

# Locomotion-Action-Manipulation: Synthesizing Human-Scene Interactions in Complex 3D Environments

Jiye Lee  
 Seoul National University  
 kay2353@snu.ac.kr

Hanbyul Joo  
 Seoul National University  
 hbjoo@snu.ac.kr




Figure 1: Our system, LAMA, produces high-quality and realistic 3D human motions that include locomotion, scene interactions, and manipulations within a given 3D scene and designated interaction cues.

## Abstract

*Synthesizing interaction-involved human motions has been challenging due to the high complexity of 3D environments and the diversity of possible human behaviors within. We present LAMA, Locomotion-Action-MANipulation, to synthesize natural and plausible long-term human movements in complex indoor environments. The key motivation of LAMA is to build a unified framework to encompass a series of everyday motions including locomotion, scene interaction, and object manipulation. Unlike existing methods that require motion data “paired” with scanned 3D scenes for supervision, we formulate the problem as a test-time optimization by using human motion capture data only for synthesis. LAMA leverages a reinforcement learning framework coupled with a motion matching algorithm for optimization, and further exploits a motion editing framework via manifold learning to cover possible variations in interaction and manipulation. Throughout extensive experiments, we demonstrate that LAMA outperforms previous approaches in synthesizing realistic motions in various challenging scenarios. Project page: <https://jiyewise.github.io/projects/LAMA/>.*

## 1. Introduction

Synthesizing interactions within real-life 3D environments has been a challenging research problem due to its

complexity and diversity. The spatial constraint arising from real-life 3D environments where many objects are cluttered makes motion synthesis highly constrained and complex. Furthermore, the nearly indefinite diversity of possible spatial arrangements of the 3D environment and human interaction behaviors makes generalization in synthesis difficult.

Due to the wide range of technical challenges involved in human-scene interactions, previous approaches have focused on sub-problems, such as (1) modeling static poses [24, 66, 68, 16, 61, 45, 69] or (2) human object interactions with a single target object or interaction type [48, 67, 64, 49, 50, 44, 63, 10]. More recent methods [55, 54, 14] extend to synthesizing dynamic interaction motions in real-world 3D scenes, where they use “scene-paired” motion datasets [15] in which motion is simultaneously captured with the surrounding 3D environment. As such paired dataset is rare and difficult to scale up, the performance of these methods is fundamentally limited in fully covering the complexity and diversity of human interaction in real-world 3D scenes.

In this paper, we present LAMA, Locomotion-Action-MANipulation, to synthesize natural and plausible long-term human motions in complex indoor environments. The key motivation of LAMA is to build a unified framework covering a series of everyday motions within real-world 3D scenes: locomotion through cluttered areas, interaction with the scene, and manipulation of objects. Unlike previous approaches [55, 14] that use a “scene-paired” motion

dataset for supervision, we formulate it as a test-time optimization by utilizing only human motion capture data. Exploiting reinforcement learning (RL) as a tool for optimization, we present an RL-based framework coupled with a motion matching algorithm [9, 5] to synthesize locomotion and scene interaction seamlessly while adapting to complex 3D scenes with collision avoidance handling. The object manipulation in our framework is performed via a motion editing approach on top, by learning an autoencoder-based motion manifold space [19]. As a test-time optimization framework, LAMA is applicable to any 3D scene scenarios (e.g., public datasets or any newly scanned scenes). Through extensive quantitative and qualitative evaluations against existing methods, we demonstrate that our method outperforms [55, 14] in various challenging scenarios.

Our contributions are summarized as follows: (1) The first method to generate realistic long-term motions combined with locomotion, scene interaction, and manipulation in complex 3D scenes without “paired” datasets; (2) A novel test-time optimization framework requiring human motion capture data only by incorporating a reinforcement learning framework coupled with motion matching, equipped with well-designed state and rewards for collision avoidance and scene interactions; (3) the state-of-the-art motion synthesis quality with longer duration (near 10 sec); (4) A newly captured and polished motion capture dataset including locomotion and action (e.g., sitting) suitable for motion matching.

## 2. Related Work

**Generating Human-Scene Interactions.** Generating natural human motion has been a widely researched topic in the computer vision community. Early methods focus on synthesizing or predicting human movements by exploiting neural networks [12, 11, 34, 34, 37, 51, 53, 43]. However, these approaches primarily address the synthesis of human motion itself, without taking into account the surrounding 3D environments. Recent approaches begin to tackle modeling and synthesizing human interactions within 3D scenes, or with objects. Many focus on statically posing humans within the given 3D environment [24, 66, 68, 15], by generating human scene interaction poses from various types of input including object semantics [16], images [20, 65, 62, 61, 23, 22], and text descriptions [45, 69].

Recently, there have been approaches to synthesize dynamic human-object interactions (e.g., sitting on chairs, carrying boxes). Starke et al. [48] introduce an autoregressive learning framework with object geometry based environmental encodings to synthesize human-object interactions. Although the encoding includes information on multiple objects within a scene, as demonstrated in [14], no explicit module for navigating through cluttered 3D scenes exists in [48]. Later work [14, 67] extends this by synthesizing




Figure 2: Overview of LAMA.

motions conditioned with variations of objects and contact points. Other approaches [64, 49, 59, 50, 44, 63, 70] focus on generating natural hand movements for manipulation, which is extended by including full body motions [49, 59]. Physics-based character control to synthesize human-object interactions has been also explored in [38, 10, 8, 44, 63, 30]. Although these methods cover a variety of human-object interactions, most of them focus on a specific interaction type or the relationship between the human and the target object without long-term navigation in cluttered 3D scenes.

More recent approaches include generating natural human scene interactions in cluttered 3D scenes [55, 54, 6, 56], closely related to ours. These methods are trained using human motion datasets paired with 3D scenes, which require both ground truth motion and simultaneously captured 3D scenes for supervision. Due to difficulties in acquiring such data, some methods exploit synthetic datasets [6, 56], data fitted from depth videos [55], or motion snapshots with short duration (1-3 sec) [57]. In previous approaches [14, 54], navigation in cluttered environments is often performed by a separate module via path planning (e.g.,  $A^*$  algorithm) by approximating the volume of a human as a cylinder. These path planning based methods approximate the spatial information of the scene and the body and therefore have limitations under highly complex conditions.

**Motion Synthesis and Editing.** Synthesizing natural human motions by leveraging motion capture data has also been a long-researched topic in computer graphics. Some approaches [25, 36] construct motion graphs, where plausible transitions are inserted as edges, and motion synthesis is done by traversing through the graph. Similar approaches [29, 46] connect motion patches to synthesize interactions in a virtual environment or multi-person interactions. Due to its versatility and simplicity, variations have been made to the graph-based approach, such as motion grammar [21] which enforces traversing rules in the motion graph. Motion matching [5, 9] can also be understood as a special case of motion graph traversal, where the plausible transitions are not precomputed but searched during runtime. Recent advances in deep learning allow to leverage motion capture data for motion manifold learning [19, 47, 18]. Autoregressive approaches based on variational autoencoders (VAE) [35, 43] and recurrent neural networks [27, 13, 39] are also used to forecast future motions based on past frames. These frameworks are generalized to synthesize a diverse set of motions including locomotion

on terrains [18] mazes [35], action-specified motions [43], and interaction-involved sports [27, 39]. Neural network-based methods are also reported to be successful in various motion editing tasks such as skeleton retargeting [2], style transfer [19, 3], and in-betweening [13].

Reinforcement learning (RL) has also been successful in combination with both data-driven and physics-based approaches for synthesizing human motions. Combined with data-driven approaches, RL serves as a control module that generates corresponding motions to a given user input by traversing motion graphs [26], latent space [33, 52, 35], and precomputed transition tables [28]. Deep reinforcement learning (DRL) has been widely used as well to synthesize physically plausible movements with a diverse set of motor skills [40, 39, 31, 4, 58, 42, 41, 30]. The key idea of these methods comes from imitation learning, where the control policy in DRL is optimized to actuate the character based on the character’s physical state, to meet the goal of tracking the given reference motion in a physically simulated environment.

### 3. Method

#### 3.1. Overview

Our system, dubbed as *LAMA*, outputs a sequence of human poses  $\mathbf{M} = \{\mathbf{m}_t\}_{t=1}^T$  by taking the 3D scene  $\mathbf{W}$ , desired interaction cues  $\Phi$ , and initial state  $\mathbf{g}_{init} = (\mathbf{p}_{root}^0, \mathbf{r}_{root}^0)$  as inputs:

$$\mathbf{M} = \mathcal{LAMA}(\mathbf{W}, \Phi, \mathbf{g}_{init}), \quad (1)$$

where  $\mathbf{p}_{root}^0 \in \mathbb{R}^3$  and  $\mathbf{r}_{root}^0 \in so(3)$  represent the global position and orientation of the character’s root respectively at initial (i.e.,  $t = 0$ ). The output posture at time  $t$ ,  $\mathbf{m}_t = (\mathbf{p}_{root}^t, \mathbf{r}_{root}^t, \mathbf{r}_1^t, \dots, \mathbf{r}_J^t) \in \mathbb{R}^{3J+6}$ , is represented by a concatenated vector of global position and orientation of the root, and the local joint orientations of  $J$  joints where each  $j$ -th joint is in angle-axis representations  $\mathbf{r}_j^t \in so(3)$ . Throughout our system, the skeleton tree structure and joint offsets are fixed as shown in Fig. 5 (a). We represent the 3D scene  $\mathbf{W} = \{\mathbf{w}_i\}$  as a set of 3D object and scene meshes, including the background scene mesh and other object meshes targeted for manipulation.

The interaction cues,  $\Phi = [\phi_A, \phi_M]$ , represent the expected goal that the output needs to fulfill, and consist of the action cue  $\phi_A$  for the action task (e.g., sitting) and the manipulation cue  $\phi_M$  for the manipulation task. The action cue  $\phi_A = \{\mathbf{q}_{root}, \mathbf{r}_{root}, \mathbf{q}_{lFoot}, \mathbf{q}_{rFoot}\}$  indicates desired position and orientation of the root, and the positions of the left foot and right foot respectively (i.e.,  $\mathbf{q}_j \in \mathbb{R}^3$ ). The foot positions are optional and can be automatically determined if not provided.  $\phi_A$  can be manually chosen to instruct the character or can be automatically given via an off-the-shelf estimator such as GoalNet in [14]. The manipulation cue

$\phi_M = \{\mathbf{q}_j^t\}_{j \in J_M}$  indicates desired locations of selected joints  $J_M$  for control, which is mainly used in the motion editing procedure (Sec. 3.6). For example,  $\phi_M$  can specify the hand joint trajectory to perform the opening laptop motion in Fig. 1 and Fig. 6. Examples of the 3D scene  $\mathbf{W}$ , action cue  $\phi_A$  and manipulation cue  $\phi_M$  are in Fig. 6 (left).

LAMA is designed via a three-level system composed of action controller **A**, motion synthesizer **S**, followed by a manifold-based motion editor **E**. The locomotion and action parts are seamlessly performed via the action controller **A** and synthesizer **S**. The essential idea in our design is to combine the RL framework with motion matching [9, 5] to synthesize realistic human motions while fulfilling the desired scene interaction tasks. By taking 3D scene  $\mathbf{W}$ , action cue  $\phi_A$ , and initial state  $\mathbf{g}_{init}$  as input, the action controller **A** makes the use of RL as a way of test-time optimization to synthesize corresponding motion. A control policy  $\pi$  is optimized<sup>1</sup> to sample an action at time  $t$ ,  $\pi(\mathbf{a}_t | \mathbf{s}_t, \mathbf{W}, \phi_A)$ , where  $\mathbf{a}_t$  indicates the plausible next action containing predicted action types and short-term future forecasting.  $\mathbf{s}_t$  is the state cue to represent the current status of the human character including its body posture, surrounding scene occupancy, and the targeting action cue. Intuitively, action controller **A** is optimized to generate plausible next action  $\mathbf{a}_t$  by considering the current character-scene state  $\mathbf{s}_t$ . The generated action signal  $\mathbf{a}_t$  from the action controller **A** is provided as input to the motion synthesizer **S**, which then determines the posture at the next time step  $\mathbf{m}_{t+1}$ , i.e.,  $\mathbf{S}(\mathbf{m}_t, \mathbf{a}_t) = \mathbf{m}_{t+1}$ . The character’s next state  $\mathbf{s}_{t+1}$  can be computed again from  $\mathbf{m}_{t+1}$ , which is subsequently taken by the action controller **A** as an input for the next time frame.

The initial output motion  $\mathbf{M}$  synthesized by **A** and **S** is followed by a motion editor  $\mathbf{E}(\mathbf{M}) = \tilde{\mathbf{M}}$ , where  $\tilde{\mathbf{M}} = \{\tilde{\mathbf{M}}_t\}_{t=1}^T$  is the edited motion. The goal of the editing module **E** is to (1) post-process  $\mathbf{M}$  to fit into diverse objects targeted for action task  $\phi_A$  (e.g., sitting on a chair with different heights), and (2) perform human-object manipulation instructed by  $\phi_M$  (e.g. moving objects, opening doors). Fig. 2 shows the overview of LAMA.

#### 3.2. Scene-Aware Action Controller

Unlike previous approaches [56, 14] that use path planning for navigation and a learning-based module trained with a scene-paired motion dataset for interaction, our action controller **A** performs locomotion and desired actions seamlessly by fulfilling the action cue  $\phi_A$  and avoiding collisions in the 3D scene  $\mathbf{W}$ . Importantly, given the scene  $\mathbf{W}$ , action  $\phi_A$  and initial state  $\mathbf{g}_{init}$  as inputs, our action controller is directly optimized to choose the most plausible motion clip in our motion database at each state, synthesizing natural human motions while taking the 3D scene into

<sup>1</sup>We use the term “optimized” rather than “learned” for the policy since we perform a test-time optimization.

account without any scene-paired motion dataset or training procedure. Intuitively, given the current state  $\mathbf{s}_t$ , the goal of the action controller is to output the best next action  $\mathbf{a}_t$  which is used to search the next motion clip in Motion Synthesizer S.

**State.** The state  $\mathbf{s}_t = \psi(\mathbf{m}_{t-1}, \mathbf{m}_t, \mathbf{W}, \phi_A)$  at time  $t$  is a feature vector representing the current status of the human character, where  $\psi$  is the function to compute the state from other inputs.  $\mathbf{s}_t = (\mathbf{s}_t^{body}, \mathbf{s}_t^{scene}, \mathbf{s}_t^{inter})$  is composed of body configuration  $\mathbf{s}^{body}$ , scene occupancy  $\mathbf{s}^{scene}$ , and desired current target interaction  $\mathbf{s}^{inter}$ . Body configuration  $\mathbf{s}^{body} = \{\mathbf{r}, \dot{\mathbf{r}}, \theta_{up}, h, \mathbf{p}_e\}$ , where  $\mathbf{r}, \dot{\mathbf{r}} \in \mathbb{R}^{J \times 6}$  are the joint rotations and velocities respectively for the  $J$  joints excluding the root in 6D representations [71],  $\theta_{up} \in \mathbb{R}$  is the up vector of the root (represented by the angle w.r.t the Y-axis),  $h \in \mathbb{R}$  is the root height from the floor, and  $\mathbf{p}_e \in \mathbb{R}^{e \times 3}$  is the end-effector position in person-centric coordinate (where  $e$  is the number of end-effectors).  $\mathbf{s}^{scene} = \{\mathbf{g}_{occ}, \mathbf{g}_{root}\}$  includes scene occupancy information in the floor plane, as shown in Fig. 4.  $\mathbf{g}_{occ} \in \mathbb{R}^{n^2}$  represents the occupancy grid on the floor plane of neighboring  $n$  cells around the agent and  $\mathbf{g}_{root} \in \mathbb{R}^2$  denotes the current global root position of the character in the discretized grid plane. Note that, while we consider the 2D floor grid for efficiency rather than 3D, the 3D scene is still considered via our collision reward term in Sec. 3.4.  $\mathbf{s}^{inter}$  represents the action cue the character is targeting, that is  $\mathbf{s}^{inter} = \phi_A$ .

**Action.** Given the current status of the character  $\mathbf{s}_t$ , the control policy  $\pi$  outputs the feasible action  $\mathbf{a}_t = (\mathbf{a}_t^{type}, \mathbf{a}_t^{future}, \mathbf{a}_t^{offset})$ .  $\mathbf{a}_t^{type}$  provides the probabilities of the next action type among possible actions (e.g., walk, sit, or stop), determining the transition timing between actions (e.g., from locomotion to sitting).  $\mathbf{a}_t^{future}$  predicts future motion cues such as plausible root position for the next 10, 20, and 30 frames. Posture offset  $\mathbf{a}_t^{offset}$  is intended to modify the raw motion data searched from the motion database in motion synthesizer module S. Intuitively, our optimized control policy generates a posture offset  $\mathbf{a}_t^{offset}$  to alter the closest plausible raw posture chosen in the database. This enables the character to perform more plausible scene-aware human poses only with human motion data. More details are addressed in Sec. 3.3.

### 3.3. Motion Synthesizer

By taking the current posture  $\mathbf{m}_t$  and actions signal  $\mathbf{a}_t$  from the action controller A as inputs, the motion synthesizer S produces the next plausible posture:  $\mathbf{S}(\mathbf{m}_t, \mathbf{a}_t) = \mathbf{m}_{t+1}$ . As the first step, the motion synthesizer searches for motion from a motion database that best matches the closest motion feature, then modifies the searched raw motion to be more suitable to the scene. To this end, the motion synthesizer’s output  $\mathbf{m}_{t+1}$  is in turn fed into the action controller recursively. We exploit a modified version of the mo-




Figure 3: Visualization of the relationship between the Action Controller and the Motion Synthesizer.

tion matching algorithm [5, 9, 17] for the first step of motion synthesis. In motion matching, motion synthesis is performed periodically by searching the most plausible next shot motion segments from a motion database, and compositing them into a long connected sequence.

**Motion features.** Motion feature  $\mathbf{y}_t$  represents the characteristic of each frame in the short motion segment and is computed as  $f(\mathbf{m}) = \mathbf{y}_t = \{\{\mathbf{p}_j\}, \{\dot{\mathbf{p}}_j\}, \theta_{up}, c, \mathbf{o}_{future}\}$ . From a posture  $\mathbf{m}$ , the positions and velocities  $\mathbf{p}_j, \dot{\mathbf{p}}_j \in \mathbb{R}^3$  are extracted for the selected joints  $j \in \{\text{Head}, \text{Hand}, \text{Foot}\}$ , which are defined in a person-centric coordinate of  $\mathbf{m}$ .  $\theta_{up} \in \mathbb{R}^3$  is the up-vector of the root joint, and  $c \in \{0, 0.5, 1\}$  indicates automatically computed foot contact cues of the left and right foot (0 for non-contact, 1 for contact, 0.5 for non-contact but close to the floor within a threshold).  $\mathbf{o}_{future} = \{\{\mathbf{p}_{root}^{\Delta t}\}, \{\mathbf{r}_{root}^{\Delta t}\}\}$  contains the cues for short-term future postures, where  $\mathbf{p}_{root}^{\Delta t}$  and  $\mathbf{r}_{root}^{\Delta t}$  are the position and orientation of root joint at  $\Delta t$  frames later from the current target frame.  $\mathbf{o}_{future}$  are computed in 2D XZ plane in person-centric coordinate of the current target motion  $\mathbf{m}$ , and thus  $\mathbf{p}_{root}^{\Delta t}, \mathbf{r}_{root}^{\Delta t} \in \mathbb{R}^2$ . The selected future frames are action-type specific, and for locomotion, we extract 10, 20, and 30 frames in the future (at 30Hz) following [9]. Intuitively, the motion feature extracts the target frame’s posture and temporal cues by considering neighboring frames. For efficiency, we pre-compute motion features  $\mathbf{y}_t$  for every frame of the motion clip in the database.

Motion feature  $\mathbf{x}_t$  of the current state of the character, or the query feature denoted, is also computed in the same way based on posture  $\mathbf{m}_{t-1}, \mathbf{m}_t$  and  $\mathbf{a}_t^{future}$  produced by the action controller, that is  $\mathbf{x}_t = f(\mathbf{m}_{t-1}, \mathbf{m}_t, \mathbf{a}_t^{type}, \mathbf{a}_t^{future})$ . The component  $\mathbf{a}_t^{future}$  serves as  $\mathbf{o}_{future}$  in the query feature, which can be understood as the action controller providing cues for predicted future postures.

**Motion searching and updating.** Given the query motion feature  $\mathbf{x}_t$  and the motion features  $\mathbf{y}_k$  in the motion database (where  $k$  is the index of the clip), motion searching finds the best matches  $k^*$  in the motion database by computing the weighted euclidean distances between the query feature and motion database features:

$$k^* = \arg \min_k \|\mathbf{w}_f^T (\mathbf{x}_t - \mathbf{y}_k)\|^2, \quad (2)$$




Figure 4: Visual representation of the occupancy grid. Grid on the right represents top view. Gray and black are occupied cells while blue indicates the root.

where  $\mathbf{w}_f$  is a fixed weight vector to control the importance of feature elements. After finding the best match  $\hat{\mathbf{m}}_{k^*}$  from the motion database, the motion synthesizer updates it with the predicted motion offset  $\mathbf{a}_t^{offset}$  from  $\mathbf{a}_t$ , that is  $\tau(\hat{\mathbf{m}}_{k^*+1}, \mathbf{a}_{offset}) = \mathbf{m}_{t+1}$ , where  $\hat{\mathbf{m}}_{k^*+1}$  is the next plausible character posture and  $\tau$  is an update function to update selected joints in  $\hat{\mathbf{m}}_{k^*+1}$ . In practice, motion searching is performed periodically (e.g., every N-th frame) to make the synthesized motion temporally more coherent.

### 3.4. Optimizing Scene-Aware Action Controller

The objective of our reinforcement learning framework is to optimize the policy by maximizing the discounted cumulative reward. In our method, we design the rewards to guide the character to perform both locomotion and desired actions (e.g., sitting) under common constraints (e.g., smooth transitions, and collision avoidance). Our reward function consists of the following terms:

$$R_{total} = w_{tr}R_{tr} + w_{act}R_{act} + w_{reg}R_{reg}, \quad (3)$$

where  $w_{tr}$ ,  $w_{act}$ , and  $w_{reg}$  are the weights to balance among reward terms. The trajectory reward  $R_{tr}$  is obtained when the character moves towards action  $\phi_A$  while meeting the spatial constraints from the 3D scene, described below:

$$R_{tr} = r_{coli} \cdot r_{pos} \cdot r_{root}, \quad \text{where} \quad (4)$$

$$r_{coli} = \exp\left(-\frac{1}{\sigma_{coli}^2} \sum_{b \in \mathbf{B}} w_b \rho(b, \mathbf{W})\right), \quad (5)$$

$$r_{pos} = \exp\left(-\frac{1}{\sigma_{root}^2} \sum_{j \in \mathbf{J}} \|\mathbf{p}_0 - \mathbf{q}_j\|^2\right), \quad (6)$$

$$r_{vel} = \begin{cases} 1 & \text{when } \dot{\mathbf{p}}_{root} \geq \sigma_{th} \\ \sigma_{vel} \|\dot{\mathbf{p}}_0\|^2 & \text{else.} \end{cases} \quad (7)$$

The collision-avoidance reward  $r_{coli}$  penalizes collisions with 3D scenes. As depicted in Fig. 5 (a), body limbs in the skeletal structure are represented as a set of box-shaped nodes  $\mathbf{B}$  with a fixed width, where each element  $b \in \mathbf{B}$  is a 3D box representation of legs and arms (we exclude torso and head). The function  $\rho(b, \mathbf{W})$  detects the collision between edges of a box-shaped node  $b$  with 3D scene  $\mathbf{W}$  and returns the number of intersection points. (Fig. 5 (b)).  $w_b$




Figure 5: (a) Skeleton with joints and box nodes. (b) Automatically detected collision points (colored as red).

is the weight to control the importance of each limb  $b$ . The collision-avoidance reward is maximized when no penetration occurs, enforcing the policy  $\pi$  to generate adequate action  $\mathbf{a}_t$  to avoid physically implausible penetrations.  $r_{pos}$  is obtained when the agent navigates toward targeting action cue  $\phi_A$ .  $r_{vel}$  encourages the character to move by penalizing when the root velocity  $\dot{\mathbf{p}}_{root}$  is less than a threshold  $\sigma_{th}$ .  $\sigma_{coli}$ ,  $\sigma_{root}$ , and  $\sigma_{vel}$  are weights to control balance.

Action reward  $R_{act}$  encourages to fulfill the given action cue  $\phi_A = \{\mathbf{q}_{root}, \mathbf{r}_{root}, \mathbf{q}_{rFoot}, \mathbf{q}_{lFoot}\}$ :

$$R_{act} = r_{inter} \cdot r_{\Delta t} \cdot r_{\Delta v}, \quad \text{where}$$

$$r_{inter} = \exp\left(-\frac{1}{\sigma_{inter}^2} \sum_{j \in J_A} \|\mathbf{p}_j - \mathbf{q}_j\|^2\right), \quad (8)$$

$$r_{\Delta t} = \exp(-\sigma_{\Delta t}^2 C_{tr}), \quad r_{\Delta v} = \exp(-\sigma_{\Delta v}^2 C_{vel})$$

where interaction reward term  $r_{inter}$  is given when the character switches from navigation to corresponding action to  $\phi_A$  and is maximized when the performed action meets the positional constraints of  $\phi_A$ . Smoothness reward terms  $r_{\Delta t}$  and  $r_{\Delta v}$  minimize the transition cost, which is based on the subpart of the feature distances defined in Eq. 2, where  $C_{tr}$  is the weighted feature distances of  $p_j$ ,  $\theta_{up}$ , and  $c$ , and  $C_{vel}$  is from  $\dot{\mathbf{p}}$ . These are intended to discourage the character from making abrupt changes.

Regularization reward  $R_{reg}$  penalizes the  $\mathbf{a}_t^{offset}$  excessively modifying the original posture searched in the motion database of  $\mathbf{S}$ , denoted as  $\hat{\mathbf{m}}_t$ , and maintains temporal consistency among frames.

$$R_{reg} = \exp\left(-\frac{1}{\sigma_{reg}^2} (\|\hat{\mathbf{m}}_t - \mathbf{m}_t\|^2 + \|\mathbf{m}_t - \mathbf{m}_{t-1}\|^2)\right).$$

As reported in [32, 39], multiplying rewards with consistent goals can enforce all reward terms to be simultaneously met. We also use early termination [40] and limited action transitions to accelerate learning. Details are in supp. mat.

### 3.5. Generalizing Action Controller

While our major focus of the use of RL is for a test-time optimization given a single target task, the optimized policy can handle variations of the task to some extent, as an advantage of the nature of RL. As shown in our experiments in Sec. 4.4, we demonstrate that our optimized controller

can be directly used for various action cues  $\phi_A$  and initials  $\mathbf{g}_{init}$  without further optimization for the same scene  $\mathbf{W}$ .

As an extension of our framework, we can make our controller more generalized by optimizing the policy with random variations of inputs,  $\mathbf{g}_{init}$  and  $\phi_A$  per each episode during policy optimization. This procedure is more similar to the usual RL framework, where the policy is “learned” in advance for the target scene  $\mathbf{W}$ , and applied to the provided inputs during inference. We also demonstrate that our controller can handle a wider range of input variations via this augmentation process. This extension of our framework can provide better efficiency for the cases where varying tasks are instructed under a fixed 3D scene  $\mathbf{W}$ . As shown in Sec. 4.4, via the generalization process we can directly use the policy for diverse inputs without further optimization. Or, if necessary, efficiently fine-tuning the policy is also possible. Note that this extension still differs from other learning-based methods [55, 14] in that we do not require any scene-paired motion datasets or other supervision.

### 3.6. Task-Adaptive Motion Editing

To cover the diversity in interactions, we include a task-adaptive motion editing module in our motion synthesis framework. In particular, in the case of object manipulation, manipulation cue  $\phi_M$  is provided to enforce an end-effector (e.g., a hand) to follow the desired trajectory expressing the manipulation task on the target object, as in Fig 6. The manipulation cue  $\phi_M$  can be provided via any possible way, and in our experiments we produce it semi-automatically. We compute the desired trajectory by simulating the target articulated object’s motion [60] by considering a contact point on the surface of the object mesh.

Not only the edited motion  $\bar{\mathbf{M}} = E(\mathbf{M})$  should fulfill the sparsely given positional constraints, it should also preserve the temporal consistency and spatial correlations among joints to maintain its naturalness. We adopt the motion manifold learning approach with convolutional autoencoders [19] to compress motion to a latent vector within a motion manifold space. Motion editing is done by searching an optimal latent vector among the manifold. For training the autoencoder, motion sequence, which we denote as  $\mathbf{X}$  converted from  $\mathbf{M}$ , is represented as a time-series of postures by concatenating joint rotations in 6D representations [71], root height, root transform relative to the previous frame projected on the XZ plane, and foot contact labels. The encoder and decoder module are trained based on reconstruction loss,  $\|\mathbf{X} - \Psi^{-1}(\Psi(\mathbf{X}))\|^2$ , where  $\Psi$  is the encoder and  $\Psi^{-1}$  is the decoder.

The latent vector from the encoder  $\mathbf{z} = \Psi(\mathbf{X})$  from the motion manifold space preserves the spatiotemporal relationship among joints and frames found in natural human motions. As demonstrated in [19], editing motions within the manifold space ensures the edited motion to be realistic




Figure 6: Visual representation of system input  $\Phi, \mathbf{W}$  and output  $\bar{\mathbf{M}}$ . On the left, action cue  $\phi_A$  and manipulation cue  $\phi_M$  are shown as red and cyan, respectively. The right is the synthesized motion  $\bar{\mathbf{M}}$ .

and coherent. The optimal latent vector  $\mathbf{z}^*$  is found by minimizing a loss function  $\mathcal{L}$  by constraining the output motions to follow the manipulation constraint  $\phi_M$ . We also include additional regularizers in  $\mathcal{L}$  so that the output motion can maintain the foot locations and root trajectories of the original motion. See supp. mat. for more details on  $\mathcal{L}$ . Finally, the edited motion  $\bar{\mathbf{M}}$  can be computed via  $\Psi^{-1}(\mathbf{z}^*)$ .

## 4. Experiments

We evaluate LAMA’s ability on synthesizing long-term motions in real-world 3D scenes with various human-scene and object interactions involved. We exploit an extensive set of quantitative metrics and perceptual studies for evaluation.

**Dataset.** For constructing the database for the motion synthesizer, we capture a new motion capture dataset involving locomotion and action. Motion is captured with IMU-based system XSens MVN Link [1]. The collected data include high quality human motion with locomotion and interaction in various scenarios, such as walking around at different angles and sitting on a chair with random starting points. Captured motion data are post-processed to be suitable for motion matching. All the data used in this system are motion capture data (in bvh format) with no scene or object related prior information. We use PROX [15] and Matterport3D [7] datasets for 3D scenes and SAPIEN [60] object meshes for manipulation. See supp. mat. for details.

### 4.1. Experimental Setup

**Evaluation metrics.** As our system does not rely on supervision for motion synthesis, quantifying synthesized quality is challenging due to the lack of ground-truth data or official evaluation metrics. We try to evaluate in terms of physical plausibility and naturalness.

- **Physical Plausibility:** We use contact and penetration metrics to evaluate the physical plausibility of the synthesized motions. Contact penalizes the foot movement when the foot is in contact. Since foot contact is a critical element in dynamics, the contact-based metric is closely related to determining the physical plausibility of motions. Penetration loss (“Penetration” in Table 1) measures implausible cases when the body penetrates the objects in the scene. We




Figure 7: Examples of motions which include locomotion, action, and manipulation. Top: opening, closing a trash can lid and sitting on a chair. Bottom: opening door and sitting on a chair.

Method	Contact	Penetration	Naturalness
Wang et al. [55]	6.32	2.75	16.27
Wang et al. [55]*	22.98	14.73	-
SAMP [14]	11.75	7.18	42.04
LAMA (ours)	<b>4.34</b>	<b>1.30</b>	<b>100</b>

Table 1: **Baseline Comparison** Foot contact (cm, ↓) averaged over all frames and penetration (percentage, ↓) score. Naturalness score (percentage, ↑) indicates selection ratio relative to LAMA (for LAMA, set to 100). Wang et al. with an asterisk indicates without post-processing.

compute penetration metric by counting frames where intersection points (Sec. 3.2) go over a certain threshold.<sup>2</sup>

- **Naturalness:** We evaluate the naturalness of the synthesized motion via perception study (A/B test) on Amazon Mechanical Turk. The motions used for testing are rendered with the exact same view and 3D characters, making them indistinguishable from the appearance side. Human observers are asked to choose a more natural motion based on two criteria: (1) the character movement is human-like and (2) the movement is plausible in the given scene. Details of the study setup are in supp. mat.

**Baselines.** We compare LAMA with the state-of-the-art methods as well as variations of ours.

- **Wang et al. [55]** is the state-of-the-art long-term motion synthesis method for human-scene interactions within a given 3D scene. We use the author’s code for evaluation. As Wang et al. post-processes synthesized motion to im-

<sup>2</sup>10 for legs and 7 for arms




Figure 8: Comparison with LAMA (left) and LAMA without collision reward (right). Without collision reward the character fails to avoid collisions with obstacles (red).

prove foot contact and reduce collisions which are directly related to our metric, we both compare Wang et al. with and without post-processing.

- **SAMP [14]** generates interactions that can be generalized not only for object variations but also random starting points within a given 3D scene. SAMP explicitly exploits path planning to navigate through cluttered 3D scenes.
- **Ablative Baselines** We perform ablation studies on the action controller and motion editing module. We perform ablation studies on the scene reward  $r_{coli}$ , and action offset  $\mathbf{a}_t^{offset}$  to present the contribution of both terms on generating scene-aware motions. We also compare our method without the transition reward  $r_{\Delta t}$  and  $r_{\Delta v}$  terms (Sec. 3.2) of the action controller. Finally, we demonstrate the strength of our motion editing module to edit motions naturally (Sec. 3.6) by comparing it with inverse kinematics (IK).

## 4.2. Comparisons with Previous Work

**Evaluation Setup and Details.** For comparison with baselines, we generate 50 motion sequences in total with random input  $\mathbf{g}_{init}$  and  $\phi_A$  from 4 PROX 3D scenes used




Figure 9: Comparison with LAMA (left) and LAMA without action offset (right). The character in original LAMA moves forward while tilting its arms to avoid collision with walls, while in LAMA without action offset does not.

in testing for Wang et al. [55]. Since our method is based on test-time optimization without explicit training and testing split, our action controller is optimized per each input, and no prior information on inputs is given before policy optimization. It takes 4 to 20 minutes to optimize a policy and 3 to 4 minutes (500 epochs) for optimization in the motion editing module. We only consider locomotion and action (walk-to-sit) motions and do not include manipulation as the baselines do not tackle manipulation. Contact metric is measured by the position difference of foot in contact, where contact is automatically labeled based on foot velocity. To compute penetration metric in a fair way, SMPL-X outputs of Wang et al. and SAMP are converted to box-shaped skeletons as in ours and intersection points are counted. Table 1 shows the results.

**Physical Plausibility.** As shown, LAMA outperforms both Wang et al. and SAMP in physical plausibility. Wang et al. post-processes the synthesized motion to ensure contact and reduce penetration, yet LAMA still outperforms. Moreover, our RL-based method with motion matching shows its advantage in collision avoidance in cluttered 3D scenes compared to path-planning based navigation in SAMP.

**Naturalness.** For perception study, we build two separate sets for comparison with Wang et al. and SAMP, and each testset is done with non-overlapping participants. For 50 motion sequences per set, 5 unique responses are collected per sequence for comparison. With Wang et al., LAMA received 215 votes while Wang et al received 35 (relative ratio 16.27%). With SAMP, LAMA received 176 votes, SAMP received 74 (relative ratio 42.04%). The results demonstrate that our method greatly outperforms baselines in terms of naturalness as well.

### 4.3. Ablation Studies

**Ablation Studies on Action Controller.** For quantitative ablations, we compare the original LAMA and the LAMA without collision reward  $r_{coli}$ . Ablation studies are performed in 5 PROX scenes. In original LAMA, penetration occurs in only 1.1% of the frames among the whole motion sequence, while the ratio is 15.7% in LAMA without  $r_{coli}$ . The result supports that the  $r_{coli}$  enforces the action controller to synthesize motions according to the given 3D scene. Example results are shown in Fig. 8. We




Figure 10: (a) Comparison with LAMA (top) and LAMA without manifold and replaced with IK (bottom) of a character opening the toilet lid. (b) Comparison with LAMA (top) and LAMA without motion editing (bottom) in sitting.




Figure 11: Examples of synthesized manipulation motions. The target object for manipulation is colored as orange, the character purple at start and aqua at the end. Left: walking and opening a toilet lid. Right: walking and opening doors.

also qualitatively compare the contribution of other components in the action controller. As seen in Fig. 9, without action offset  $a_t^{offset}$  the character does not tilt its limbs to avoid penetration with objects or walls, as the raw motion brought from the motion database does not have any information about the scene. This shows that  $a_t^{offset}$  also plays a role in generating detailed scene-aware poses. Moreover, the results without smoothness rewards  $r_{\Delta t}$  and  $r_{\Delta v}$  are not smooth enough, showing unnatural and abrupt movements.

**Ablation Studies on Task-Adaptive Motion Editing.** We ablate our motion editing module by replacing it with an alternative approach via IK. Same as  $\phi_M$ , only the trajectory of a joint in contact (e.g., the right hand) is given to the IK solver. As shown in Fig. 10 (left), LAMA with motion editing module shows natural moves such as bending knees and tilting hips to make contact. However, results with IK show awkward poses as such spatiotemporal correlations in natural human motions are not reflected in the IK solver. Furthermore, as seen in Fig. 10 (right), the motion editing module makes the character properly sit in chairs with different shapes.

### 4.4. Robustness Test of Action Controller

As described in Sec. 3.5, utilizing RL for test-time optimization allows the optimized policy to handle variations in input. In this experiment, we aim to measure the extent






Figure 12: Visualization of the range of a policy can cover, with and without generalization. Colored points indicate initial starting point  $\mathbf{g}_{init}$  where the policy can synthesize motions meeting the action cue (white).

to which a policy optimized for a single task  $\phi_A$  and initial  $\mathbf{g}_{init}$  can generalize to varying inputs. To test the robustness with varying initials and tasks, we apply the optimized policy to all possible input variations in the scene and count the number of inputs the policy succeeds in synthesizing. From all possible initials sampled, the colored points in Fig 12 illustrate initial starting locations where the policy can synthesize motions meeting the given action cue  $\phi_A$ . As shown in Table 2, a policy initially optimized for a single set of inputs (red in Fig. 12) can successfully synthesize motions even with distinct set of inputs without any additional optimization. Furthermore, to test the robustness of the generalized policy (described in Sec. 3.5), we perform the same test with the policy trained with our augmentation strategy during optimization. As shown above, it shows even more robustness in variations as expected.

We further demonstrate the generalization ability among unseen scenes with a policy optimized with the augmentation strategy. The generalized policy (Sec. 3.5) optimized in scene  $\mathbf{W}_0$  (scene in Fig. 12) are tested on two unseen scenes  $\mathbf{W}_1$  and  $\mathbf{W}_2$  from PROX [15] shown in Fig. 13. As demonstrated in Table 3, an generalized policy (Sec. 3.5) optimized with scene  $\mathbf{W}_0$  can be generalized to some extent to scene  $\mathbf{W}_1$ , as  $\mathbf{W}_0$  and  $\mathbf{W}_1$  shares a similar structure (sofa and chairs around a table). However, as expected, the generalization ability decreases when tested on a totally distinct scene  $\mathbf{W}_2$ .

Note that the inference time here is about 0.2-3 sec per input as no further policy optimization is required. Details of the test setup are in supp. mat.

## 5. Discussion

We present a unified framework to synthesize human motions within complex real-world 3D scenes with motion-only datasets. We formulate it as a test-time optimization,

Method	$\phi_A^1$	$\phi_A^2$	$\phi_A^3$	$\phi_A^4$
Original Policy	10.5%	25.9%	25.5%	13.9%
Generalized Policy	40.2%	81.7%	77.0%	71.7%

Table 2: **Robustness Test.** Ratio of  $\mathbf{g}_{init}$  (percentage, out of total valid  $\mathbf{g}_{init}$  within the scene) which the optimized policy succeeds in synthesizing motion fulfilling  $\phi_A^n$ .




Figure 13: Robustness test on unseen scenes  $\mathbf{W}_1$  and  $\mathbf{W}_2$ . Colored points represent initials where the policy can synthesize motions meeting the action cue.

Method	$\mathbf{W}_1, \phi_A^1$	$\mathbf{W}_1, \phi_A^2$	$\mathbf{W}_2, \phi_A^1$	$\mathbf{W}_2, \phi_A^2$
-	70.5%	53.4%	11.9%	21.2%

Table 3: **Robustness Test in Unseen Scenes.** Ratio of  $\mathbf{g}_{init}$  (percentage, out of total valid  $\mathbf{g}_{init}$  within the scene) which the optimized policy succeeds in synthesizing motion fulfilling  $\phi_A^n$  in scene  $\mathbf{W}_n$ .

leveraging RL with motion matching for realistic motion synthesis, and also utilize motion manifold to further cover the diversity of manipulation behaviors. Our method has been thoroughly evaluated in diverse scenarios, outperforming previous approaches [55, 14].

Despite RL is used for test-time optimization, a single policy can cover variations in input and can also be generalized for extensive variations. Combining this framework with supervised learning for further efficiency increase can be an interesting future research direction. Furthermore, although we assume a fixed skeleton throughout the system, interaction motions may change depending on the character’s body shapes and sizes. We leave synthesizing motions on varying body shapes as future work.

**Acknowledgements.** This work was supported by SNU-Naver Hyperscale AI Center, SNU Creative-Pioneering Researchers Program, NRF grant funded by the Korea government (MSIT) (No. 2022R1A2C2092724), and IITP grant funded by the Korea government (MSIT) (No.2022-0-00156 and No.2021-0-01343). H. Joo is the corresponding author.

## References

- [1] Xsens mvn link. <https://www.movella.com/products/motion-capture/xsens-mvn-link>. 6
- [2] Kfir Aberman, Peizhuo Li, Dani Lischinski, Olga Sorkine-Hornung, Daniel Cohen-Or, and Baoquan Chen. Skeleton-aware networks for deep motion retargeting. *ACM Trans. Graph.*, 39(4), 2020. 3
- [3] Kfir Aberman, Yijia Weng, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Unpaired motion style transfer from video to animation. *ACM Trans. Graph.*, 39(4), 2020. 3
- [4] Kevin Bergamin, Simon Clavet, Daniel Holden, and James Richard Forbes. Drecon: data-driven responsive control of physics-based characters. *ACM Trans. Graph.*, 38(6), 2019. 3
- [5] Michael Büttner and Simon Clavet. Motion matching - the road to next gen animation. In *Proc. of Nucl.ai*, 2015. 2, 3, 4
- [6] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *ECCV*, 2020. 2
- [7] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *3DV*, 2017. 6
- [8] Yu-Wei Chao, Jimei Yang, Weifeng Chen, and Jia Deng. Learning to sit: Synthesizing human-chair interactions via hierarchical control. In *AAAI*, 2021. 2
- [9] Simon Clavet. Motion matching and the road to next-gen animation. In *Proc. of GDC*, 2016. 2, 3, 4
- [10] Haegwang Eom, Daseong Han, Joseph S Shin, and Junyong Noh. Model predictive control with a visuomotor system for physics-based character animation. *ACM Trans. Graph.*, 39(1), 2019. 1, 2
- [11] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *ICCV*, 2015. 2
- [12] Ikhsanul Habibie, Daniel Holden, Jonathan Schwarz, Joe Yearsley, and Taku Komura. A recurrent variational autoencoder for human motion synthesis. In *BMVC*, 2017. 2
- [13] Félix G Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion in-betweening. *ACM Trans. Graph.*, 39(4), 2020. 2, 3
- [14] Mohamed Hassan, Duygu Ceylan, Ruben Villegas, Jun Saito, Jimei Yang, Yi Zhou, and Michael Black. Stochastic scene-aware motion prediction. In *ICCV*, 2021. 1, 2, 3, 6, 7, 9
- [15] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J. Black. Resolving 3D human pose ambiguities with 3D scene constraints. In *ICCV*, 2019. 1, 2, 6, 9
- [16] Mohamed Hassan, Partha Ghosh, Joachim Tesch, Dimitrios Tzionas, and Michael J Black. Populating 3d scenes by learning human-scene interaction. In *CVPR*, 2021. 1, 2
- [17] Daniel Holden, Oussama Kanoun, Maksym Perepichka, and Tiberiu Popa. Learned motion matching. *ACM Trans. Graph.*, 39(4), 2020. 4
- [18] Daniel Holden, Taku Komura, and Jun Saito. Phase-functioned neural networks for character control. *ACM Trans. Graph.*, 36(4), 2017. 2, 3
- [19] Daniel Holden, Jun Saito, and Taku Komura. A deep learning framework for character motion synthesis and editing. *ACM Trans. Graph.*, 35(4), 2016. 2, 3, 6
- [20] Chun-Hao P Huang, Hongwei Yi, Markus Höschle, Matvey Safroshkin, Tsvetelina Alexiadis, Senya Polikovsky, Daniel Scharstein, and Michael J Black. Capturing and inferring dense full-body human-scene contact. In *CVPR*, 2022. 2
- [21] Kyunglyul Hyun, Kyungho Lee, and Jehee Lee. Motion grammars for character animation. In *Computer Graphics Forum*, volume 35, 2016. 2
- [22] Nan Jiang, Tengyu Liu, Zhexuan Cao, Jieming Cui, Yixin Chen, He Wang, Yixin Zhu, and Siyuan Huang. Chairs: Towards full-body articulated human-object interaction. *arXiv preprint arXiv:2212.10621*, 2022. 2
- [23] Yuheng Jiang, Suyi Jiang, Guoxing Sun, Zhuo Su, Kaiwen Guo, Minye Wu, Jingyi Yu, and Lan Xu. Neuralhofusion: Neural volumetric rendering under human-object interactions. In *CVPR*, 2022. 2
- [24] Vladimir G Kim, Siddhartha Chaudhuri, Leonidas Guibas, and Thomas Funkhouser. Shape2pose: Human-centric shape analysis. *ACM Trans. Graph.*, 33(4), 2014. 1, 2
- [25] Jehee Lee, Jinxiang Chai, Paul SA Reitsma, Jessica K Hodgins, and Nancy S Pollard. Interactive control of avatars animated with human motion data. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, 2002. 2
- [26] Jehee Lee and Kang Hoon Lee. Precomputing avatar behavior from human motion data. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, 2004. 3
- [27] Kyungho Lee, Seyoung Lee, and Jehee Lee. Interactive character animation by learning multi-objective control. *ACM Trans. Graph.*, 37(6), 2018. 2, 3
- [28] Kyungho Lee, Sehee Min, Sunmin Lee, and Jehee Lee. Learning time-critical responses for interactive character control. *ACM Trans. Graph.*, 40(4), 2021. 3
- [29] Kang Hoon Lee, Myung Geol Choi, and Jehee Lee. Motion patches: building blocks for virtual environments annotated with motion data. In *ACM SIGGRAPH 2006 Papers*. 2006. 2
- [30] Seunghwan Lee, Phil Sik Chang, and Jehee Lee. Deep compliant control. In *ACM SIGGRAPH 2022 Conference Proceedings*, 2022. 2, 3
- [31] Seyoung Lee, Sunmin Lee, Yongwoo Lee, and Jehee Lee. Learning a family of motor skills from a single motion clip. *ACM Trans. Graph.*, 40(4), 2021. 3
- [32] Seunghwan Lee, Moonseok Park, Kyoungmin Lee, and Jehee Lee. Scalable muscle-actuated human simulation and control. *ACM Trans. Graph.*, 38(4), 2019. 5
- [33] Sergey Levine, Jack M Wang, Alexis Haraux, Zoran Popović, and Vladlen Koltun. Continuous character control with low-dimensional embeddings. *ACM Trans. Graph.*, 31(4), 2012. 3
- [34] Ruilong Li, Shan Yang, David A. Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In *ICCV*, 2021. 2

- [35] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. Character controllers using motion vae. *ACM Trans. Graph.*, 39(4), 2020. 2, 3
- [36] Kovar Lucas, Gleicher Michael, and Pighin Frédéric. Motion graphs. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, 2002. 2
- [37] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, 2017. 2
- [38] Josh Merel, Saran Tunyasuvunakool, Arun Ahuja, Yuval Tassa, Leonard Hasenclever, Vu Pham, Tom Erez, Greg Wayne, and Nicolas Heess. Catch & carry: reusable neural controllers for vision-guided whole-body tasks. *ACM Trans. Graph.*, 39(4), 2020. 2
- [39] Soohwan Park, Hoseok Ryu, Seyoung Lee, Sunmin Lee, and Jehee Lee. Learning predict-and-simulate policies from un-organized human motion data. *ACM Trans. Graph.*, 38(6), 2019. 2, 3, 5
- [40] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4), 2018. 3, 5
- [41] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Trans. Graph.*, 41(4), 2022. 3
- [42] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.*, 40(4), 2021. 3
- [43] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *ICCV*, 2021. 2, 3
- [44] Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. Dexmv: Imitation learning for dexterous manipulation from human videos. In *ECCV*, 2022. 1, 2
- [45] Manolis Savva, Angel X Chang, Pat Hanrahan, Matthew Fisher, and Matthias Nießner. Pigraphs: learning interaction snapshots from observations. *ACM Trans. Graph.*, 35(4), 2016. 1, 2
- [46] Hubert PH Shum, Taku Komura, Masashi Shiraishi, and Shuntaro Yamazaki. Interaction patches for multi-character animation. *ACM Trans. Graph.*, 27(5), 2008. 2
- [47] Sebastian Starke, Ian Mason, and Taku Komura. Deepphase: periodic autoencoders for learning motion phase manifolds. *ACM Trans. Graph.*, 41(4), 2022. 2
- [48] Sebastian Starke, He Zhang, Taku Komura, and Jun Saito. Neural state machine for character-scene interactions. *ACM Trans. Graph.*, 38(6), 2019. 1, 2
- [49] Omid Taheri, Vasileios Choutas, Michael J Black, and Dimitrios Tzionas. Goal: Generating 4d whole-body motion for hand-object grasping. In *CVPR*, 2022. 1, 2
- [50] Omid Taheri, Nima Ghorbani, Michael J Black, and Dimitrios Tzionas. Grab: A dataset of whole-body human grasping of objects. In *ECCV*, 2020. 1, 2
- [51] Graham W Taylor and Geoffrey E Hinton. Factored conditional restricted boltzmann machines for modeling motion style. In *ICML*, 2009. 2
- [52] Adrien Treuille, Yongjoon Lee, and Zoran Popović. Near-optimal character animation with continuous control. In *ACM SIGGRAPH 2007 papers*. 2007. 3
- [53] Ruben Villegas, Jimei Yang, Yuliang Zou, Sungryull Sohn, Xunyu Lin, and Honglak Lee. Learning to generate long-term future via hierarchical prediction. In *ICML*, 2017. 2
- [54] Jingbo Wang, Yu Rong, Jingyuan Liu, Sijie Yan, Dahua Lin, and Bo Dai. Towards diverse and natural scene-aware 3d human motion synthesis. In *CVPR*, 2022. 1, 2
- [55] Jiashun Wang, Huazhe Xu, Jingwei Xu, Sifei Liu, and Xiaolong Wang. Synthesizing long-term 3d human motion and interaction in 3d scenes. In *CVPR*, 2021. 1, 2, 6, 7, 8, 9
- [56] Jingbo Wang, Sijie Yan, Bo Dai, and Dahua Lin. Scene-aware generative network for human motion synthesis. In *CVPR*, 2021. 2, 3
- [57] Zan Wang, Yixin Chen, Tengyu Liu, Yixin Zhu, Wei Liang, and Siyuan Huang. Humanise: Language-conditioned human motion generation in 3d scenes. *NeurIPS*, 2022. 2
- [58] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. A scalable approach to control diverse behaviors for physically simulated characters. *ACM Trans. Graph.*, 39(4), 2020. 3
- [59] Yan Wu, Jiahao Wang, Yan Zhang, Siwei Zhang, Otmar Hilliges, Fisher Yu, and Siyu Tang. Saga: Stochastic whole-body grasping with contact. In *ECCV*, 2022. 2
- [60] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *CVPR*, 2020. 6
- [61] Xianghui Xie, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Chore: Contact, human and object reconstruction from a single rgb image. In *ECCV*, 2022. 1, 2
- [62] Xiang Xu, Hanbyul Joo, Greg Mori, and Manolis Savva. D3d-hoi: Dynamic 3d human-object interactions from videos. *arXiv preprint arXiv:2108.08420*, 2021. 2
- [63] Zeshi Yang, Kangkang Yin, and Libin Liu. Learning to use chopsticks in diverse gripping styles. *ACM Trans. Graph.*, 41(4), 2022. 1, 2
- [64] He Zhang, Yuting Ye, Takaaki Shiratori, and Taku Komura. Manipnet: Neural manipulation synthesis with a hand-object spatial representation. *ACM Trans. Graph.*, 40(4), 2021. 1, 2
- [65] Jason Y. Zhang, Sam Pepose, Hanbyul Joo, Deva Ramanan, Jitendra Malik, and Angjoo Kanazawa. Perceiving 3d human-object spatial arrangements from a single image in the wild. In *ECCV*, 2020. 2
- [66] Siwei Zhang, Yan Zhang, Qianli Ma, Michael J Black, and Siyu Tang. Place: Proximity learning of articulation and contact in 3d environments. In *3DV*, 2020. 1, 2
- [67] Xiaohan Zhang, Bharat Lal Bhatnagar, Sebastian Starke, Vladimir Guzov, and Gerard Pons-Moll. Couch: Towards controllable human-chair interactions. In *ECCV*, 2022. 1, 2
- [68] Yan Zhang, Mohamed Hassan, Heiko Neumann, Michael J Black, and Siyu Tang. Generating 3d people in scenes without people. In *CVPR*, 2020. 1, 2

- [69] Kaifeng Zhao, Shaofei Wang, Yan Zhang, Thabo Beeler, , and Siyu Tang. Compositional human-scene interaction synthesis with semantic control. In *ECCV, 2022*. [1](#), [2](#)
- [70] 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](#)
- [71] Yi Zhou, Connelly Barnes, Lu Jingwan, Yang Jimei, and Li Hao. On the continuity of rotation representations in neural networks. In *CVPR, 2019*. [4](#), [6](#)