing your commands via motion diffusion in latent space. In *CVPR*, 2023. 1, 2
- [4] Enric Corona, Albert Pumarola, Guillem Alenya, Francesc Moreno-Noguer, and Grégory Rogez. Ganhand: Predicting human grasp affordances in multi-object scenes. In *CVPR*, 2020. 2
- [5] Zicong Fan, Omid Taheri, Dimitrios Tzionas, Muhammed Kocabas, Manuel Kaufmann, Michael J Black, and Otmar Hilliges. Arctic: A dataset for dexterous bimanual hand-object manipulation. In *CVPR*, 2023. 2, 6
- [6] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. Imos: Intent-driven full-body motion synthesis for human-object interactions. In *Computer Graphics Forum*, 2023. 2, 6, 7
- [7] Patrick Grady, Chengcheng Tang, Christopher D Twigg, Minh Vo, Samarth Brahmabhatt, and Charles C Kemp. Contactopt: Optimizing contact to improve grasps. In *CVPR*, 2021. 2
- [8] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *CVPR*, 2022. 1, 2, 6, 7
- [9] Chuan Guo, Xinxin Zuo, Sen Wang, and Li Cheng. Tm2t: Stochastic and tokenized modeling for the reciprocal generation of 3d human motions and texts. In *ECCV*, 2022. 1, 2
- [10] Shreyas Hampali, Mahdi Rad, Markus Oberweger, and Vincent Lepetit. Honnotate: A method for 3d annotation of hand and object poses. In *CVPR*, 2020. 2
- [11] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NIPS*, 2020. 2, 3, 4
- [12] Biao Jiang, Xin Chen, Wen Liu, Jingyi Yu, Gang Yu, and Tao Chen. Motiongpt: Human motion as a foreign language. *arXiv preprint arXiv:2306.14795*, 2023. 2
- [13] Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong Wang. Hand-object contact consistency reasoning for human grasps generation. In *ICCV*, 2021. 2, 5
- [14] Korrawe Karunratanakul, Jinlong Yang, Yan Zhang, Michael J Black, Krikamol Muandet, and Siyu Tang. Grasping field: Learning implicit representations for human grasps. In *2020 International Conference on 3D Vision (3DV)*, 2020. 2
- [15] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. Flame: Free-form language-based motion synthesis & editing. In *AAAI*, 2023. 1, 2
- [16] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. In *ICCV*, 2021. 2, 6
- [17] Taeryung Lee, Gyeongsik Moon, and Kyoung Mu Lee. Multiact: Long-term 3d human motion generation from multiple action labels. In *AAAI*, 2023. 2
- [18] Haoming Li, Xinzhuo Lin, Yang Zhou, Xiang Li, Yuchi Huo, Jiming Chen, and Qi Ye. Contact2grasp: 3d grasp synthesis via hand-object contact constraint. *arXiv preprint arXiv:2210.09245*, 2022. 2, 3
- [19] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. Intergen: Diffusion-based multi-human motion generation under complex interactions. *arXiv preprint arXiv:2304.05684*, 2023. 2, 5
- [20] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6-dof graspnet: Variational grasp generation for object manipulation. In *ICCV*, 2019. 2
- [21] Mathis Petrovich, Michael J Black, and Gül Varol. Temos: Generating diverse human motions from textual descriptions. In *ECCV*, 2022. 1, 2
- [22] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. 3
- [23] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 3, 6
- [24] Javier Romero, Dimitrios Tzionas, and Michael J Black. Embodied hands: modeling and capturing hands and bodies together. *ACM Transactions on Graphics (TOG)*, 2017. 3
- [25] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *ECCV*, 2020. 2, 6
- [26] Omid Taheri, Yi Zhou, Dimitrios Tzionas, Yang Zhou, Duygu Ceylan, Soren Pirk, and Michael J Black. Grip: Generating interaction poses using latent consistency and spatial cues. *arXiv preprint arXiv:2308.11617*, 2023. 5
- [27] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *ICLR*, 2023. 1, 2, 4, 5, 6, 7
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NIPS*, 2017. 3
- [29] He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. Manipnet: neural manipulation synthesis with a hand-object spatial representation. *ACM Transactions on Graphics (ToG)*, 2021. 2, 3, 6
- [30] Jianrong Zhang, Yangsong Zhang, Xiaodong Cun, Yong Zhang, Hongwei Zhao, Hongtao Lu, Xi Shen, and Ying Shan. Generating human motion from textual descriptions with discrete representations. In *CVPR*, 2023. 1, 2
- [31] Mingyuan Zhang, Zhongang Cai, Liang Pan, Fangzhou Hong, Xinying Guo, Lei Yang, and Ziwei Liu. Motiiondiffuse: Text-driven human motion generation with diffusion model. *arXiv preprint arXiv:2208.15001*, 2022. 1, 2
- [32] Juntian Zheng, Qingyuan Zheng, Lixing Fang, Yun Liu, and Li Yi. Cams: Canonicalized manipulation spaces for category-level functional hand-object manipulation synthesis. In *CVPR*, 2023. 3
- [33] Chongyang Zhong, Lei Hu, Zihao Zhang, and Shihong Xia. Att2m: Text-driven human motion generation with multi-perspective attention mechanism. In *ICCV*, 2023. 2
- [34] Keyang Zhou, Bharat Lal Bhatnagar, Jan Eric Lenssen, and Gerard Pons-Moll. Toch: Spatio-temporal object-to-hand correspondence for motion refinement. In *ECCV*, 2022. 2
- [35] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *CVPR*, 2019. 3